



The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search

<http://ageconsearch.umn.edu>

aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

No endorsement of AgEcon Search or its fundraising activities by the author(s) of the following work or their employer(s) is intended or implied.

Crop price comovements during extreme market downturns

David M. Zimmer[†]

This study develops and estimates mixture models of crop price comovements using copula functions, which allow for departures from normality during extreme market circumstances. The models also account for unique time-series patterns inherent in crop price data. The results point to two main conclusions. First, mixture models appear to provide an easy-to-estimate approach for capturing real-life crop price movements. Second, mixture models find that, during extreme market downswings, correlations in price movements strengthen by several orders of magnitude. These results suggest that structured securities assembled from different crops tend to lose diversified protection during extreme market downswings, the exact times when such protection is needed most.

Key words: Clayton, finite mixture, latent class, structured finance.

1. Introduction

While most markets experience unexpected shifts in supply and demand, few markets face the vast array of uncertainties present in agricultural markets. Financial traders of crop-based assets, seeking diversified protection from such uncertainties, create structured securities that pool multiple crops into unified assets, which then are traded in secondary markets. However, understanding risk built into such structured securities requires a detailed understanding of correlations in crop prices. This study develops and estimates mixture models, which reveal several previously hidden features of crop price movements. The main finding is that, in the midst of a sharp market downturn, correlations in crop prices strengthen substantially, thus weakening the extent of diversified protection during the exact times when such diversification is needed most.

In addition to obvious sources of uncertainty, such as temperature, flooding, drought, disease and infestation, crop markets also face uncertainties in the political arena. Support for programs such as ethanol mandates, farm subsidies and crop insurance often shifts depending on budgetary pressures, as well as which political party currently controls key agricultural committees. Currently, the U.S. federal government pays on average 62 per cent of farmers' crop insurance premiums, but that percentage is frequently adjusted, often unpredictably so, when farm bills come up for renewal. Such

[†]David M. Zimmer (email: david.zimmer@wku.edu) is at Department of Economics, Western Kentucky University, Bowling Green, Kentucky, USA.

uncertainties affect not only crop prices, but also a myriad of downstream markets, from animal feed to motor fuel, that directly or indirectly use crops (Moss 2010). In fact, uncertainty in crop markets generates such price volatility that, in calculating measures of market-wide inflation for guiding monetary policy, the Federal Reserve prefers a 'core' measure of inflation that removes food (and fuel) prices.

Structured securities that combine different crops offer diversified protection from such uncertainty. Of course, such instruments are not unique to crop markets. Structured mortgage-based securities are a common, albeit shrinking, presence in housing finance and are often cited as a root cause of the housing crisis (Zimmer 2015). And much like mortgage-based structured securities, assessing the riskiness of such assets requires estimates of the extent to which crop prices move in synchrony. If comovements are reasonably small, then assets that pool multiple crops should have diversified protection in the event that one crop experiences a sudden collapse in price.

A convenient choice for estimating price comovements is the Gaussian copula (Woodard *et al.* 2011), which owes its popularity, in part, to its connection to the familiar multivariate normal distribution. The Gaussian copula also conveniently connects to linear regression set-ups, in that the measure of dependence in the Gaussian copula captures the same linear correlation reflected in the coefficients in linear regressions (Goodwin 2014). But unfortunately, the Gaussian copula, much like the related normal distribution, imposes certain restrictions that might not apply to real-life crop price movements. First, the Gaussian copula imposes symmetry, such that correlations in price movements must be the same during market upswings and downswings. Second, the Gaussian copula imposes asymptotic independence, such that, regardless of the magnitude of correlations in price movements, during extreme markets swings, prices must move independently. Goodwin and Hungerford (2014) provide evidence that real-life crop price movements violate both of these restrictions. They develop more flexible, higher-dimensional models of revenues received from crop sales, with particular emphasis on what those correlations imply for federally backed crop insurance premiums.

The statistical implication is that jointly related asset prices potentially exhibit departures from normality during unique market circumstances. Financial economists have long recognised that univariate distributions of financial returns depart from normality (Blanchard and Watson 1982). Recent evidence finds similar results in multivariate distributions of stock prices (Lux and Sornette 2002), hedge fund returns (Li and Kazemi 2007; Boyson *et al.* 2010) and housing prices (Zimmer 2015). Rodriguez (2007), using procedures similar to those employed in this paper, confirms departures from normality are especially prevalent in the midst of financial turmoil.

The fact that crop prices tend to move in synchrony, coupled with possible departures from normality, calls for a multivariate model capable of accommodating flexible dependence patterns. Existing literature offers several examples of such models constructed using copula functions (Chen and Fan

2006; Hu 2006; Patton 2006; Goodwin *et al.* 2012). In this paper, versions of those methods are used to investigate comovements in crop prices. The approach is consistent with common-shock models, correlated heterogeneity, and single- and multifactor models often used to explain theoretical underpinnings of financial dependence and contagion (Corsetti *et al.* 2005).

The results point to three main conclusions. First, mixture models, by freeing up each part of the ‘mixture’ to capture the part of the distribution to which it is best suited, provide an easy-to-estimate approach for flexibly modelling real-life crop price movements. Second, mixture models that allow for departures from normality during extreme market downswings find that, during such extreme market downswings, correlations in price movements strengthen by several orders of magnitude. Third, the mixture model finds that, compared to findings from the Gaussian copula, the probability of two crops experiencing simultaneous price collapses is 38–50 per cent larger. Altogether, these results suggest that structured securities assembled from different crops tend to lose diversified protection during extreme market downturns, the exact times when such protection is needed most.

2. Data

Data on crop prices come from the USDA National Agricultural Statistics Service, compiled on the website *Understanding Dairy Markets*, maintained by Brian Gould of the Agricultural and Applied Economics Department at the University of Wisconsin.¹ Each data point represents the monthly average price, measured as dollars per bushel, received by farmers for three crops – corn, soya beans and wheat – which represent, by far, the most acreage among food grains in the USA.² Reflecting the emphasis on price movements, the series are transformed such that each data point represents the percentage change in price from the previous month. The series cover the period February 1949 through January 2013, for a total of $T = 768$ observations.³

Figure 1 provides a graphical depiction of the three series. The figure demonstrates the relatively high volatility in the price series. Certain periods, such as the mid-1970s and late 2000s, show especially high volatility, with monthly prices occasionally fluctuating by more than 30 per cent.

Overall, the figure points to a history of crop prices characterised by large and unpredictable price swings. Ignoring such time-series details tends to lead to incorrect findings of spurious correlation (Granger and Newbold 1974;

¹ <http://future.aae.wisc.edu/index.html>

² Acreage numbers are available at <http://www.epa.gov/oecaagct/ag101/>

³ All methods in this paper also were applied to a daily data set of prices for corn-based, soya bean-based and wheat-based investment funds operated by London-based ETF Securities. (Whereas the monthly data used in this paper begin in 1949, these daily data begin in 2006.) The goal was to explore whether daily data produce different findings from the monthly data used in this paper. Qualitative findings using those daily data were similar to those reported in this paper.

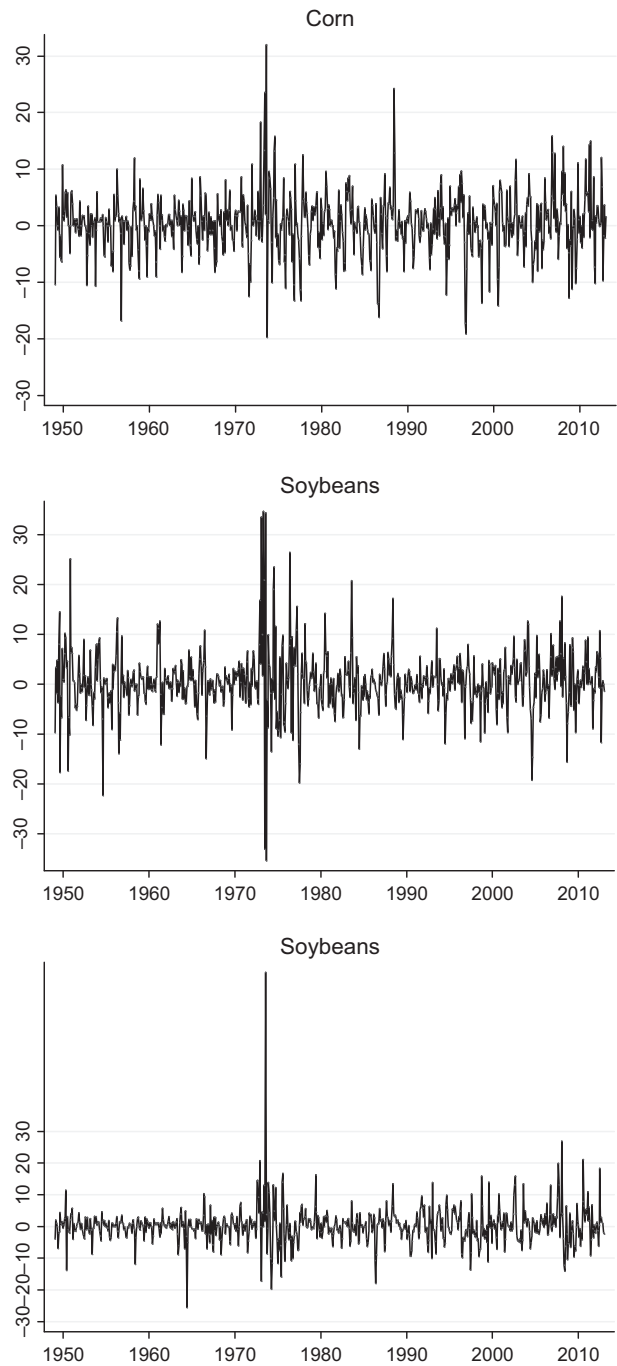


Figure 1 Monthly per cent change in crop prices.

Deb *et al.* 1996; Chen and Fan 2006). The following sections present econometric models capable of accommodating those time-series features inherent in crop price data.

3. Marginal distributions

The proposed model of price comovements proceeds in two steps. This section presents the first of those two steps, in which univariate price series are ‘filtered’ of their autoregressive and heteroscedastic components, in order to ensure that measures of dependence are not clouded by those time-series patterns. The following section then turns to the second step, in which the filtered price series are jointly analysed using copula functions. The two-step procedure resembles models used in finance to study exchange rate dynamics (Chen and Fan 2006; Patton 2006).

The first step estimates univariate AR-GARCH models on the three series. Yet, considering the relatively lengthy time series used in this paper, the possibility of structural breaks during the series potentially renders GARCH-type models badly biased (Kokoszka and Leipus 1999, 2000; Hillebrand 2005). Therefore, to check for possible structural breaks, Figure 2 presents CUSUM tests, for which the null hypothesis is of no structural breaks (Johnston and Dinardo 1997, pp. 119–120). For none of the three price series does the CUSUM test support rejection of the null hypothesis.

Thus, it is assumed in this paper that each price series y_{it} (for $j = 1, 2, 3$), where each data point represents the percentage change in price from the previous month, and where

$j = 1$ denotes corn

$j = 2$ denotes soya beans

$j = 3$ denotes wheat

follows a univariate AR(1)-GARCH(1,1) specification

$$y_{j,t} = \beta_{j,0} + \beta_{j,1}y_{j,t-1} + \varepsilon_{j,t}$$

where the error term $\varepsilon_{j,t}$ follows a normal distribution with mean zero and conditional variance given by

$$\sigma_{j,t}^2 = \alpha_{j,0} + \alpha_{j,1}\varepsilon_{j,t-1}^2 + \delta_j\sigma_{j,t-1}^2$$

Coefficient estimates of the AR-GARCH models appear in Table 1.

After estimation, new series $\tilde{y}_{i,j}$ are calculated as

$$\tilde{y}_{j,t} = \frac{\tilde{\varepsilon}_{j,t}}{\sqrt{\hat{\alpha}_{j,0} + \hat{\alpha}_{j,1}\hat{\varepsilon}_{j,t-1}^2 + \hat{\delta}_j\hat{\sigma}_{j,t-1}^2}}$$

The newly created series $\tilde{y}_{1,t}, \tilde{y}_{2,t}, \tilde{y}_{3,t}$, hereafter referred to as ‘filtered’ price series, remove the presence of autoregressive and GARCH patterns from the data, ensuring that remaining correlations calculated between the series are not a consequence of those patterns. As a formal check, Table 2 presents P -values for Breusch–Godfrey tests (Breusch 1978; Godfrey 1978) of zero

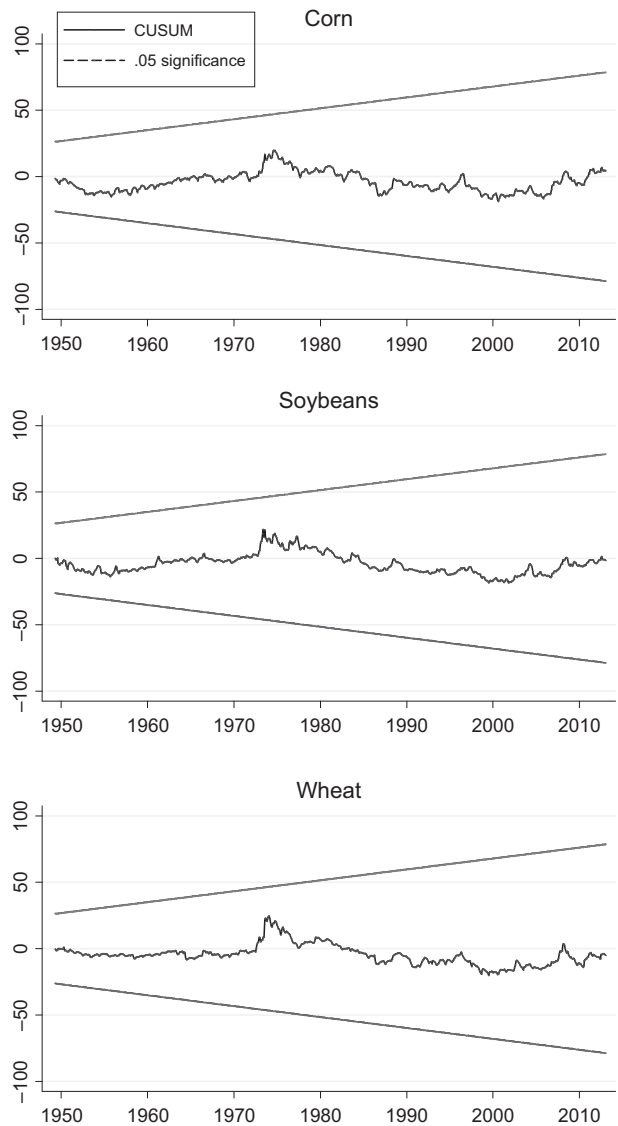


Figure 2 CUSUM tests for structural breaks.

autocorrelation, and Bollerslev (1986) tests of zero autoregressive conditional heteroscedasticity. The tests confirm that the AR(1)-GARCH(1,1) models remove those contaminants.

After calculating the filtered price series, they are plugged into the cdf and pdf of the standard normal distribution, denoted, respectively, as $F_j(\tilde{y}_{j,t})$ and $f_j(\tilde{y}_{j,t})$. These are the marginal distributions that appear in the following section.

Of course, specifying the marginal distributions as normal requires that the filtered price series approximate normality. However, as shown in the right-

Table 1 Estimates of marginal AR(1)-GARCH(1,1) models

	Corn		Soya beans		Wheat	
	Est.	St. Err.	Est.	St. Err.	Est.	St. Err.
$\beta_0(\text{constant})$	0.473	0.279	0.121	0.249	0.095	0.219
β_0 (AR 1)	0.377	0.037	0.363	0.040	0.471	0.034
α_0 (constant)	10.681	1.511	2.374	0.417	2.782	0.299
α_0 (ARCH 1)	0.256	0.054	0.230	0.027	0.702	0.079
δ (GARCH 1)	0.299	0.091	0.704	0.032	0.456	0.035

Table 2 Specification tests for filtered price series (*P*-values)

	Breusch–Godfrey test H_0 : zero autocorrelation	Bollerslev test H_0 : zero ARCH	Jarque–Bera test H_0 : normal
Corn	0.47	0.63	<0.01
Soya beans	0.16	0.41	<0.01
Wheat	0.77	0.48	<0.01

hand column of Table 2, Jarque–Bera tests (Jarque and Bera 1987) reject normality for all the three series.⁴ Figure 3, which presents kernel density estimates of the filtered price series, shows that rejection stems not from skewness, but rather from the fact the filtered price series show excess kurtosis. One solution is to specify marginals that allow leptokurtotic patterns. Another possibility is to use nonparametric marginals. Such an approach was attempted in this paper using kernel density estimates as suggested by Goodwin (2013). Unfortunately, the mixture model described in the following section, which uses parametrically specified copulas, achieved convergence for only one of the crop pairs (corn/soya beans). Results from that estimation were similar to those reported below.

Therefore, the marginals have been left as normal, but are switched to a fully nonparametric model in Section 7 as a way to test whether the misspecified marginals produce erroneous conclusions. Results of that robustness check, discussed in more detail in Section 7, suggest that assuming normality for the marginals, while obviously incorrect, does not impact the paper's main results.

4. Copula model

Let $\tilde{y}_{j,t}$ ($j = 1, 2, 3$) denote the filtered series created in the previous section. The empirical goal is to quantify the magnitude of comovements between the

⁴ Adding lags to the AR-GARCH specification did not change this result, nor did specifying alternative distributions for $\varepsilon_{j,t}$, such as t-distributions with various degrees of freedom or generalised error distributions with various shape parameters.

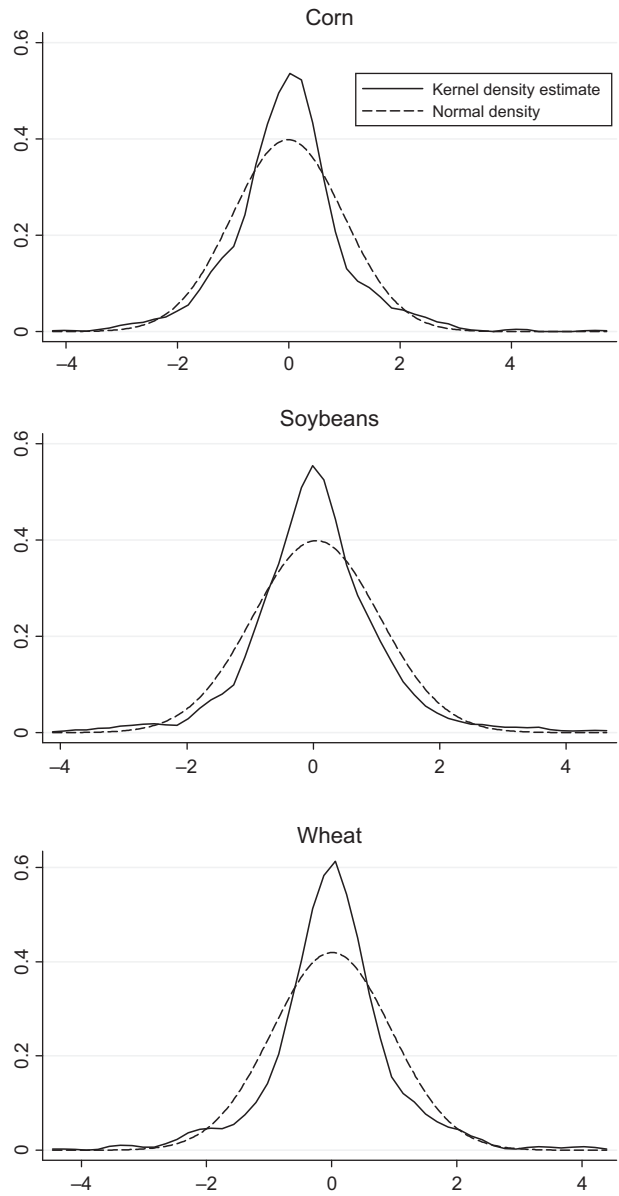


Figure 3 Kernel density estimates of filtered price changes.

three filtered price series, which ideally could be accomplished by specifying the trivariate distribution $F(\tilde{y}_{1,t}, \tilde{y}_{2,t}, \tilde{y}_{3,t}; \Theta)$, where Θ captures the correlation structure. However, working directly with that trivariate distribution is difficult for several reasons. First, most distributional forms, besides the familiar elliptical ones such as the normal and t, do not readily extend to higher-than-univariate settings, let alone the trivariate one needed here. Second, the normal and t impose unrealistic patterns on correlation

structures. Specifically, the normal imposes tail independence, such that price movements show zero correlation during extreme market swings. And although the t does accommodate tail dependence, it restricts the magnitude of that tail dependence to be identical for each pair of crops, which does not seem realistic for crop price data.

Instead, this paper considers the three pairwise distributions: $F(\tilde{y}_{1,t}, \tilde{y}_{2,t}; \theta_{12})$, $F(\tilde{y}_{1,t}, \tilde{y}_{3,t}; \theta_{13})$ and $F(\tilde{y}_{1,t}, \tilde{y}_{3,t}; \theta_{23})$. The main interest lies in the dependence terms θ_{12} , θ_{13} and θ_{23} . Considering separate bivariate models, although seemingly inferior to a unified trivariate set-up, draws inspiration from ‘composite likelihood’ methods, pioneered by Lindsay (1988), which hold that, in likelihood settings, higher-dimensional likelihood functions can be pieced together from individual lower-dimensional pieces. To give a simple motivation for this method for multivariate densities, the trivariate density $f(\tilde{y}_{1,t}, \tilde{y}_{2,t}, \tilde{y}_{3,t}; \Theta)$ can be approximated by the product of its three pairwise bivariate densities, $f(\tilde{y}_{1,t}, \tilde{y}_{2,t}; \theta_{12}) f(\tilde{y}_{1,t}, \tilde{y}_{3,t}; \theta_{13}) f(\tilde{y}_{1,t}, \tilde{y}_{3,t}; \theta_{23})$. When this expression is logged for purposes of maximum-likelihood estimation, it becomes $\ln f(\tilde{y}_{1,t}, \tilde{y}_{2,t}; \theta_{12}) f(\tilde{y}_{1,t}, \tilde{y}_{3,t}; \theta_{13}) f(\tilde{y}_{1,t}, \tilde{y}_{3,t}; \theta_{23})$. Then, because the three dependence terms, θ_{12} , θ_{13} and θ_{23} , are contained within additively separable pieces in the likelihood function, each can be consistently estimated by considering the individual bivariate pieces. (See Bhat (2014) for a detailed discussion of these methods, including proofs of consistency theorems.)

The implication, therefore, is that the bivariate pairwise terms, θ_{12} , θ_{13} and θ_{23} , provide consistent estimates for the ‘true’ trivariate correlation structure contained in Θ , even though the procedure ignores possible higher-dimensional linkages. The main advantage of composite methods is that bivariate distributions offer (slightly) simpler modelling options. Composite methods generally involve efficiency loss, but they might offer protection from incorrect specification of higher-dimensional likelihood functions (Bhat 2014).

However, working with separate bivariate distributions is not substantially easier than working with the larger trivariate counterpart. As a solution, this paper assembles the three bivariate distributions using copula functions, which have well-developed and well-understood properties in bivariate settings. (See Goodwin *et al.* (2012) for a discussion of ‘vine copula’ methods for forming higher-dimensional copulas.) A copula function is a multivariate cumulative distribution function for which all univariate marginal distributions are uniformly distributed on (0,1) (Sklar 1973). The key detail is that, because the range of any univariate cumulative distribution functions (cdf) follows a uniform (0,1) distribution, regardless of the underlying distribution itself, any multivariate cdf can be expressed in terms of its individual marginal cdfs and a copula function that glues the marginals together. For ease of notation, the following derivation is shown for the first of the three distributions, $F(\tilde{y}_{1,t}, \tilde{y}_{2,t}; \theta_{12})$; the other two bivariate distributions are derived similarly. Sklar’s Theorem implies that this bivariate distribution with individual marginal distributions F_1 and F_2 can be expressed as

$$F(\tilde{y}_{1,t}, \tilde{y}_{2,t}) = C(F_1(\tilde{y}_{1,t}), F_2(\tilde{y}_{2,t}); \theta_{12}) \tag{1}$$

where θ_{12} measures dependence. If the marginals are continuous, as is the case in this paper, then the corresponding copula in Eqn (1) is unique. (See Genest and Neslehova (2007) for complications that may arise in discrete settings.)

The bivariate density, useful for maximum-likelihood estimation, comes from differentiating,

$$f(\tilde{y}_{1,t}, \tilde{y}_{2,t}) = c(F_1(\tilde{y}_{1,t}), F_2(\tilde{y}_{2,t}); \theta_{12})f_1(\tilde{y}_{1,t})f_2(\tilde{y}_{2,t})$$

where $c(\cdot)$, called the ‘copula density’, is calculated as $\frac{\partial}{\partial(F_1\partial F_2)} C(F_1(\tilde{y}_{1,t}), F_2(\tilde{y}_{2,t}); \theta_{12})$, and f_1 and f_2 denote the densities of F_1 and F_2 .

The dependence parameter θ_{12} , which represents the main parameter of interest, does not translate conveniently between different copulas. To facilitate comparison, copula dependence parameters are converted to measures of concordance such Kendall’s τ (Nelsen 2006), which then can be compared across copulas. Kendall’s τ is bounded on the region $(-1,1)$ with $-1,0$, and 1 corresponding to perfect negative dependence, independence and perfect positive dependence. For copulas that allow tail dependence, separate measures of tail dependence also can be computed.

The copula approach offers two advantages. First, it separates the marginal behaviour of each variable from the dependence structure linking them together. This is important in the current setting, as the marginals must account for autoregressive and heteroscedastic time-series features of crop price data, which are easier to address when working with each individual series separately. The second advantage of the copula approach is that, with a variety of off-the-shelf copulas from which to choose, a researcher can find the one that appropriately reflects dependence patterns present in the data.

The main points of this paper can be made using two specific copulas, as well as a mixture of those two. The following table summarises those copulas.

Copula	Functional form	Kendall’s τ
Gaussian	$\Phi_2(\Phi^{-1}(F_1), \Phi^{-1}(F_2); \theta_{12})$	$2/\pi \arcsin \theta_{12}$
Clayton	$\left(F_1^{-\theta_{12}} + F_2^{-\theta_{12}} - 1\right)^{-1/\theta_{12}}$	$\theta_{12}/(1 + \theta_{12})$
Gaussian/Clayton mixture	$\pi \cdot \Phi_2(\Phi^{-1}(F_1), \Phi^{-1}(F_2); \theta_{12,G})$ $+ (1 - \pi) \cdot \left(F_1^{-\theta_{12,C}} + F_2^{-\theta_{12,C}} - 1\right)^{-1/\theta_{12,C}}$	

The term Φ^{-1} denotes the quantile function of the standard normal distribution, Φ_2 is the standard bivariate normal distribution, and $\pi \in (0, 1)$ is an estimable parameter indicating the proportional contribution of the

Gaussian copula. The mixture copula permits more flexible dependence patterns than either of the two individual component copulas. It does so by freeing up each individual component to capture the part of the distribution to which it is best suited. Specifically, the mixture copula has two dependence parameters, with the term $\theta_{12,G}$ capturing dependence during ‘normal’ times and the term $\theta_{12,C}$ capturing dependence during extreme market downturn.⁵

The copulas considered in this paper specify the dependence parameter as constant for the duration of the series. A more flexible approach would be to consider time-varying copulas (see Manner and Reznikova (2012) for a review of such methods). Unfortunately, in time-varying copulas, it can be difficult to untangle whether a particular market outcome is a tail event or the consequence of a changing bivariate distribution. As an informal check of the appropriateness of using a time-constant approach, bivariate dynamic conditional correlation (DCC) models were estimated with the same marginal distributions outlined above [AR(1)-GARCH(1,1)]. The errors in the DCC set-up were specified as bivariate t with three degrees of freedom, in order to accommodate (some) tail dependence. The time-varying correlations from those models, not shown but available upon request, did not show any obvious structural breaks, suggesting that dependence across the series can be approximated using a time-constant copula approach.⁶

5. Results

This section first presents evidence that, among the aforementioned copulas, the Gaussian/Clayton mixture appears to provide the best fit to the data. The section then proceeds to present and discuss estimates of dependence.

5.1 Tests of copula fit

The aforementioned copulas are not nested, nor do they contain the same numbers of estimated parameters. Therefore, to compare the different copulas, Table 3 presents Vuong (1989) tests of non-nested likelihood-based models. The first row reveals that the Gaussian copula is superior to its Clayton counterpart. Based upon such evidence, a researcher might conclude that the Gaussian offers an appropriate fit and that crop prices do not tend to follow Clayton patterns in which prices simultaneously collapse. However, as shown in more detail below, while the Gaussian indeed does describe the majority of the distribution, a fraction of the distribution – in fact, the part

⁵ In the statistics literature, mixture copulas are often referred to using ‘BB’ notation. For example, a Clayton/Gumbel mixture is often called BB1.

⁶ The bivariate DCC model for the soya beans/wheat pair failed to converge.

that is most important for characterising market collapses – adheres closer to the Clayton form.

The next two rows in Table 3 demonstrate that the Gaussian/Clayton mixture performs better than either of its individual components by themselves, although the mixture copula performs only marginally better than the Gaussian for the corn/wheat and soya beans/wheat pairs. Overall, the tests presented in Table 3 suggest that, among the three copulas considered, the Gaussian/Clayton mixture provides the most accurate fit to real-life crop price movements.

Despite providing a convenient method for ranking copulas, Vuong tests do not inform upon whether any of the three copulas conform to the true data-generating process. To that end, Table 4 presents Cramer-von Mises goodness-of-fit tests, with *P*-values calculated by parametric bootstrap (Genest *et al.* 2009). For each of the three crop pairs, the test statistics reject each of the three copulas, indicating that none of the three copulas, including the mixture specification preferred by the Vuong tests, adheres to the true data-generating process.

It is worth noting that the same Genest *et al.* (2009) study also attests to the impressive statistical power of the Cramer-von Mises test. Their power results, although reflecting positively on the statistical usefulness of the test, draw into question whether a *statistically* significant rejection of a particular copula is *economically* meaningful. Similar to the discussion of the marginal distributions above, the mixture copula is chosen as the preferred specification, even though it fails its goodness-of-fit test. To check whether the mixture copula affects the paper’s main conclusions, a nonparametric copula has been estimated (with nonparametric marginals) and is reported in Section 7. Results from that robustness check are very similar to those obtained from the mixture specification.

5.2 Estimates of dependence

Table 5 presents estimates of dependence, where for ease of comparison across copulas, all dependence parameters are converted to measures of Kendall’s τ , with standard errors calculated by the delta method. The table also presents estimates of lower tail dependence for the Clayton parts.

The top two rows show estimates for the single-component Gaussian and Clayton models. What is most noteworthy about the top two rows is that the

Table 3 Vuong tests

	Corn/Soya bean	Corn/Wheat	Soya bean/Wheat
H ₀ : Gaussian not superior to Clayton	4.45 (<0.01)	3.66 (<0.01)	3.52 (<0.01)
H ₀ : Mixture not superior to Gaussian	2.14 (0.02)	1.30 (0.10)	1.34 (0.09)
H ₀ : Mixture not superior to Clayton	5.82 (<0.01)	4.04 (<0.01)	3.92 (<0.01)

Note: *P*-values in parentheses.

Table 4 Cramer-von Mises goodness-of-fit tests

	Corn/Soya bean	Corn/Wheat	Soya bean/Wheat
H ₀ : Gaussian is true model	1.13 (<0.01)	1.62 (<0.01)	1.41 (<0.01)
H ₀ : Clayton is true model	1.95 (<0.01)	2.26 (<0.01)	1.67 (<0.01)
H ₀ : Mixture is true model	1.23 (<0.01)	1.63 (<0.01)	1.45 (<0.01)

Notes: *P*-values in parentheses calculated by parametric bootstrap (Genest *et al.* 2009).

Gaussian copula finds far stronger linkages between price movements than its Clayton counterpart. In fact, for two of the crop pairs (corn/wheat and soya beans/wheat), the Clayton estimates nearly zero dependence in the lower tails. A researcher might interpret this as lack of significant Clayton patterns in the data. Further, a researcher might conclude based on these findings that the Gaussian is a conservative choice for risk assessment, in that it finds stronger comovements. However, neither the Gaussian nor the Clayton allows distributional shifts in the event of severe market downturns.

The mixture copula, shown in the bottom panel of the table, allows each copula to capture the part of the distribution to which it is best suited. For each pair, the model finds that the majority of the distribution follows a Gaussian pattern: 77 per cent for corn/soya beans, 92 per cent for corn/wheat and 91 per cent for soya beans/wheat. These large proportions explain why the single-component Gaussian appears to outperform its Clayton counterpart according to the Vuong tests. Note that, in the mixture copula, estimates of dependence for the Gaussian component resemble those obtained from the single-component Gaussian model.

On the other hand, the mixture copulas reveal that nontrivial parts of the distribution adhere closer to Clayton patterns. And in those Clayton parts, dependence is far stronger in magnitude: 0.66 for corn/soya beans, 0.77 for corn/wheat and 0.75 for soya beans/wheat. The measures of lower tail dependence are even larger: 0.84, 0.90 and 0.89, respectively. The single-component Clayton is not capable of finding those parts of the distribution, as it must account for the entire distribution, most of which appears relatively Gaussian.

To visualise what these findings mean for crop price movements, Figure 4 shows contour plots of the Gaussian and mixture copulas, with estimated dependence parameters and the Gaussian proportional term (π) set equal to their converged values. For each crop price pair, the Gaussian distribution shows the familiar elliptical shape of the multivariate normal distribution. On the other hand, the mixture copulas, which retain a mostly Gaussian shape, reveal a tightly packed ridge in the south-west quadrant. The implication is that, when price movements land on that ridge during severe market downturns, correlations strengthen substantially, and as a consequence, portfolios based on those crops suddenly lose much of their diversified protection.

Table 5 Estimates of eependence

Copula	Corn/Soya bean	Corn/Wheat	Soya bean/Wheat
Gaussian τ	0.34 (0.02)	0.23 (0.02)	0.17 (0.02)
Clayton τ	0.23 (0.02) [0.31]	0.10 (0.02) [0.04]	0.06 (0.02) [0.01]
Mixture			
Gaussian τ	0.30 (0.02)	0.21 (0.02)	0.14 (0.02)
Clayton τ	0.66 (0.05) [0.84]	0.77 (0.07) [0.90]	0.75 (0.08) [0.89]
Gaussian proportion (π)	0.77 (0.06)	0.92 (0.04)	0.91 (0.04)

Notes: Standard errors, calculated by the delta method, appear in parentheses. Estimates of tail dependence appear in parentheses.

5.3 Probabilities of joint price collapses

Investors interested in diversified protection care about joint probabilities that prices simultaneously collapse. Copulas allow simple calculations of such joint probabilities. Recalling that the filtered price series $\tilde{y}_{j,t}$ follow distributions that are treated as standard normal, a price is said to ‘collapse’ if the filtered price falls below -1.645 (i.e. if the price decrease falls in the bottom 5 per cent of the distribution). Then, the probability of simultaneous price collapses is given by

$$Pr(\tilde{y}_{1,t} < -1.645 \text{ and } \tilde{y}_{2,t} < -1.645) = C(F_1(-1.645), F_2(-1.645); \hat{\theta}_{12})$$

for the Gaussian copula, and

$$\begin{aligned} Pr(\tilde{y}_{1,t} < -1.645 \text{ and } \tilde{y}_{2,t} < -1.645) \\ = \hat{\pi} \times C_{\text{gaussian}}(F_1(-1.645), F_2(-1.645); \hat{\theta}_{12,G}) + (1 - \hat{\pi}) \\ \times C_{\text{clayton}}(F_1(-1.645), F_2(-1.645); \hat{\theta}_{12,C}) \end{aligned}$$

for the mixture copula. The circumflexes denote converged parameter estimates.

Calculated probabilities appear in Table 6. Taking the corn/soya beans pair as an example, the Gaussian copula estimates a 1.2 per cent chance that both crops experience simultaneous collapses in prices. The mixture, on the other hand, estimates that risk of simultaneous collapse to be 1.8 per cent, which represents a 50 per cent larger probability compared to the Gaussian. The other two pairs reveal similar patterns, with the mixture finding a 38 per cent higher probability of simultaneous collapses in corn and wheat prices and a 50 per cent higher probability of simultaneous collapses in soya bean and wheat prices.

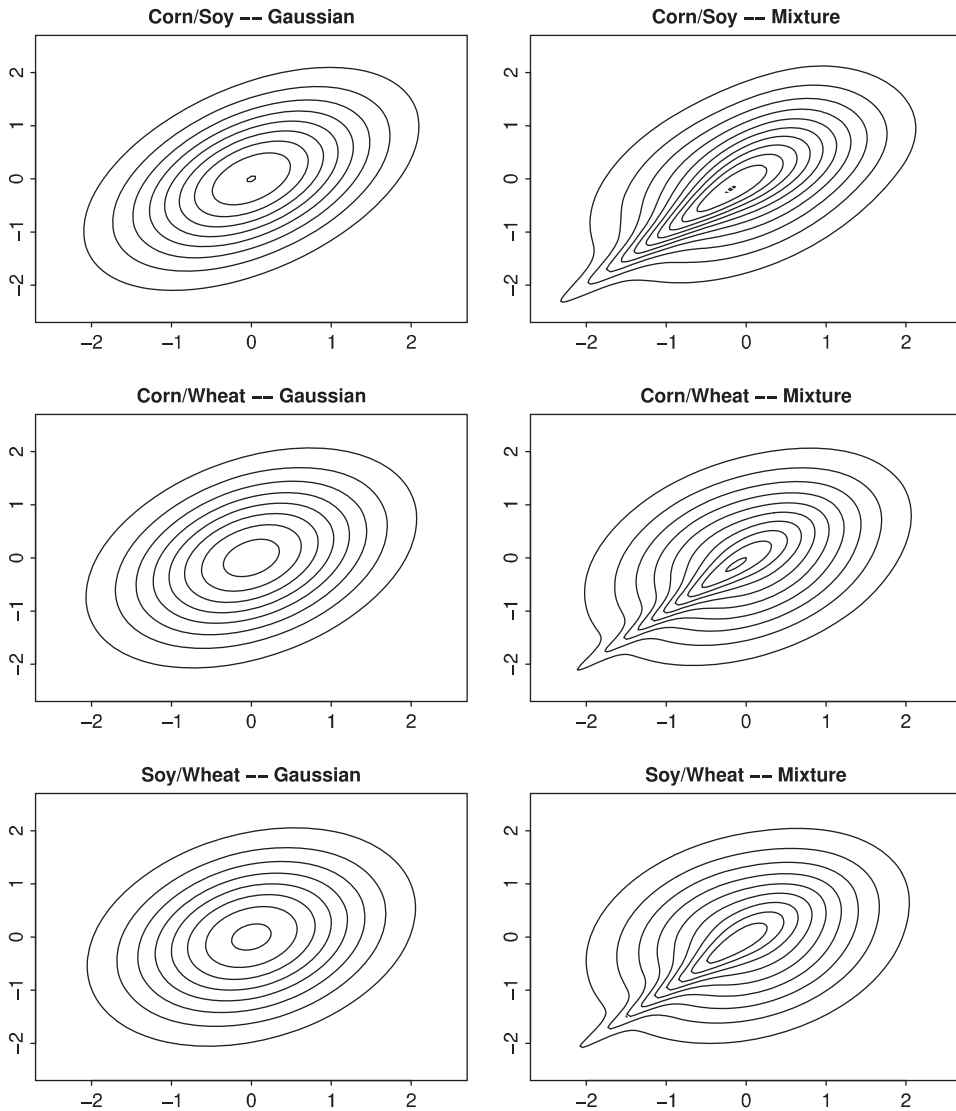


Figure 4 Estimated contour plots.

6. Investigating parametric assumptions

The results in this paper rely on normal marginals, but as shown in Table 2, Jarque–Bera tests reject normality for the filtered price series. Furthermore, as shown in Table 4, goodness-of-fit tests reject all three copulas, including the preferred mixture specification. Consequently, a valid concern is whether the main conclusions stem from incorrect parametric assumptions.

As a robustness check, this section reports estimation of a fully nonparametric specification involving calculation of both the marginals and the

copula empirically. That is, using the filtered price series, the marginals are calculated as

$$F_j(\tilde{y}_{j,t}) = \frac{1}{T+1} \sum_{s=1}^T 1(\tilde{y}_{j,s} < \tilde{y}_{j,t})$$

for $j = 1,2,3$, and where $1(\cdot)$ is the indicator function. The empirical copula function is calculated as

$$C(u, v) = \frac{1}{T} \sum_{t=1}^T 1(F_{t,j} \leq u, F_{t,k} \leq v)$$

for $u, v \in (0, 1)$ and $j, k = 1,2,3$. This empirical approach is not used as the baseline specification, because empirical calculations require lots of data, especially if one wishes to focus on specific parts of the joint distribution. Nonetheless, as a robustness check, this empirical approach should inform upon whether the parametric assumptions in this paper contaminate its main conclusions.

Figure 5 plots estimated probabilities of simultaneous price decreases for the empirical approach and compares those to the parametric mixture copula. The empirical copula uncovers probabilities that are similar to those obtained from the parametric model. Consequently, although the parametric model assumes forms that are rejected by specification tests, those assumptions do not appear to drive the main conclusions of this paper.

7. Conclusion

This paper estimates models of crop price comovements using several different copula function specifications. In comparing one-component Gaussian and Clayton models, the Clayton, which accommodates stronger dependence during extreme market downturns, does not appear to outperform its Gaussian counterpart. But this finding derives not from any probity attached to the Gaussian copula, but rather from the fact that real-life crop price movements appear to follow a mixture of distributions. Specifically, crop prices usually

Table 6 Probabilities of simultaneous price collapses

	Corn/Soya bean	Corn/Wheat	Soya bean/Wheat
Gaussian	0.012	0.008	0.006
Mixture			
Gaussian part	0.011	0.008	0.006
Clayton part	0.042	0.045	0.045
Gaussian proportion	0.77	0.92	0.91
Overall mixture	0.018	0.011	0.009
Mixture versus Gaussian % difference	50%	38%	50%




Figure 5 Probabilities of simultaneous price decreases.

follow Gaussian patterns, but when extreme market downturns arrive, the distribution appears to switch to something closer to the Clayton distribution.

The implication is that during extreme market downturns, correlations in crop price movements appear to strengthen. Consequently, structured securities that combine different crops might lose diversified protection

during times when such protection is needed most. This paper does not contend that a Gaussian/Clayton mixture suffices for all occasions. Indeed, financial analysts who study crop prices also might wish to consider mixtures that capture upper tail dependence or even dynamically changing dependence. Rather, this paper seeks to offer an endorsement of flexible models capable of capturing hidden, but potentially important, parts of distributions of crop price comovements.

References

- Bhat, C. (2014). The Composite Marginal Likelihood (CML) inference approach with applications to discrete and mixed dependent variable models, *Foundations and Trends in Econometrics* 7, 1–117.
- Blanchard, O. and Watson, M. (1982). Bubbles, Rational expectations and speculative markets, in Wachtel, P. (ed.), *Crisis in Economic and Financial Structure: Bubbles, Bursts, and Shocks*, Lexington Books, Lexington, pp. 295–316.
- Bollerslev, T. (1986). Generalized autoregressive conditional heteroscedasticity, *Journal of Econometrics* 31, 307–327.
- Boyson, N., Stahel, C. and Stulz, R. (2010). Hedge fund contagion and liquidity shocks, *Journal of Finance* 65, 1789–1816.
- Breusch, T. (1978). Testing for autocorrelation in dynamic linear models, *Australian Economic Papers* 17, 334–355.
- Chen, X. and Fan, Y. (2006). Estimation and model selection of semiparametric copula-based multivariate dynamic models under copula misspecification, *Journal of Econometrics* 135, 125–154.
- Corsetti, G., Pericoli, M. and Sbracia, M. (2005). Some contagion, some interdependence: more pitfalls in tests of financial contagion, *Journal of International Money and Finance* 24, 1177–1199.
- Deb, P., Trivedi, P. and Varangis, P. (1996). Excess co-movement in commodity prices reconsidered, *Journal of Applied Econometrics* 11, 275–291.
- Genest, C. and Neslehova, J. (2007). A primer on copulas for count data, *The Astin Bulletin* 37, 475–515.
- Genest, C., Remillard, B. and Beaudoin, D. (2009). Goodness-of-fit tests for copulas: a review and a power study, *Insurance: Mathematics and Statistics* 44, 199–213.
- Godfrey, L. (1978). Testing against general autoregressive and moving average error models when the regressors include lagged dependent variables, *Econometrica* 46, 1293–1302.
- Goodwin, B. (2013). A note on a simplified and general approach to simulating from multivariate copula functions, *Applied Economics Letters* 20, 910–915.
- Goodwin, B. (2014). Agricultural Policy Analysis: The Good, the Bad, and the Ugly. 2014 AAEA Presidential Address.
- Goodwin, B. and Hungerford, A. (2014). Copula-based models of systemic risk in U.S. Agriculture: implications for crop insurance and reinsurance contracts, *American Journal of Agricultural Economics* 97, 879–896.
- Goodwin, B., Holt, M., Onel, G. and Prestemon, J. (2012). Copula-based Nonlinear Models of Spatial Market Linkages. Working paper, North Carolina State University.
- Granger, C. and Newbold, P. (1974). Spurious regressions in econometrics, *Journal of Econometrics* 2, 111–120.
- Hillebrand, E. (2005). Neglecting parameter changes in GARCH models, *Journal of Econometrics* 129, 121–138.
- Hu, L. (2006). Dependence patterns across financial markets: a mixed copula approach, *Applied Financial Economics* 16, 717–729.

- Jarque, C. and Bera, A. (1987). A test for normality of observations and regression residuals, *International Statistical Review* 55, 163–172.
- Johnston, J. and Dinardo, J. (1997). *Econometric Methods*, 4th edn. New York, McGraw-Hill.
- Kokoszka, P. and Leipus, R. (1999). Testing for parameter changes in ARCH models, *Lithuanian Mathematical Journal* 39, 182–195.
- Kokoszka, P. and Leipus, R. (2000). Change-point estimation in ARCH Models, *Bernoulli* 6, 513–539.
- Li, Y. and Kazemi, H. (2007). Conditional properties of hedge funds: evidence from daily returns, *European Financial Management* 13, 211–238.
- Lindsay, B. (1988). Composite likelihood methods, *Contemporary Mathematics* 80, 221–239.
- Lux, T. and Sornette, D. (2002). On rational bubbles and fat tails, *Journal of Money, Credit, and Banking* 34, 589–610.
- Manner, H. and Reznikova, O. (2012). A survey on time-varying copulas: specification, simulations, and application, *Econometric Reviews* 31, 654–687.
- Moss, C. (2010). *Risk, Uncertainty, and the Agricultural Firm*. World Scientific.
- Nelsen, R. (2006). *An Introduction to Copulas*, 2nd edn. New York, Springer.
- Patton, A. (2006). Modelling asymmetric exchange rate dependence, *International Economic Review* 47, 527–556.
- Rodriguez, J. (2007). Measuring financial contagion: a copula approach, *Journal of Empirical Finance* 14, 401–423.
- Sklar, A. (1973). Random variables, joint distributions, and copulas, *Kybernetika* 9, 449–460.
- Vuong, Q. (1989). Likelihood ratio tests for model selection and non-nested hypotheses, *Econometrica* 57, 307–333.
- Woodard, J., Paulson, N., Vedenov, D. and Power, G. (2011). Impact of copula choice on the modeling of crop yield basis risk, *Agricultural Economics* 42, 101–112.
- Zimmer, D. (2015). Analyzing comovements in housing prices using vine copulas, *Economic Inquiry* 53, 1156–1169.