SAFEGUARDING FOOD SECURITY IN VOLATILE GLOBAL MARKETS

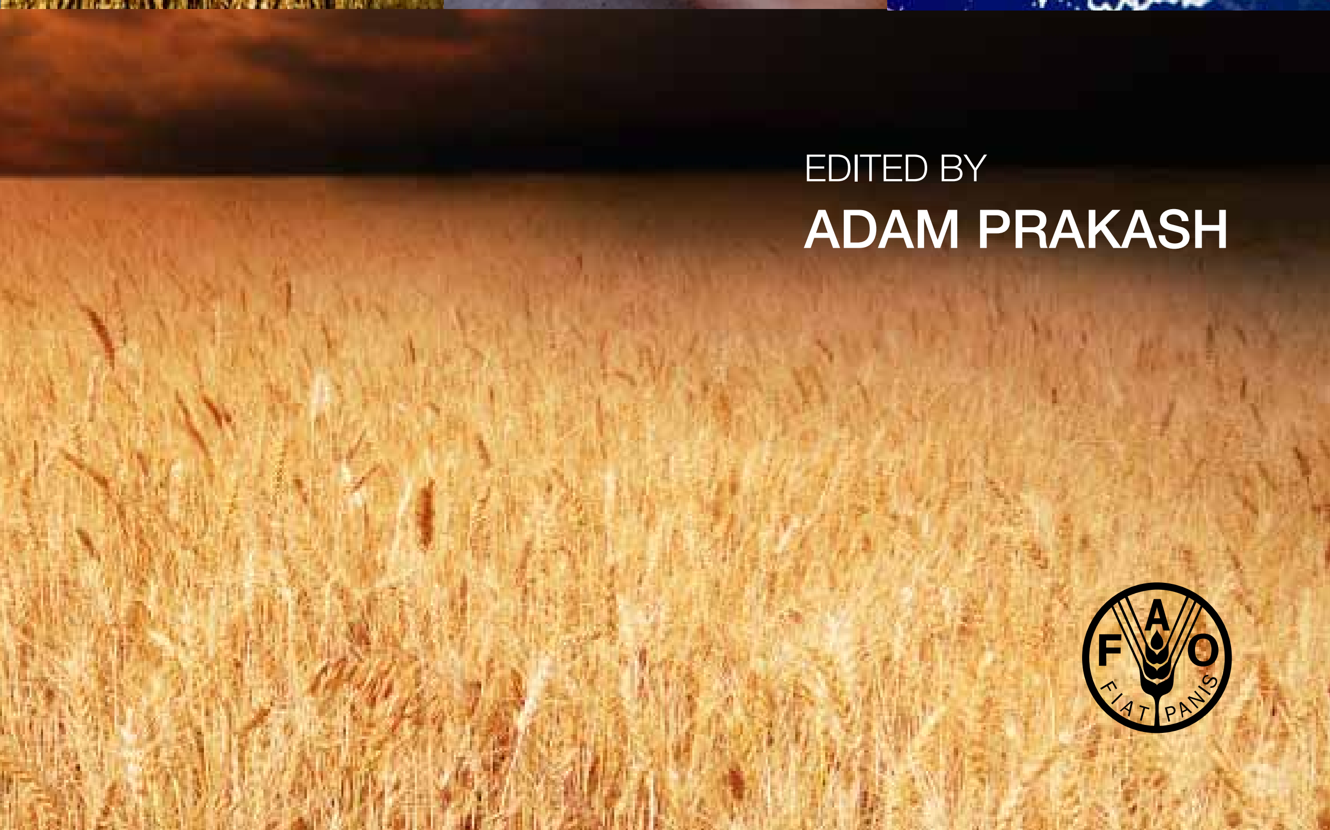

EDITED BY
ADAM PRAKASH

ISBN 978-92-5-106803-8

Cover design by Giancarlo de Pol
Cover photo (centre) ©FAO/Ami Vitale
Typesetting and layout by River Valley Technologies (river-valley.com)

Safeguarding food security in volatile global markets

Edited by Adam Prakash

Food and Agriculture Organization of the United Nations, Rome, 2011

Contents

Chapter 16

The role of low stocks in generating volatility and panic

Matthieu Stigler and Adam Prakash[1]

The purpose of this study is to investigate non-linearities in the relationship between agricultural commodity prices, using wheat futures from the Chicago Mercantile Exchange (CME) and wheat inventories of the world's leading wheat exporter - the United States of America - relative to disappearance.[2] The underlying notion of our research is that in modelling volatility, one should take into account the presence of different volatility regimes, and to determine in what manner and to what extent do expectations about future inventories (as an indicator of prospective market "tightness") influence those regimes.

For this purpose, we employ a volatility regime-switching model. Results show that in the absence of market tightness, commodity prices do not appear to be influenced by inventories. However, when inventories fall low, commodity prices become highly linked to information on stocks, and especially to supply and demand disturbances that reduce the stocks-to-disappearance ratio further. Conversely, low volatility regimes emerge when stocks are in abundance.

Introduction

As seen in earlier chapters, among potential sources of market volatility, stocks have been instrumental in moderating or amplifying volatility, and are viewed by policy-makers as key to buffering market turbulence. Indeed, stocks have played an important role in price stabilization policies in the past, and remain topical today in discussions about achieving food security. For example, the announcement of the release of Japanese rice security stocks is thought to have acted as a depressant during the rice spike (Dawe, 2010). Therefore, it is not surprising that the usefulness of holding public stocks has been the subject of debate by scholars in recent years - see Timmer (2010) and von Braun & Torero (2009). This chapter focuses on the impact of new information on stocks and its subsequent impact on price dynamics. The analysis contributes to understanding how periods of extreme volatility can arise in commodity markets.

In discussions of the role of stocks in generating market turmoil, a comprehensive understanding of the relationship between inventories and price dynamics is required. The

1 Statistics Division (FAO). We want to give special thanks to David Ardia for providing guidance on this chapter. We also acknowledge Jean-Yves Pittarakis for his comments on an early draft.

2 "Disappearance" is measured by domestic utilization plus exports.

econometric approach of time series decomposition in Chapter 5 supported the significance of this link, but several economists remain less than convinced about the empirical importance of the relationship, both in the short- and long-term. For instance, Dawe (2009) concludes that the link between volatility and world rice stock levels was rather weak during the 2006-08 event. Roache (2010), who used data from the last one-hundred years, came to a similar finding that long-term volatility in commodity prices is not influenced by commodity stock levels.

More generally, there is a significant body of theoretical literature, centred on the *competitive storage model*, which views inventories as the main determinant in commodity price behaviour. While this book devotes due attention to the storage model, with a comprehensive description in Chapter 15 and a discussion of its predictive power for the dynamic properties of commodity prices in Chapter 2, our task here is to investigate the storage model's implications for price volatility.

The storage model, introduced in the pioneering work of Gustafson (1958) and further developed by Samuelson (1971), Scheinkman & Schechtman (1983), Wright & Williams (1982, 1984) and Deaton & Laroque (1992), studies whether or not speculators[3] will store a commodity depending on its expected price at the next period. A key issue is recognizing that storage cannot be negative, i.e. one can subtract a commodity from the present to deliver it in the future, but one cannot borrow a commodity produced in the future and deliver it in the present.

This constraint introduces a non-linearity, where price behaviour radically changes between periods where stocks are held and periods when they are not. Periods of positive stocks appear when the actual price is below its future expected value. In this regime, speculators store the commodity; by doing so, they introduce auto-correlation in the price although the supply is assumed to be independent and identically distributed (i.i.d.). But when the price is unusually high, and hence expected to be lower in the next period, incentives to store vanish, leading to a "stock-out" during which prices simply follow the assumed i.i.d. process.

Figure 16.1 illustrates this phenomenon in the Chicago Mercantile Exchange (CME) wheat market, data for which are used in this study. The non-linear behaviour of the (futures) prices tending to be exceptionally high with forecasts of low stocks (relative to expected disappearance) can be clearly seen.

Volatility dynamics follow a similar scheme, as they are seen to differ between regular and stock-out regimes. In the regular regime, volatility is found to increase with the price level. Conversely, in the stock-out regime dominated by i.i.d supply, the conditional variance of prices is constant regardless of how high prices are.

The storage model's prediction that volatility is constant in stock-out periods may at first seem surprising, given that one would expect that in a regime in which stocks are absent, any small supply or demand disturbances would lead to exceptional upward price movements. However, it is important to stress that the competitive storage model's basis is in explaining inter-crop year fluctuations only, as it models a world in which production occurs at each and every period. Thus, the prediction of constant volatility stemming from the competitive storage model concerns the comparison between years, and not within years.

In addition, using an annualized data set leads to a significant reduction in the number of observations, which potentially discards key information revealed by more frequent data that

[3] In this chapter, the term speculator refers to those agents who physically store the commodity, including commercial agents.

Figure 16.1: Stocks-to-disappearance forecasts and futures prices (US wheat)

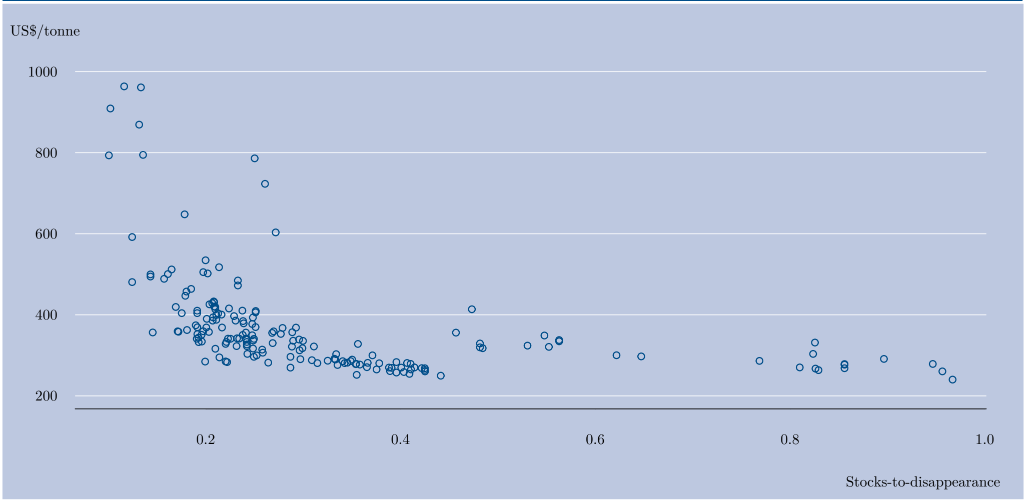

are often at one's disposal. A second problem concerns the measurement of stocks. Because stocks are held by speculators, their levels are not publicly known, and available annual estimates are only imperfect approximations. Moreover, the imprecision of approximated stock levels is compounded by the fact that they are often derived as a residual that balances the supply and demand identity.

As a consequence, we move our analysis in a slightly different direction from the competitive storage model. Rather than asking whether stocks affect (annual) wheat price volatility, we examine whether the reaction of traders to official announcements on expected wheat stocks-to-disappearance affect price volatility, and more importantly, whether responses to announcements are conditioned by market sentiment. Simply put, we focus on how the market reacts to stock forecasts. There are several advantages to our approach. First, it addresses and circumvents the issue of using annual stock variables in balance sheets, as analysts typically do. Annual measures of stocks are misleading, because being *ex post*, they are unknown by market participants at the time, and thus have no direct influence on agents' current behaviour.[4] By contrast, through using stock forecasts instead of *ex post* stock numbers, we are able to claim that these variables were used in trading decisions and therefore pertinent to price determination over the sample, including volatility outcomes. In taking this methodological approach, we follow the advice of Schwager (1984), who notes that:

> *It is frequently possible to build a more accurate model using past estimates rather than actual statistics as the price explanatory variables. For example, if we are trying to construct a model to explain and predict October-December prices for a given commodity, we might find that past supply estimates released during the October-December period are more helpful than the actual supply statistics in explaining historical price variation. Such price behaviour would merely reflect that what the market thought was true in the past was more important in determining prices than what was actually true (as defined by the final revised estimates).* (Schwager, 1984, p. 58)

[4] This is certainly true for end-of-year (season) stocks, while for the previous year's (season's) stocks, the information content and its relevance would be conceivably "stale".

We employ stocks-to-disappearance forecasts published by the official institution of the world's leading grain exporter: the World Supply and Demand Estimates (WASDE) published by the United States Department of Agriculture (USDA). A further advantage of using the WASDE forecasts is that higher frequency - monthly- data may be employed, to arrive at more meaningful hypotheses. But, even though using data of a higher than annual frequency is more appropriate to investigate volatility, by using daily price data we still end up with a frequency discrepancy in our sample. The importance of using daily price data is that it embodies precise information on market behaviour at the time of a specific WASDE announcement. Therefore, as our primary focus is to assess the market reaction to new information on stocks, we retain daily price data, and choose to treat monthly stock expectation data as a latent variable.

Moving forward to the empirical analysis, the following section presents our modelling framework and methodological issue. After which, we explain how the data series employed in the analysis are constructed, and then we report on our results. Finally, some concluding thoughts are presented.

How do stock forecasts affect the market volatility?

Econometric practitioners usually model the conditional volatility of returns by employing the Autoregressive Conditional Heteroskedasticity (ARCH) model initially introduced by Engle (1982), and its variants. In its basic form, the ARCH model specifies the conditional variance as dependent on the previous squared innovation of the series. Bollerslev (1986) generalized the ARCH model by allowing the conditional variance to depend also on previous conditional variance levels. This modification, which leads to a more parsimonious parameterization of the model[5], is known as GARCH (1,1):

$$\begin{aligned} y_t &= f(y_{t-1}) + u_t \\ u_t &= \epsilon_t \sigma_t \\ \epsilon_t &\sim \mathcal{D}(0,1) \\ \sigma_t^2 &= \omega + \alpha u_{t-1}^2 + \beta \sigma_{t-1}^2, \end{aligned} \tag{1}$$

where, to ensure a positive variance, one constrains $\omega \geq 0$ and$\alpha, \beta > 0$. Stationarity of the variance process is imposed through: $\alpha + \beta < 1$. $\mathcal{D}(0,1)$ is an arbitrary independent and identical distribution[6] with mean and variance equal to 0 and 1, respectively. f is a simple "filtering" function to remove possible auto-correlation of the log-returns, and y_t is the log-return series.

The GARCH framework has spurred significant interest in theoretical and empirical research, and this model class (along with its extensions) is now used widely to forecast commodity price volatility. If stock data were available at the same frequency as prices, it would be straightforward to investigate the effect of stocks on volatility by adding a further term in equation (1). This is not a feasible option though, as only monthly stock data are available. While a simple solution would be to insert a dummy variable taking value 1 when stock forecasts are available, and 0 if not, this approach lacks a suitable theoretical underpinning as it does not acknowledge the possibility of regime-switching behaviour as predicted by the storage model.

5 It can be shown that a GARCH(1,1) is equivalent to an ARCH(∞).

6 Typical choices for the distribution include the normal or the student, the latter is employed in the analysis.

The approach we therefore adopt is more theoretically consistent: because we acknowledge that the storage model predicts the presence of two regimes, we employ a regime-switching GARCH model when estimating volatility. That is, we investigate how volatility differs in periods of low stocks compared with periods of high stocks. This could be accomplished by using a threshold GARCH model, where the stock series variable determines the switching between two GARCH regimes. But because stock forecasts are observed only at a monthly frequency, we instead treat them as a latent variable. This requires us to model the switching process between regimes through a latent process that is assumed to follow a first-order Markov process. From hereon, we investigate the relationships between the stock series and the estimated switching process. Importantly, this enables us to conclude whether the model supports the hypothesis that changes in stocks lead to regime-switching.

We recognize a drawback of the approach in that it models the switch in regimes as the outcome of an unobserved a latent process, and so results are less precise about the determinants of the switching process. Indeed, it is not possible to infer with confidence whether stock changes lead to regime-switching, or whether this phenomenon is attributable to other variables.

The Markov-switching GARCH Model The Markov-switching model, introduced in econometrics by Hamilton (1989), describes a regime-switching regression where the transition between regimes is driven by a latent discrete Markov chain. In other words, the model parameters switch from one regime to another according to an unobservable process that is assumed to follow a first-order Markov process. Formally, the probability of the *state variable* S_t being in regime $i \in \{1,\dots,N\}$ only depends on its previous state:

$$P(S_t = i|S_{t-1} = j, S_{t-2}, \dots, S_0, \Omega_{t-1}) = P(S_t = i|S_{t-1} = j) \equiv p_{ji}, \tag{2}$$

where Ω_t is an information set at time t containing variables other than S_t. Typically, one considers only a restricted number of regimes N, in the present case only two.

Markov-switching models have been used in linear regression and autoregressive frameworks (Hamilton, 1990, 1991). An extension of the Markov-switching AR to the GARCH framework was provided by Hamilton & Susmel (1994). In this context, coefficients of Equation (1) can change depending on the state of the latent variable S_t. While appealing, the approach poses a considerable challenge for estimation. Hamilton & Susmel focused on a Markov-switching ARCH, because introducing a GARCH component creates a complicated path of dependence where the variance at time t depends on the entire history of the process. Haas et al. (2004), however, circumvented this problem by using a different specification, where the switching is assumed to occur between several conditional volatility processes. Because we are now dealing with two different volatility processes, the path-dependence problem is avoided allowing us to include GARCH components. This leads to the following equation:

$$\sigma^2_{t,i} = \omega_i + \alpha_i u^2_{t-1} + \beta_i \sigma^2_{t-1,i}. \tag{3}$$

Several insights can be gained by employing the Markov-switching GARCH (hereafter MS-GARCH) model. Firstly, because its specification allows periods of high and low unconditional volatility to be clearly identified, we can accurately estimate both regime probabilities as well as of the average duration of each regime. Secondly, regime probability estimates can be further used for comparison with the stocks variable, which allows an investigation on whether the stocks is attributed as the cause of switching behaviour.

A third advantage is that the MS-GARCH model is robust against changes in market conditions. Indeed, as Lamoureux & Lastrapes (1990) have shown, the single-regime GARCH model (1) tends to significantly overestimate volatility persistence in the presence of structural changes. This is precisely what motivated Hamilton & Susmel (1994) to develop the MS-ARCH model. The model avoids the bias of GARCH effects by allowing the coefficient on the unconditional volatility level to switch between regimes. When applied to a history of equity returns in the United States of America, during which equity markets underwent a significant crash in October 1987, the MS-ARCH model led to a significantly lower volatility persistence. Furthermore, the model was able to provide more accurate forecasts than those provided by a range of single-regime GARCH models.

MS-GARCH: Extension to GJR Another issue raised by the analysis is the question of asymmetry in volatility responses. As Chapter 2 illustrates, volatility tends to respond differently to positive shocks than to negative shocks. Interestingly, while in traditional financial asset markets negative shocks impact to a greater extent volatility (referred to as the leverage effect), in agricultural markets it is positive shocks that tend to drive greater volatility. This phenomenon, discussed more thoroughly in Chapters 1 and 2, is explained by the storage model, which posits that positive shocks will assume a contraction in stocks, in turn generating an increase in volatility.

One of the simplest ways to measure such asymmetric effects is to use the GARCH-GJR model (named after Glosten et al., 1993). These authors estimated positive and negative shocks in a separate manner through the following equation:

$$\sigma_t^2 = \omega + \alpha^+ u_{t-1}^2 \mathrm{I}(u_{t-1} > 0) + \alpha^- u_{t-1}^2 \mathrm{I}(u_{t-1} \leq 0) + \beta\sigma_{t-1,}^2 \tag{4}$$

where α^+ and α^- are positive and 'I' is an indicator function.

Asymmetry is tested for by comparing the α^+ and α^- coefficients. Integrating the so-called "GJR effect" into each regime of the MS-GARCH model is amenable. Doing so provides insights about whether the asymmetric effect is different in periods of either low and high volatility. Intuitively, the effect should be less pronounced in the low volatility regime, where a price shock does not ostensibly affect the level of inventories.

Modelling with MS-GARCH While Equation (3) of the MS-GARCH specifies that all coefficients can change between regimes, this need not be the case. It is indeed possible to use a simpler model where, for example, only the constant in the GARCH equation switches between regimes. Using this model has several advantages. Firstly, it allows for easier comparison and interpretation of regime dynamics, as there is only one parameter switching. Moreover, it reduces the computational burden encountered by the fully flexible specification of (3).

The model is estimated by direct maximization of the log-likelihood function, which is obtained by using the BEKK filter (Krolzig, 1997). Arriving at an estimation procedure is nevertheless a challenging task as both the parameters and the regime probabilities need to be estimated simultaneously. A further complication arises when the desire is to test for the presence of Markov-switching effects. Indeed, one is then confronted with the so-called "problem of non-identified parameters" under the null hypothesis (Andrews & Ploberger, 1994), as well as zero scores (Garcia, 1998).

While Carrasco & Hu (2004) and Hu & Shin (2008) have proposed several solutions that allow testing for Markov-switching, it is unclear whether their solutions can be applied to the framework of Haas et al. (2004) we adopt here. Therefore, we resort to a methodology that uses a more straightforward information-criterion based comparison, which has been shown to produce robust results in many different settings (see among others: Gonzalo & Pitarakis, 2002 in the similar case of threshold regime-switching models, and Aznar & Salvador, 2002 for determining the cointegration rank). In the same vein, we also employ an information-criterion procedure instead of standard statistical tests to compare MS-GARCH models with different parameterizations (i.e. the presence of GJR effects and whether all coefficients switch or not).

Finally, a few words ought to be said regarding our distributional assumptions. An interesting feature of GARCH models is that even if a symmetric normal distribution is assumed for $\mathcal{D}$ in (1), the unconditional distribution can exhibit excess kurtosis. Nevertheless, the normal distribution inadequately describes the fat-tails of the error distribution that is typically observed in financial variables. A simple solution therefore is to use the Student distribution instead (see Bollerslev, 1987), which is better suited for fat-tailed distributions. The distribution's degrees of freedom are estimated from the data, and could even be assumed to switch between regimes, as in Dueker (1997). However, this flexibility is not without cost, namely in the difficulty in interpreting and comparing results. Having described our modelling approach, we now turn to a discussion of the data employed and the results from estimation.

Data and estimation

Data

As previously mentioned, our inventory data constitute end-of-season forecasts for both the current and following year published in the USDA's monthly WASDE report. Our analysis uses the stock-to-disappearance ratio. Thus by "stocks" we refer formally to this ratio.

Figure 16.2 shows stocks-to-disappearance forecasts for wheat over the period 1970-2010. Because there can be two forecasts per month (actual and following year), years are shown in different colours – dark and light blue – for the sake of clarity. The final realized value (corresponding to the end-of-year stocks-to-disappearance) is indicated by a black circle. Interestingly, WASDE forecast accuracy does not seem to have been affected by volatility in business cycles, especially economic crises: typically, forecast errors do not seem to differ from the previous years.

Figure 16.3 shows USDA forecasts for the current year against the futures price for wheat reported on the day of the forecast release. Visually, there appears a highly negative relationship between the two series.

When constructing the stock series, we always used the subsequent end-of-season forecast: published from January to August (since September is considered to mark the end of the season) for the current year, and then September-December reports for the following year. Admittedly, this implies a certain heterogeneity in the forecast horizons: while the August report forecasts prices for the next month, the November or December report forecasts prices for almost one year ahead. This heterogeneity means that forecasts are likely not to have the same impact: the market probably reacts more to forecasts close to the end season than to those for longer horizons.

Regarding the price series, as shown in Figure 16.3, we use daily futures data from CME from January 1985 to January 2009 (6 000 observations in total). Here, construction of the series is more standard, although one should, when constructing such price series, keep in

Figure 16.2: WASDE wheat stocks-to-disappearance forecasts for US: 1970-2010

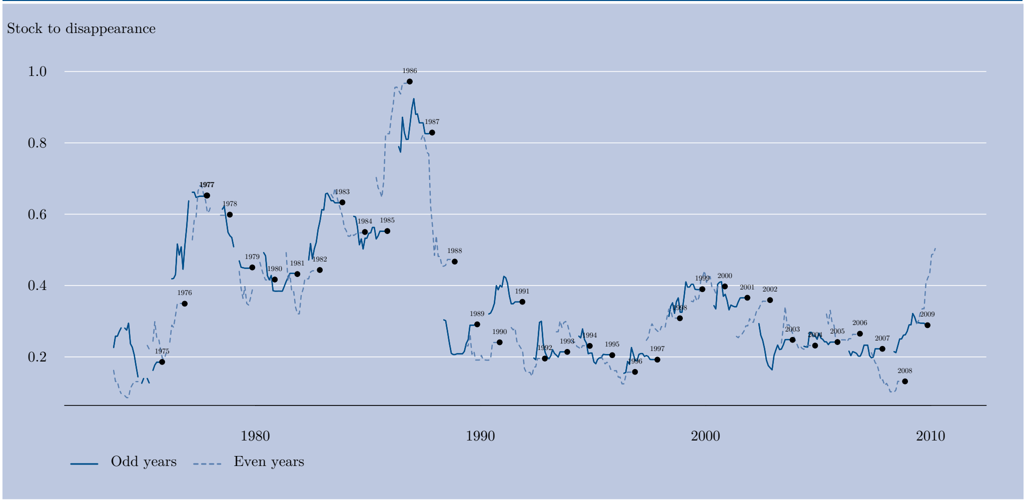

Figure 16.3: WASDE wheat stocks-to-disappearance forecasts for US and CME wheat futures prices

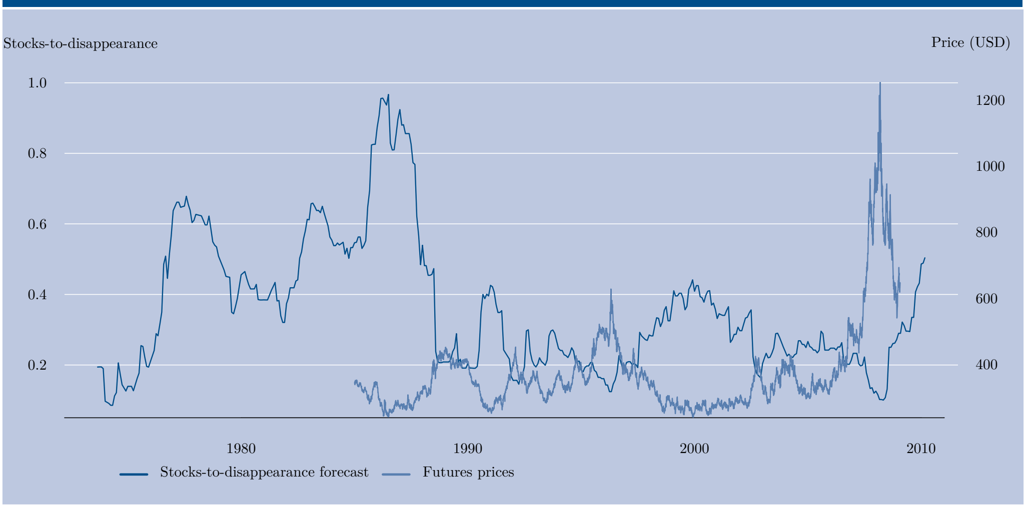

mind the so-called Samuelson (or maturity) effect. The maturity effect states that volatility of the futures price tends to increase when the maturity date approaches, see Chapter 2 for a more detailed discussion. To avoid the "artificial volatility" introduced by the maturity effect, we created a synthetic futures series with a constant maturity of 100 days rather than the nearby maturity.[7]

[7] This is done by spline-interpolating the 100 day-maturity price based on available maturities for each observation.

Table 16.1: GARCH model results

	Estimate	Std. Error	t value	Pr (>|t|)
ω	0.02	0.01	3.73	0.00
α	0.05	0.01	8.27	0.00
β	0.94	0.01	123.46	0.00
ν	12.33	1.58	7.81	0.00
AIC:	21805.69			
NLL	10648.85			
Persistence	0.991			
Uncond. variance	0.02			

Results

As conventional unit root tests have indicated our price series is non-stationary, we investigate the volatility of the log-returns, as is commonly employed in the financial literature. We first estimate as a benchmark, a simple GARCH model with Student errors:

$$\sigma_t^2 = \omega + \alpha u_{t-1}^2 + \beta \sigma_{t-1}^2 \tag{5}$$

Results are shown in Table 16.1. All parameters were found to be statistically significant. The persistence of the estimated variance (given by $\alpha+\beta$) is close to 1 at 0.991, as is typically found for high frequency series. The corresponding unconditional variance is then 0.02.

When GJR parameters are introduced, we find that the coefficient α^+ for positive shocks is higher than α^- for negative shocks (0.06 compared with 0.03), confirming the results of Carpantier (2010). Testing for the inequality of $\alpha^+ > \alpha^-$ is complicated, as inequality tests that involve two or more coefficients have non-standard distributions. Hence, we resort to a simple comparison of the 99% confidence intervals of both parameters. These intervals do not overlap, indicating α^+ is statistically higher than α^-.

Introducing now Markov-switching effects in the GARCH model, we turn to the simpler specification where only ω_i in (3) can switch, while the other parameters remain constant:

$$\sigma_{t,i}^2 = \omega_i + \alpha u_{t-1}^2 + \beta \sigma_{t-1,i}^2 \tag{6}$$

In this model, the persistence (given by $\alpha+\beta$) is the same in each regime, but the unconditional volatility $\omega_i/(1-\alpha-\beta)$, can differ depending on the regime.

Results of the MS-GARCH with ω switching are shown in Table 16.2. Interestingly, we see that the unconditional variance is much higher (0.19) in the second regime compared with the first (0.009). Turning to the transition probabilities, we obtain the following matrix:

$$\begin{pmatrix} 0.991 & 0.009 \\ 0.133 & 0.867 \end{pmatrix}$$

While the first regime appears to dominate (with a probability of 99 percent of dwelling in this regime), the second regime is also persistent with a probability of remaining within, as high as 86 percent.

Table 16.2: MS-GARCH with ω switching				
	Estimate	Std. Error	t value	Pr (>|t|)
ω_1	0.0094575	0.0029223	3.2363	0.0012108
ω_2	0.1938141	0.0535027	3.6225	0.0002918
α	0.0278933	0.0044838	6.2210	4.941e-10
β	0.9634873	0.0052218	184.5139	<2.2e-16
p_{11}	0.9913151	0.0035032	282.9707	<2.2e-16
p_{22}	0.8668888	0.0428881	20.2128	<2.2e-16
ν	28.9465468	13.4929700	2.1453	0.0319284
NLL	10636.31			
AIC	21286.62			
	Regime1		Regime 2	
Persistence	0.9913		0.9913	
Uncond. variance	0.0095		0.1955	

An interesting extension is to consider the inclusion of GJR effects and further allow all GARCH parameters to switch, which leads to:

$$\sigma^2_{t,i} = \omega_i + \alpha^+_i u^2_{t-1} \mathrm{I}(u_{t-1} > 0) + \alpha^-_i u^2_{t-1} \mathrm{I}(u_{t-1} \leq 0) + \beta_i \sigma^2_{t-1,i}. \tag{7}$$

Results are shown in Table 16.3. It is of interest to compare the different GJR dynamics between low and high regimes. The asymmetry is still found to be present, but in this instance with a much stronger impact in the high regime. Indeed, while in the low regime the impact is 0.014 (α^+_1), it switches to 0.116 (α^+_2), which represents a highly important difference. Somewhat surprising is that in the high volatility regime, negative (price-decreasing) shocks do not have any influence on volatility whatsoever: α^-_2 - the coefficient on negative shocks - is not significantly different from zero.

Figure 16.4 shows the regime-dependent news-impact curve. Two facts surface from this figure. First, the unconditional volatility level is seen to be very different in the low-volatility regime (light blue line, right axis) than in the high-volatility regime (dark blue line, left axis). Secondly, the asymmetric effect is much stronger in the high volatility regime: when volatility is already high, “bad news” (positive shocks) have a dramatic impact on volatility, where they will increase conditional volatility by 0.11 (compare to 0.02 in the low-volatility regime). Put simply, "bad news" has more than 4 times a greater impact than "good news".

Turning now to the transition probabilities, the second regime appears to be much more persistent than in the previous model, where the probability to remain in the second regime is now close to 98 percent, similar to the probability of staying in the first regime. This can also be seen in Figure 16.5, which shows the smoothed probability of the high volatility state over time, together with the original price series.

Finally, we compare the three models used based on the AIC criterion. The AIC criterion favours the last model –the MS-GARCH with all coefficients allowed to switch– to the simpler MS-GARCH and the single-regime GARCH. This suggests that indeed a Markov-switching

Table 16.3: MS-GARCH with all coefficients switching

	Estimate	Std. Error	t value	Pr (>|t|)
ω_1	1.8491e-03	8.0236e-04	2.3046	0.021188
ω_2	6.6999e-02	2.5600e-02	2.6171	0.008868
α_1^+	1.3706e-02	4.2598e-03	3.2175	0.001293
α_1^-	6.4696e-03	4.9853e-03	1.2977	0.194380
α_2^+	1.15953-01	2.5424e-02	4.5605	5.102e-06
α_2^-	2.4797e-02	1.6529e-02	1.5003	0.133549
β_1	9.8650e-01	3.3663e-03	293.0537	<2.2e-16
β_2	9.1111e-01	2.3490e-02	38.7874	<2.2e-16
p_{11}	9.8517e-01	2.3490e-02	38.7874	<2.2e-16
p_{22}	9.8239e-01	8.2463e-03	119.1310	<2.2e-16
ν	1.5357e+01	2.6222e+00	5.8565	4.728e-09

NLL	10618.21			
AIC	21258.42			
	Regime1		Regime 2	
Persistence	0.9966		0.9815	
Uncond. variance	0.0019		0.0683	

Figure 16.4: Regime-dependent news impact curve

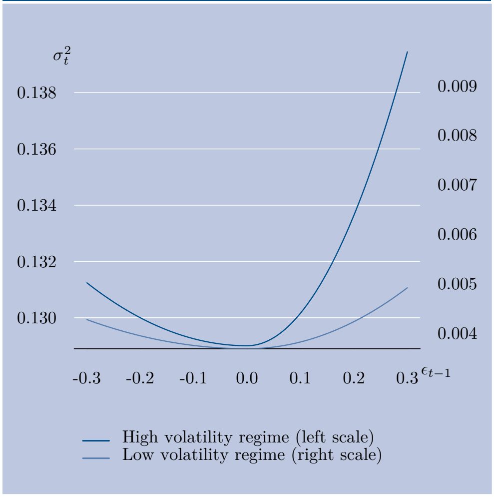

Figure 16.5: High-volatility regime probabilities

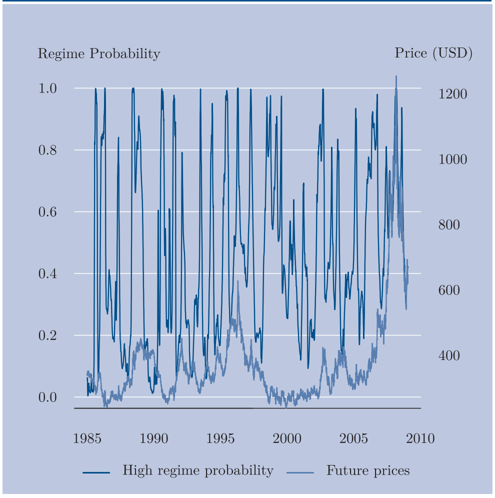

model is best suited to capture the volatility, in accordance with the prediction from the storage model.

In conclusion, using Markov-switching GARCH models, we observe two regimes characterized by significant differential volatility levels. Furthermore, once we take

Figure 16.6: Probability of remaining in high-volatility regime versus wheat stock disappearance

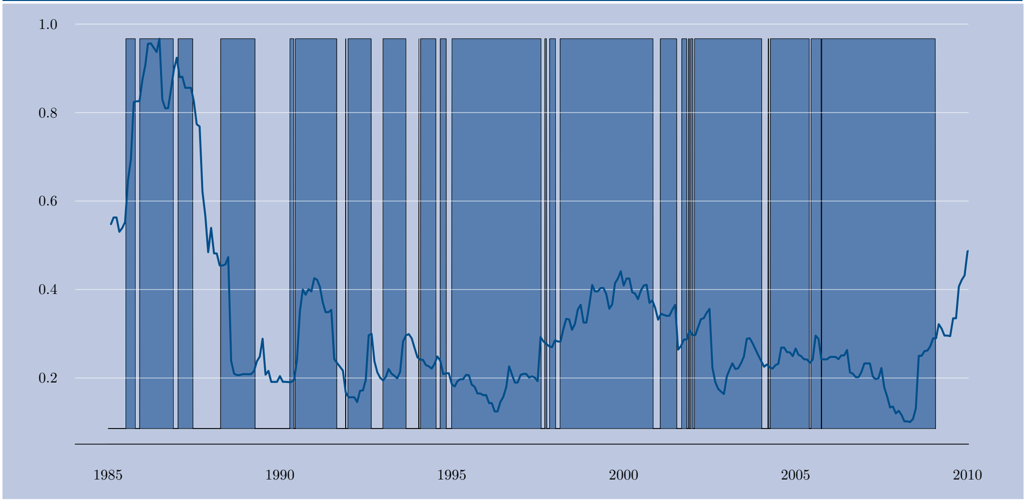

Note: We use here the value of 0.8 as the threshold above, which the smoothed probablity is assumed to be one, and conversely, zero below.

asymmetric effects into account, we observe that in the extreme volatility regime, bad news have a much more dramatic impact than under the quiescent regime. This suggests that the regimes reflect “market sentiment”: in periods of high volatility, even small surprises can exacerbate market tension thereby fuelling panic.

Now that we have identified different regimes of volatility, a question remains on what determines the "switch" between regimes. Within the Markov framework, this is assumed to be triggered by an unobserved latent variable. To address this, the next question is whether information on stocks-to-disappearance can be associated with the regime switches.

Comparisons of switching regimes with the stocks variable

Given that we observed Markov-switching between regimes of significant differing volatility levels in the futures price of wheat, we now seek to understand whether information on United States wheat stocks can generate the observed switching. To do so, we employ an informal approach, where a graphical comparison between stock forecasts with regime transition probabilities is first made, and then from which a simple probit model is applied. Figure 16.6 shows the graphical comparison of stock forecasts with the probability of being in the high volatility regime (in blue).

Though it is difficult to establish a direct link between the two variables, one can nevertheless observe that periods of low stocks are only present in the high regime. What is more surprising is to see that there are also periods of high stocks in the high volatility regime such as in 2000, but for the most part, periods of high stocks and high volatility have been characterized by a strong decrease, for instance in 2003 and 2004.

We now turn to the probit model, where we use only the regime probabilities in the day for which the forecast was published. It is observed in Table 16.4 that there is a strong

Table 16.4: Probit model results

	Estimate	Std. Error	z value	Pr (>|z|)
Intercept	1.12	0.16	6.95	0.00
Stocks	-0.92	0.40	-2.28	0.02

Table 16.5: Average effects model results

	Intercept	Stocks
Average effect	0.32	-0.26
At average values	0.33	-0.27
At Q3 values	0.35	-0.29
At Q1 values	0.29	-0.24

and significant negative coefficient on the stock forecast variable. Because the coefficient has no direct interpretation in the probit framework, we assess the average effects, as shown in Table 16.5.

The negative effect is confirmed, and seems robust to measures used, while varying only little among quantiles. Thus, the probit model suggests that stocks indeed have an impact on the regime-switching process: lower stocks-to-disappearance increases the probability of being in the high volatility regime.

In summary, our hypothesis that stock forecasts influence the observed switching is confirmed, albeit under this rather informal approach, in that downward forecast revisions augment the probability of being in the high volatility regime.

Conclusions

This chapter investigated the impact of USDA's forecasts of end-of-season stocks-to-disappearance on volatility in the CME wheat futures market, which is a reference point for price discovery in the global wheat market. Our enquiry was conducted in two steps. Firstly, we modelled volatility in a Markov regime-switching GARCH framework. Within the two regimes identified, we observed a significant difference in unconditional volatility levels, where volatility in the high regime ranges between 20 and 36 times greater than in the low volatility regime. Secondly, we estimated how each regime reacts to the arrival of positive and negative news on stocks-to-disappearance, drawing upon the important body of literature investigating the so-called *leverage effect*. Our results show indeed a strong asymmetry between the low and high volatility regimes. In the high volatility regime, "bad news" (i.e. a positive shock) will result in an increase in volatility, which confirms the result that volatility is 36 times stronger in order than under the low volatility regime. We propose to interpret this as "panic", where in periods of high volatility, bad news will have a much more dramatic impact on the market than otherwise.

In the second step, we enquired whether switching between regimes could be a result of changes in the USDA's stocks-to-disappearance forecasts. This was done through resorting to a graphical comparison, followed by applying a simple probit model. Based on our investigations, we observed that stock forecasts are likely to generate regime-switching: when forecasts of stocks depletion are announced by the USDA, the probability of being in the extreme volatility regime increases significantly.

The approach we adopted appears promising in shedding light on the behaviour of agricultural commodity prices and warrants deeper and more extensive enquiry. One line of research, for instance, would be to apply the analysis to other storable commodities, including those in non-agricultural markets. It would also be useful to apply the approach to asset prices, to identify if the phenomenon of volatility regimes are observable in financial markets. Another interesting avenue to explore would be to employ other proxies for inventory tightness, such as the spread between futures and spot prices as in Ng & Pirrong (1994).

As for policy, the results reveal the importance of expected stocks held by major grain exporting countries in determining episodes of elevated price volatility in food markets. It might be tempting to infer that the corollary of this conclusion would be to increase inventories *per se* to prevent turmoil. While this may be true to diffuse the prospect of isolated turbulence in domestic markets, this chapter demonstrates that ample and highly liquid commercial stocks held by major international suppliers appear a necessary and sufficient condition to instil confidence in world markets and to lessen the probability of future bouts of extreme global volatility and crises from occurring.

References

Andrews, D. W. K. & Ploberger, W. 1994. Optimal tests when a nuisance parameter is present only under the alternative, *Econometrica*, 62(6): 1383–1414.

Aznar, A. & Salvador, M. 2002. Selecting the rank of the cointegration space and the form of the intercept using an information criterion, *Econometric Theory*, 18(4): 926–947.

Bollerslev, T. 1986. Generalized autoregressive conditional heteroskedasticity, *Journal of Econometrics*, 31(3): 307–327.

——— 1987. A conditionally heteroskedastic time series model for speculative prices and rates of return, *The Review of Economics and Statistics*, 69(3): 542–47.

Carpantier, J.-F. 2010. *Commodities inventory effect,* Center for Operations Research and Econometrics.

Carrasco, M. & Hu, L. 2004. Optimal test for Markov switching, *2004 Meeting Papers 374*, Society for Economic Dynamics.

Dawe, D. 2009. *The unimportance of "low" world grain stocks for recent world price increases,* FAO, Rome.

——— (ed.) 2010. *The Rice Crisis*. London, Earthscan.

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

Dueker, M. J. 1997. Markov switching in GARCH processes and mean-reverting stock-market volatility, *Journal of Business & Economic Statistics*, 15(1): pp. 26–34.

Engle, R. F. 1982. Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation, *Econometrica*, 50(4): 987–1007.

Garcia, R. 1998. Asymptotic null distribution of the likelihood ratio test in Markov switching models, *International Economic Review*, 39(3): 763–788.

Glosten, L. R., Jagannathan, R. & Runkle, D. E. 1993. On the relation between the expected value and the volatility of the nominal excess return on stocks, *Journal of Finance*, 48(5): 1779–1801.

Gonzalo, J. & Pitarakis, J.-Y. 2002. Estimation and model selection based inference in single and multiple threshold models, *Journal of Econometrics*, 110(2): 319 – 352.

Gustafson, R. 1958. *Carryover levels for grains,* Washington DC: USDA, Technical bulletin 1178.

Haas, M., Mittnik, S. & Paolella, M. 2004. A new approach to Markov-switching GARCH models, *Journal of Financial Econometrics*, 2(4): 493–530.

Hamilton, J. D. 1989. A new approach to the economic analysis of nonstationary time series and the business cycle, *Econometrica*, 57(2): pp. 357–384.

——— 1990. Analysis of time series subject to changes in regime, *Journal of Econometrics*, 45(1-2): 39–70.

——— 1991. A quasi-Bayesian approach to estimating parameters for mixtures of normal distributions, *Journal of Business & Economic Statistics*, 9(1): pp. 27–39.

Hamilton, J. D. & Susmel, R. 1994. Autoregressive conditional heteroskedasticity and changes in regime, *Journal of Econometrics*, 64(1-2): 307–333.

Hu, L. & Shin, Y. 2008. Optimal test for Markov switching GARCH models, *Studies in Nonlinear Dynamics & Econometrics*, 12(3).

Krolzig, H.-M. 1997. *Markov-switching vector autoregressions modelling, statistical inference, and application to business cycle analysis*. Springer.

Lamoureux, C. G. & Lastrapes, W. D. 1990. Persistence in variance, structural change, and the GARCH model, *Journal of Business & Economic Statistics*, 8(2): 225–34.

Ng, V. K. & Pirrong, S. C. 1994. Fundamentals and volatility: Storage, spreads, and the dynamics of metals prices, *The Journal of Business*, 67(2): pp. 203–230.

Roache, S. K. 2010. What explains the rise in food price volatility?, *IMF Working Papers 10/129*, International Monetary Fund.

Samuelson, P. 1971. Stochastic speculative price, *Proceedings of the National Academy of Sciences*, 68: 335–337.

Scheinkman, J. & Schechtman, J. 1983. A simple competitive model with production and storage, *Review of Economic Studies*, 50: 427–441.

Schwager, J. D. 1984. *A complete guide to the futures markets: fundamental analysis, technical analysis, trading, spreads and options*. Wiley-Interscience.

Timmer, C. P. 2010. Management of rice reserves stocks in Asia: Analytical issued and country experience, in *Commodity Market Review 2009-2010*, pp. 87–120. FAO.

von Braun, J. & Torero, M. 2009. Implementing physical and virtual food reserves to protect the poor and prevent market failure, *Policy briefs 10*, International Food Policy Research Institute (IFPRI).

Wright, B. D. & Williams, J. C. 1982. The economic role of commodity storage, *Economic Journal*, 92(367): 596–614.

——— 1984. The welfare effects of the introduction of storage, *The Quarterly Journal of Economics*, 99(1): 169–92.