Pricing Options on Commodity Futures:
The Role of Weather and Storage

Marin Bozic
Department of Agricultural and Applied Economics
University of Wisconsin-Madison
427 Lorch Street, Madison, WI 53705
Email: bozic@wisc.edu

T. Randall Fortenbery
Department of Agricultural and Applied Economics
University of Wisconsin-Madison
427 Lorch Street, Madison, WI 53705
Email: trforten@wisc.edu

PRICING OPTIONS ON COMMODITY FUTURES:

THE ROLE OF WEATHER AND STORAGE

Marin Bozic†
T. Randall Fortenbery
University of Wisconsin-Madison

April 24, 2011

ABSTRACT. Options on agricultural futures are popular financial instruments used for agricultural price risk management and to speculate on future price movements. Poor performance of Black's classical option pricing model has stimulated many researchers to introduce pricing models that are more consistent with observed option premiums. However, most models are motivated solely from the standpoint of the time series properties of futures prices and need for improvements in forecasting and hedging performance. In this paper we propose a novel arbitrage pricing model motivated from the economic theory of optimal storage, and consistent with implications of plant physiology on the importance of weather stress. We introduce a pricing model for options on futures based on a Generalized Lambda Distribution (GLD) that allows greater flexibility in higher moments of the expected terminal distribution of futures price. We use times and sales data for corn futures and options for the period 1995-2009 to estimate the implied skewness parameter separately for each trading day. An economic explanation is then presented for inter-year variations in implied skewness based on the theory of storage. After controlling for changes in planned acreage, we find a statistically significant negative relationship between ending stocks-to-use and implied skewness, as predicted by the theory of storage. Furthermore, intra-year dynamics of implied skewness reflect the fact that resolution of uncertainty in corn supply is resolved between late June and middle of October, i.e. during corn growth phases that encompass corn silking through grain maturity. Impacts of storage and weather on the distribution of terminal futures price jointly explain upward sloping implied volatility curves.

JEL Codes: G13, Q11, Q14

Keywords: arbitrage pricing model, options on futures, generalized lambda distribution, theory of storage, skewness

† Corresponding author. E-mail: bozic@wisc.edu

1. Introduction

Options written on commodity futures have been investigated from several aspects in the commodity economics literature. For example, Lence (1994), Vercammen (1995), Lien and Wong (2002), and Adam-Müller and Panaretou (2009) considered the role of options in optimal hedging. Use of options in agricultural policy was examined by Gardner (1977), Glauber and Miranda (1989), and Buschena (2008). The effects of news on options prices has been investigated by Fortenbery and Sumner (1993), Isengildina-Massa, Irwin, Good, and Gomez (2008) and Thomsen (2009). The informational content of options prices has been looked into by Fackler and King (1990), Sherrick, Garcia and Tirupattur (1996), and Egelkraut, Garcia, and Sherrick (2007). Some of the most interesting work done in this area considers modifications to the standard Black-Scholes formula that accounts for non-normality (skewness, leptokurtosis) of price innovations, heteroskedasticity, and specifics of commodity spot prices (e.g. mean-reversion). Examples include Kang and Brorsen (1995), and Ji and Brorsen (2009).

In this article we revisit the well-known fact that the classical Black's (1976) model is inconsistent with observed option premiums. Previous studies like Fackler and King (1990) and Sherrick et al. (1996) address this puzzle by identifying properties of futures prices that deviate from assumptions of Black's model, i.e. leptokurtic and skewed distributions of the logarithm of terminal futures prices and stochastic volatility. A common feature of past studies is the grounding of their arguments in the time-series properties of stochastic processes for futures prices and the distributional properties of terminal futures prices. In other words, their arguments are primarily statistical. In contrast to previous studies, we offer an economic explanation for the observed statistical characteristics. In this paper we analyze in detail options on corn futures. The focus is on presenting an alternative pricing model that is not motivated by improving the

3

forecasts of options premiums compared to Black's or other models, but by linking option pricing models with the economics of supply for annually harvested storable agricultural commodities. In particular, we demonstrate the effect of storability and crop physiology (i.e. susceptibility to weather stress) on higher moments of the futures price distribution. Only by understanding these fundamental economic forces can we truly explain why classical option pricing models work so poorly for commodity futures.

The article is organized as follows. In the next section we examine in detail the implications of Black's classical option pricing model on the shape and dynamics of the futures price distribution. We follow by presenting the rational expectations competitive equilibrium model with storage, and a testable hypothesis on conditional new crop price distributions that follows from it. In addition to storage, we present the agronomical research on the impact of weather on corn yields. We then develop a novel arbitrage pricing model for options on commodity futures based on the Generalized Lambda Distribution (GLD) which we propose to use in calibrating skewness of new crop futures price to match observed option premiums. The third section describes the econometric model. In the fourth section we summarize the data used in econometric analysis. Finally, we describe the estimation procedure and present results of statistical inference, followed by a set of conclusions and directions for further research.

2. Theory

2.1. Foundations of arbitrage pricing theory for options on futures

Black (1976) was the first to offer an arbitrage pricing model for options on futures contracts. Despite numerous extensions and modifications proposed in the literature, and the inability of the model to explain observed option premiums, traders still use this model in practice. This is likely

due to its simplicity and ability to forecast option premiums after appropriate "tweaks" are put in

place. Black proposes that futures prices follow a stochastic process as described below:

$$dF = \sigma F dz \tag{1}$$

where F stands for futures price, σ for volatility, and dz is an increment of Brownian motion.

The implication is that futures prices are unbiased expectations of terminal futures prices (ideally

equal to the spot price at expiration), and the stochastic process followed by futures prices is a

geometric Brownian motion.

Under this scenario the option premium V is equal to the present value of the expected option

payoff under a risk-neutral distribution for terminal prices. For example, for a call option with

strike K, volatility σ, risk-free interest rate r and time left to maturity T:

$$V\left(K,F_0,T,\sigma,r\right) = e^{-rT} \int_0^\infty Max\left(F_T - K, 0\right) f\left(F_T; F_0, \sigma, r, T\right) dF_T \tag{2}$$

Because delta hedging with futures does not require a hedger to pay the full value of the futures

contract due to margin trading, a risk-neutral terminal distribution for futures prices is equivalent

to a risk-neutral terminal distribution for a stock that pays a dividend yield equal to the risk-free

interest rate:

$$\ln F_T \sim N\left(\ln F_0 - \frac{1}{2}\sigma^2, \sigma^2 T\right) \tag{3}$$

Thus, Black's model postulates that the distribution of terminal futures prices, conditional on

information known at time zero, is lognormal with the first four moments fully determined by the

current futures price and volatility parameter σ. In particular, the first four moments of the risk-

neutral terminal distribution are equal to:

$$\tilde{\mu} = F_0 \quad \tilde{\sigma}^2 = F_0^2 \left(e^{\sigma^2 t} - 1 \right) \quad SKEW = (e^{\sigma^2 t} + 2)\sqrt{e^{\sigma^2 t} - 1} \quad KURT = e^{4\sigma^2 t} + 2e^{3\sigma^2 t} + 3e^{2\sigma^2 t} \quad (4)$$

For example, if a futures price is $2.50, volatility is 30%, and there are 160 days left to maturity, the standard deviation of the terminal distribution would be $0.50, skewness would be 0.60 and kurtosis would be 3.64. Therefore, the standard Black's model implies that the expected distribution of terminal prices would be positively skewed, and leptokurtic. When complaints are raised that Black's model imposes normality restrictions, it is the logarithm of the terminal price that the critique refers to.

The standard way to check if Black's model is an appropriate pricing strategy is to exploit the fact that for a given futures price, strike price, risk-free interest rate, and time to maturity, the model postulates a one-to-one relationship between the volatility coefficient and the option premium. Thus, the pricing function can be inverted to infer the volatility coefficient from an observed option premium. Such coefficients are referred to as *implied volatility* and the principal testable implication of Black's model is that implied volatility does not depend on how deep in-the-money or out-of-money an option is. If the logarithm of terminal price is not normally distributed, then Black's model is not appropriate, and implied volatility (IV) will vary with option moneyness – a flagrant violation of the model's assumptions. Black's model gives us a pricing formula for European options on futures, i.e. options that can only be exercised at contract maturity. Prices of American options on futures that are assumed to follow the same stochastic process as in Black's model must also account for the possibility of early exercise. For that reason, their prices cannot be obtained through a closed-form formula, but must be estimated through numerical methods such as the Cox, Ross and Rubinstein (CRR) (1979) binomial trees.

Implied volatility curves for storable commodity products are almost always upward sloping. As an example consider the December 2006 corn contract. The futures price on June 26, 2006 was

$2.49/bu. As seen in Figure 1, the implied volatility curve associated with calculating IV using various December option strikes is strongly upward sloping, with the implied volatility coefficients for the highest strike options close to 15 percentage points higher than the implied volatility for options with lower strikes.

Geman (2005) calls this phenomenon an "inverse leverage effect," after the "leverage effect" proposed to explain downward sloping implied volatility curves for individual company stocks. However, this is a complete misnomer. As Black (1976b) explains, the leverage effect arises from the fact that as stock price declines, the ratio of a company's debt to equity value, its leverage, increases. If the volatility of company assets is constant, then as the equity share of assets declines, volatility in equity will increase. While the leverage effect has a coherent causal model to justify the term, nothing explains "inverse leverage effect."

We can gain further insight as to how Black's model performs if we plot the implied volatility curve for a single contract at different time-to-maturity horizons. As an example, consider December corn contracts in the years 2004 and 2006. As Figure 2 shows, three distinct patterns are noticeable. First, except when options are very near maturity, we always see an upward sloping implied volatility curve. Second, implied volatility of at-the-money options, i.e. options that have the strike price equal to the current futures price, rises almost linearly until the end of June, declines throughout the summer months, and then starts rising again. Finally, near maturity, volatility skews give way to symmetric volatility smiles. The implied volatility coefficient measures volatility on an annual basis, and the variance of the terminal price, conditional on time remaining to maturity, is $\sigma^2(T-t)$. So if uncertainty about the terminal price is uniformly resolved as time passes, implied volatility will not decrease, but will stay the same. Likewise, when the same amount of uncertainty needs to be resolved in a shorter time

interval implied volatility will increase. Therefore, linear increases in implied volatility from distant horizons up until June is best interpreted not as increases in day to day volatility of futures price changes, but a market consensus that the conditional variance of terminal prices is not much reduced before June.

While CRR binomial trees preserve the basic restrictions of Black's model, i.e. the normality of the log-prices terminal distribution, Rubinstein (1994, 1998) shows how that can be relaxed to allow for non-normal skewness and kurtosis. To illustrate the effect of skewness and kurtosis on Black's implied volatility we used Edgeworth binomial trees (Rubinstein, 1998). This allows for pricing options that exhibit skewed and leptokurtic distributions of terminal log-prices. As can be seen in panel 1 in Figure 3, zero skewness and no excess kurtosis (S=0, K=3) corresponds to a flat IV curve, i.e. CRR implied volatility estimated from options premiums is the same no matter what strike is used to infer it, just like Black's model would have it. A leptokurtic distribution will cause so called "smiles", i.e. options with strikes further away from the current futures price will produce higher implied volatility coefficients. Positive skewness creates an upward sloping curve, and negative skewness a downward sloping IV curve.

Faced with the inability of Black's model to explain observed option premiums, researchers and traders have pursued three different approaches to address this issue:

1) Start from the end: relax the assumptions concerning risk-neutral terminal distributions of underlying futures prices, i.e. allow for non-lognormal skewness and kurtosis. As long as delta hedging is possible at all times (i.e. markets are complete), it is still possible to calculate option premiums as the present value of expected option payoffs. Examples of this approach include Jarrow and Ruud (1982), Sherrick et al. (1996), and Rubinstein

(1998). While the formulas that derive option premiums as discounted expected payoffs assume that options are European, one can still price American options using implied binomial trees calibrated to the terminal distribution of choice (Rubinstein, 1994).

2) Start from the beginning: start by asking what kind of stochastic process is consistent with a non-normal terminal distribution? By introducing appropriate stochastic volatility and/or jumps, one might be able to fit the data just as well as by the approach above. Examples of this approach are Kang and Brorsen (1995), Hilliard and Reis (1998) and Ji and Brorsen (2009).

3) "Tweak it so it works good enough" approach: if one is willing to sacrifice mathematical elegance, the coherence of the second approach, and insights that might emerge from the first approach, and if the only objective is the ability to forecast day-ahead option premiums one can simply tweak Black's model. An example of such an approach would be to model the implied volatility coefficient as a quadratic function of the strike. Even though it makes no theoretical sense (this is like saying that options with different strikes live in different universes), this approach will work good enough for many traders. Just as in that famous saying by Yogi Berra (2010): "In theory, there is no difference between theory and practice. In practice, there is." A seminal article that evaluates the hedging effectiveness of such an approach is Dumas, Fleming and Whaley (1998). The authors find that for hedging purposes such an ad-hoc approach seems to work equally well compared to the more sophisticate and theoretically coherent models they evaluate.

In this article we take the first approach, and modify the Black's model by modifying the terminal distribution of futures price. Instead of a lognormal, we propose a generalized lambda distribution (GLD) developed by Ramberg and Schmeiser (1974) and introduced to options

pricing by Corrado (2001). An alternative would be to use Edgeworth binomial trees, but preliminary analysis showed that such an approach may not be adequate for situations where skewness and kurtosis are rather high. In addition, Edgeworth trees work with the skewness of terminal log-prices, while we prefer to have implied parameters for the skewness of terminal futures prices directly, not their logarithms. In addition, the GLD pricing model allows for a higher degree of flexibility in terms of skewness and kurtosis, i.e. its' parameters are rather easy to calibrate from observed options prices and it is straightforward to develop a closed-form solution for pricing options. While these are all favorable characteristics, it is in fact the ability to gain additional economic insight that truly justifies yet another option pricing model. GLD allows us to get an explicit estimate of skewness and kurtosis of the terminal distributions, that can used to make a strong connection between the economics of supply for storable agricultural commodities and financial models for pricing options on commodity futures.

2.2.   Theory of storage and time-series properties of commodity spot and futures prices

Deaton and Laroque (1992) used a rational expectations competitive storage model to explain nonlinearities in the time series of commodity prices: skewness, rare but dramatic substantial increases in prices, and a high degree of autocorrelation in prices from one harvest season to the next. The basic conclusion of their work was that the inability to carry negative inventories introduces a non-linearity in prices that manifests itself in the above characteristics.

This is an example of theory being employed in an attempt to replicate patterns of observed price data. In a similar fashion, but subtly different, Williams and Wright (1991) postulate that the moments of expected price distributions at harvest time vary with the current (pre-harvest) price and available carryout stocks, as shown in Figure 4. According to them, when observed at annual or quarterly frequency, spot prices exhibit positive autocorrelation that emerges because storage

allows unusually high or low excess demand to be spread out over several years. Furthermore, the variance of price changes depends on the level of inventory. When stocks are high, and the spot price is low, the abundance of stored stocks serves as a buffer to price changes, and variance is low. When stocks are low, and thus the spot price is high, stocks are not sufficient to buffer price changes. Finally, the third moment of the price change distribution also varies with inventories. Since storage can always reduce the downward price pressure of a windfall harvest, but cannot do as much for a really bad harvest, large price increases are more common than large decreases. The magnitude of this cushioning effect of storage depends on the size of the stocks. In conclusion, one should expect commodity prices to be mean-stationary, heteroskedastic and with conditional skewness, where both the second and third moments depend on the size of the inventories.

Testing the theory proceeds with this argument: if we can replicate the price pattern using a particular set of rationality assumptions, then we cannot refute the claim that markets indeed behave as described above. That is the road taken by Deaton and Laroque (1992) and Miranda and Rui (1995). However, since in the spot price series we only see the realizations of prices, not the conditional expectations of them, we cannot use spot price data to directly test what the market *expected* to happen. As such, predictions from storage theory focused on the scale and shape of expected distributions of new harvest spot prices have remained untested. In this paper we use options data to infer the conditional expectations of terminal futures prices, and therefore test the following prediction of the theory of storage:

- The lower inventories are, the more positive will be the skewness of the conditional harvest futures price distribution

This is tested using an options pricing formula based on the generalized lambda distribution to calibrate the skewness and kurtosis of expected (conditional) harvest futures price distributions. Implied parameters from the model are then used to test the hypotheses above.

2.3. The role of weather in intra-year resolution of price uncertainty

As illustrated in section 2.1., a very small share of uncertainty concerning the terminal price of a new crop futures contract is resolved before June. A large part of the uncertainty is resolved between late June and early October. The reason lies in corn physiology and the way weather stress impacts corn throughout the growing season. In the major corn producing areas of the U.S., corn is planted starting the last week of April. It takes about 80 days after planting for a plant to reach its reproduction stage, also known as corn silking. At this juncture the need for nutrients is highest, and moisture stress has a large impact on final yield. Weather continues to play an important role through the rest of the growing cycle, as summarized by Figure 5, taken from Shaw et al. (1988).

Beginning in July, the United States Department of Agriculture (USDA) publishes updated forecasts of corn yield per acre. At the beginning of the growing season, before corn starts silking, production forecasts ae generally based on estimated acres and historical trend yields. As can be seen in Figure 6, June forecasts of final yield deviated from the historical trend value essentially the same in both what was at the time the record-setting yield year 2004/2005 when final yield was 15 bushels above the trend, and the major draught year of 1988/89 when final yields were 32 bushels below the trend. However, uncertainty is quickly resolved in July and August. As shown in Figure 7, whereas June forecasts deviated from final estimates from the low of -11% in 1994/95 to high of 45% in 1988/89, the September estimate deviations ranged only from -7% to 12%. Besides weather, more precise methods used by USDA from August onwards

12

estimate final yields also contribute to decrease in uncertainty. Starting in late July, and first reported in August edition of the *Crop Production* report, final yields are estimated not only based on statistical models that control for trend and crop condition, but also include information obtained through grower-reported yield survey and objective measurement survey.

A testable hypothesis that emerges from these stylized facts concerns the fundamental role of seasonality in uncertainty resolution, as well as pronounced negative skewness in deviations of final yields from trend values. In other words, do seasonal yield deviations contribute to a positive skewness of the terminal price distribution and the dynamics of skewness throughout the marketing year? In particular, we might expect implied skewness to decrease throughout the growing season.

2.4.   Option pricing formula using generalized lambda distribution

The generalized lambda distribution (GLD) was developed by Ramberg and Schmeiser (1974), with Ramberg et al. (1979) further describing its properties. It was introduced to options pricing by Corrado (2001) who derived a formula for pricing options on non-dividend paying stocks. Here we review the properties of GLD and adopt Corrado's formula to options on futures.

GLD is most easily described by a percentile function[1] (i.e. inverse cumulative density function):

$$F(p) = \lambda_1 + \frac{p^{\lambda_3} - (1-p)^{\lambda_4}}{\lambda_2} \tag{5}$$

For example, to say that for $p = 0.90, F(p) = 4.5$ means that the market expects with a 90% probability that the terminal futures price will be lower than or equal to \$4.50/bu.

---

[1] F here stands for futures price, not for cumulative density function.

GLD has four parameters: λ_1 controls location, λ_2 determines variance, and λ_3 and λ_4 jointly determine skewness and kurtosis. In particular the mean and variance are calculated as follows:

$$\mu = \lambda_1 + A / \lambda_2$$
$$\sigma^2 = \left(B - A^2\right) / \lambda_2^2 \tag{6}$$

with $A = \dfrac{1}{1+\lambda_3} - \dfrac{1}{1+\lambda_4}$ and $B = \dfrac{1}{1+2\lambda_3} + \dfrac{1}{1+2\lambda_4} - 2\beta\left(1+\lambda_3, 1+2\lambda_4\right)$, where $\beta(\ )$ stands for the

complete beta function. We see that the λ_3 and λ_4 parameters influence both location and variance, however λ_1 influences only the first moment, and λ_2 influences only the first two moments. Thus, skewness and kurtosis do not depend on λ_1 and λ_2.

The skewness and kurtosis formulas are:

$$\alpha_3 = \frac{\mu_3}{\sigma^3} = \frac{C - 3AB + 2A^3}{\lambda_2^2 \sigma^3}$$

$$\alpha_4 = \frac{\mu_4}{\sigma^4} = \frac{D - 4AC + 6A^2 B - 3A^4}{\lambda_2^4} \tag{7}$$

where expressions for C and D are:

$$C = \frac{1}{1+3\lambda_3} - \frac{1}{1+3\lambda_3} - 3\beta\left(1+2\lambda_3, 1+\lambda_4\right) + 3\beta\left(1+\lambda_3, 1+2\lambda_4\right)$$

$$D = \frac{1}{1+4\lambda_3} + \frac{1}{1+4\lambda_3} - 4\beta\left(1+3\lambda_3, 1+\lambda_4\right) - 4\beta\left(1+\lambda_3, 1+3\lambda_4\right) + 6\beta\left(1+2\lambda_3, 1+2\lambda_4\right)$$

A standardized GLD has a zero mean and unit variance, and has a percentile function of the form:

$$F(p) = \frac{1}{\lambda_2(\lambda_3, \lambda_4)}\left(p^{\lambda_3} - (1-p)^{\lambda_4} + \frac{1}{\lambda_4+1} - \frac{1}{\lambda_3+1}\right) \tag{8}$$

with $\lambda_2(\lambda_3,\lambda_4)=sign(\lambda_3)\times\sqrt{B-A^2}$

From here, we can move more easily to an options pricing environment. We wish to make GLD

an approximate generalization of the log-normal distribution so we keep the mean and the

variance the same as in (4), while allowing skewness and kurtosis to be separately determined by

the λ_3 and λ_4 parameters. Therefore, the percentile function relevant for option pricing will be

$$F(p)=F_0\left(1+\frac{\sqrt{e^{\sigma^2 t}-1}}{\lambda_2(\lambda_3,\lambda_4)}\left(p^{\lambda_3}-(1-p)^{\lambda_4}+\frac{1}{\lambda_4+1}-\frac{1}{\lambda_3+1}\right)\right) \tag{9}$$

Note that this is equivalent to (5) with $\lambda_1=F_0+\dfrac{\sqrt{e^{\sigma^2 t}-1}}{\lambda_2(\lambda_3,\lambda_4)}\left(\dfrac{1}{\lambda_4+1}-\dfrac{1}{\lambda_3+1}\right)$

and $\lambda_2=\dfrac{\lambda_2(\lambda_3,\lambda_4)}{\sqrt{e^{\sigma^2 t}-1}}$. This will guarantee that the first two moments of the terminal distribution

will be $\tilde{\mu}=F_0 \quad \tilde{\sigma}^2=F_0^2\left(e^{\sigma^2 t}-1\right)$, just as in Black's model.

The pricing formula for European calls is

$$V(K,F_0,T,\sigma,r,\lambda_3,\lambda_4)=e^{-rT}\int_0^\infty Max(F_T-K,0)dp(F) \tag{10}$$

As shown by Corrado (2001), we can simplify this through a change-of-variable approach where

$F(p)=F_T$:

$$\int_0^\infty Max(F_T-K,0)dp(F)=\int_K^\infty(F_T-K)dp(F)=\int_{p(K)}^1(F(p)-K)dp \tag{11}$$

Here $p(K)$ stands for the cumulative density function, evaluated at K. While there is no closed

form formula for the function, values can be easily found with numerical approaches by using

the percentile function.

Integrating $F(p)$ we get

$$G_1 = \int_{p(K)}^1 F(p)\,dp = F_0\left(p + \frac{\sqrt{e^{\sigma^2 t}}-1}{\lambda_2(\lambda_3,\lambda_4)}\left(\frac{1}{\lambda_3+1}p^{\lambda_3+1} + \frac{(1-p)^{\lambda_4+1}}{\lambda_4+1} + \frac{1}{\lambda_4+1}p - \frac{1}{\lambda_3+1}p\right)\right)\Bigg|_{p(K)}^1$$

$$= F_0\left(1-p(K) + \frac{\sqrt{e^{\sigma^2 t}}-1}{\lambda_2(\lambda_3,\lambda_4)}\left(\frac{p(K)-p(K)^{\lambda_3+1}}{\lambda_3+1} + \frac{1-p(K)-(1-p(K))^{\lambda_4+1}}{\lambda_4+1}\right)\right)$$

with the final European call pricing formula being:

$$V\left(K,F_0,T,\sigma,r,\lambda_3,\lambda_4\right) = F_0 e^{-rt}G_1 - e^{-rt}KG_2 \tag{12}$$

where G_1 is defined above and $G_2 = 1 - p(K)$

In a similar way it can be shown that the price for a put is

$$V_P\left(K,F_0,T,\sigma,r,\lambda_3,\lambda_4\right) = e^{-rt}K\left(1-G_2\right) - F_0 e^{-rt}\left(1-G_1\right) \tag{13}$$

3. Econometric Model

3.1. Estimating implied skewness

Implied skewness is used as a dependent variable in subsequent econometric models, thus the first task at hand is to estimate implied higher moments of the terminal futures price distribution for a particular underlying futures contract. The GLD option pricing model can be used to price only European options, that is, options that can only be exercised at contract maturity. As mentioned before, options on corn futures are American options, i.e. they can be also exercised at any time before contract maturity. Therefore, for each option trade we use in fitting implied GLD higher moments, we first need to calculate the price at which such an option would trade if it indeed were of the European type. To do this, for each data point, we separately estimate

16

implied volatility using CRR binomial trees with 500 steps. Then, for each observation

separately, we use Black's model to calculate the price of a European option with same futures

price, strike, interest rate and time to maturity as that record for actually traded American option.

Using calibrated premiums for European options on corn futures, we then fit the following

option pricing model to options of a particular contract month:

$$O_i^E = V\left(K_i, F_{0i}, \tau, r, \sigma, \lambda_3, \lambda_4\right) + \varepsilon_i \tag{14}$$

where function used is as in (12) for calls or (13) for puts, O_i^E would be the previously calibrated

option premium for trade i for an option with strike K_i and with F_{0i} being the last observed

traded futures price prior to this trade. Observed parameters common to all options of the same

contract month traded on the same day include the interest rate r and the time to maturity

measured in calendar days, denoted as τ.

The unobserved generalized lambda distribution parameters $\sigma, \lambda_3, \lambda_4$ jointly determine variance,

skewness and kurtosis of the implied terminal distribution of futures prices, and are assumed to

be the same for all trades occurring on a single trading day. Implied parameters are fitted by a

nonlinear least squares model, minimizing squared differences between calibrated option

premiums for European options, and option premiums that arise from the GLD option pricing

model. Models are estimated separately for each trading day and each contract month traded at

that day.

3.2. Modeling intra-year dynamics of implied skewness

As we postulated in section 2.3., corn physiology in conjunction with weather patterns should

play a major role in governing the intra-year dynamics of implied skewness. The panels in Figure

8 present scatter diagrams of estimated implied skewness over the life of particular contract months. Each dot represents the estimated implied skewness on a particular trading day, with bolded diamonds being averages for a particular time-to-maturity horizon over the 15 marketing years used in estimation (1995-2009). Visual inspection does not contradict patterns we expected to see. In particular, new-crop contracts (September and December), exhibit near flat average implied skewness until late June, followed by a concave decrease for the September contract, and linear downward trend for December. Patterns for carry contracts (March, May and July) share strong and concave decreases in implied skewness over the last four months of contract life, with the effects on implied skewness during corn growth period not as distinct as for new-crop contracts. All five patterns stand in stark contrast to Black's model where variance of the terminal futures price distribution is assumed to be decreasing linearly in time. Given that Black's model stipulates the terminal distribution to be lognormal, a linear decrease in variance would correspond to a slightly convex and smooth decline in implied skewness.

If skewness in options on corn futures arises due to asymmetry in the ability of old-crop stocks to mitigate price effects of unexpected weather events during the growing season then skewness should exhibit different dynamics before corn silking, during the growing season, and post-harvest. To test this hypothesis, we fit implied skewness as a function of time using several models. Let implied skewness be denoted with IS_t. If options expire at time T, then the remaining time to maturity $T - t$ is denoted as τ. The models we test can then be written as

Linear model:

$$IS_t = \alpha + \beta_1 \tau + \varepsilon_t \tag{15}$$

Quadratic model:

18

$$IS_t = \alpha + \beta_1 \tau + \beta_2 \tau^2 + \varepsilon_t \qquad (16)$$

Linear model with one change in regime (timing is estimated endogenously):

$$IS_t = [\alpha_1 + \beta_1 \tau][\tau > \tilde{\tau}_1] + [\alpha_2 + \beta_2 \tau][\tau \le \tilde{\tau}_1] + \varepsilon_t$$
$$s.t. \quad \alpha_1 + \beta_1 \tilde{\tau}_1 = \alpha_2 + \beta_2 \tilde{\tau}_1 \qquad (17)$$

Quadratic model with one change in regime (timing is estimated endogenously):

$$IS_t = [\alpha_1 + \beta_1 \tau + \gamma_1 \tau^2][\tau > \tilde{\tau}_1] + [\alpha_2 + \beta_2 \tau + \gamma_2 \tau^2][\tau \le \tilde{\tau}_1] + \varepsilon_t$$
$$s.t. \quad \alpha_1 + \beta_1 \tilde{\tau}_1 + \gamma_1 \tilde{\tau}_1^2 = \alpha_2 + \beta_2 \tilde{\tau}_1 + \gamma \tilde{\tau}_1^2 \qquad (18)$$

Quadratic model with two changes in regime (timing is estimated endogenously):

$$IS_t = [\alpha_1 + \beta_1 \tau + \gamma_1 \tau^2][\tau > \tilde{\tau}_1] + [\alpha_2 + \beta_2 \tau + \gamma_2 \tau^2][\tilde{\tau}_2 \le \tau \le \tilde{\tau}_1] + [\alpha_3 + \beta_3 \tau + \gamma_3 \tau^2][\tau \le \tilde{\tau}_2] + \varepsilon_t$$
$$s.t. \quad \alpha_1 + \beta_1 \tilde{\tau}_1 + \gamma_1 \tilde{\tau}_1^2 = \alpha_2 + \beta_2 \tilde{\tau}_1 + \gamma_2 \tilde{\tau}_1^2 \qquad (19)$$
$$\alpha_2 + \beta_2 \tilde{\tau}_2 + \gamma_2 \tilde{\tau}_2^2 = \alpha_3 + \beta_3 \tilde{\tau}_2 + \gamma_3 \tilde{\tau}_2^2$$

Simple linear (15) and quadratic models (16) are used as benchmarks. In particular, it is

interesting to compare the performance of model (16) to more complicated models as model (16)

together with a restriction that β_2 be positive (i.e. IS exhibiting a convex pattern over time)

follows as an implication of Black's option pricing model. Different skewness dynamics through

a marketing year would be captured either by estimating higher polynomial or multiple-regime

models. In the multiple-regime models fit here, the restrictions listed above result in continuity of

predicted implied skewness at points of regime change, but smoothness at those points is not

imposed.

The points at which regimes changes, i.e. $\tilde{\tau}_1$ in models (17) and (18) and $\tilde{\tau}_1, \tilde{\tau}_2$ in model (19) are

also treated as parameters that need to be estimated, rather than being pre-determined.

Conditional on a particular choice of these parameters, the rest of the model can be estimated

using restricted least squares. For one-switch models, similar to Hansen (1999), denote the sum

of square errors for restricted least squares estimates conditional on a particular value of τ_1 as $SSE(\tau_1)$. The optimal point for the regime switching time is found as the minimizer of the conditional restricted sum of square errors:

$$\tilde{\tau}_1 = \arg \min_{\tau \in \Gamma} SSE(\tau) \tag{20}$$

For models with two switches, we can find the optimal switching points through a three-step minimization. First, conditional on particular values of τ_1, τ_2 we can find the optimal slope coefficients by restricted least squares estimation. Then, like above

$$\tilde{\tau}_2 \mid \tau_1 = \arg \min_{\tau_2 \in \Gamma(\tau_1)} SSE(\tau_1, \tau_2)$$
$$\tilde{\tau}_1 = \arg \min_{\tau_1 \in \Gamma_1} SSE(\tau_1, \tilde{\tau}_2 \mid \tau_1) \tag{21}$$

where $\Gamma_1 = \{\tau_1 : \tau_{MAX} - 20 \leq \tau_1 \leq 50\}$ and $\Gamma_2(\tau_1) = \{\tau_2 : \tau_1 - 30 \leq \tau_2 \leq 20\}$.

To implement this when estimating optimal points for regime switching, conditional on stipulating the number of regime switch points, simple grid search is used, and then the sum of squared errors (SSE) from the estimated restricted least squares are ranked. We stipulate that regime switching cannot be less than 20 days to expiry or closer than 20 days to the maximum time to maturity. For models with two switch dates, we also stipulate that the two switch dates cannot be less than 30 days apart. The model with the lowest SSE is chosen as best in its class.

Models are estimated separately for each contract month (March, May, July, September and December), using daily values of implied skewness for the period 1995-2009. To repeat, implied skewness is itself estimated using high-frequency data as described in the previous section. As such, implied skewness estimates become very unstable on a day-to-day basis for very high time to maturity horizons. One reason could be a lack of liquidity in options markets for options far

from expiry, and another the low number of years for which options with such long horizons have even been traded. To eliminate the effect of noise in the estimation of implied skewness for long time to maturity horizons, we truncate the maximum allowable time to maturity for each contract separately at the point where simple visual inspection indicates noise starts to dominate.

In selecting the optimal model specification among the five models listed, we have used the theory developed by Hansen (1996, 1999, 2000) and used in Cox, Hansen and Jimenez (2004). As Hansen (1996) explains, the problems of inference in the presence of nuisance parameters (i.e. regime switching times) is that they are not identified under the null hypothesis of no-regime change. If we fixed the regime switching-time to a particular value, we could perform a standard Wald test to see if parameters for intercept and slopes are equal for observations occurring before and after $\tilde{\tau}$ days to maturity. However, since we cannot restrict the possible threshold time a priori, as Hansen (1996) explains, the asymptotic distribution of standard tests are nonstandard and nonsimilar, which means that tabulation of critical values is impossible. The finite sample distribution of the Wald statistic under the null hypothesis is calculated by simulation and the null hypothesis is rejected if the test statistic is higher than the desired percentile of the simulated Wald statistic distribution under the null. Details of the bootstrapping method used in testing for the optimal model class are presented in section 5.

3.3. Inter-year variation of implied skewness

Finally, we turn to explaining the inter-year variations in implied skewness. As argued in the previous section, skewness will likely be impacted by weather once corn silking starts. Therefore, if we are to infer an impact of storage on skewness across many years, each with its own weather peculiarities, we should choose the time before the reproductive growth phase starts, i.e. no later than third week of June. If we were to choose skewness observed much earlier

than that, we would risk falling in the endogeneity trap. Before a marketing year is close to the end, consumption can react to changes in futures price, possibly even to changes in options premiums, thus increasing or decreasing carryout stocks. It would make little sense then to use expected ending stocks-to-use as a predetermined explanatory variable and implied skewness as a dependent variable. To avoid this problem, the expected ending stocks-to-use ratio of the previous marketing year, as reported in June edition of World Agricultural Supply and Demand Estimates (WASDE) report[2] is employed for explanatory variable for storage adequacy.

If the price elasticity of supply for corn is not zero, we would expect producers to react to tighter expected stocks and higher new crop prices with an increase in planted acreage, so acreage response is the second variable we need to include in the model. Specifically, we use the measure of change between intended plantings for a given year as reported in the USDA Prospective Plantings[3] report published at the end of March, and the actual acreage planted in the previous marketing year.

In addition to supply side covariates, we need to address possible asymmetries in uncertainty of demand. Domestically, corn is used as a livestock feed, an industrial sweetener and as an input in ethanol production. All three of these derived demand categories are likely impacted by macroeconomic shocks. Therefore, as a measure of demand uncertainty we use the June-to-June change in the national unemployment rate as published by the Bureau of Labor Statistics.

The final econometric model has the following form:

---

[2] WASDE is produced by World Agricultural Outlook Board, inter-agency body at United States Department of Agriculture. Historical WASDE reports can be accessed at
http://usda.mannlib.cornell.edu/MannUsda/viewDocumentInfo.do?documentID=1194

[3] Prospective Plantings is a government report produced annually by the National Agricultural Statistics Service, an agency of the United States Department of Agriculture. Historical Prospective Plantings reports can be accessed at
http://usda.mannlib.cornell.edu/MannUsda/viewDocumentInfo.do?documentID=1136

$$IS_t = \alpha + \beta_1 E_t \Delta A_T + \beta_2 E_t \left[S_T / D_t \right] + \beta_3 \Delta U_t \tag{22}$$

Where IS_t stands for implied skewness for a December contract of year t estimated as the

average of implied skewness for the 10 trading days following the June WASDE report. The

change in acreage planted is ΔA_T. Since in June we only observe intended plantings, this is

written as the expected change in acreage. Expected ending stocks-to-use is $E_t \left[S_T / D_t \right]$ and ΔU_t

is the June-to-June change in the U.S. unemployment rate. Theory predicts that all coefficients

except the constant should be negative. A stronger acreage response and higher carryout stocks

relative to demand imply more ability to buffer adverse weather shocks, and will thus reduce

skewness. Likewise, a more unstable macroeconomic environment will decrease demand for fuel

and possibly even for meat, thus reducing upward pressure on corn prices.

4. Data

Commodity futures for corn as well as options on futures are traded on the Chicago Mercantile

Exchange (formerly the Chicago Board of Trade). A dataset comprising all recorded

transactions, i.e. times and sales data (also known as "tick data") for both futures and options on

futures, for the period 1995 through 2009, was obtained. It includes data for both the regular and

electronic trading sessions. The total number of transactions exceeds 30 million, including 22

million observations on futures contract trades, and about 10 million trades in options contracts.

Options data were matched with the last preceding futures transaction. LIBOR interest rates were

obtained from British Bankers' Association, and represent the risk-free rate of return. Overnight,

1 and 2 weeks, and 1 through 12 months of maturity LIBOR rates for period the 1995 through

2009 were used to obtain the arbitrage-free option pricing formulas. In particular, each options

transaction was assigned the weighted average of interest rates with maturities closest to the

contract traded. To avoid serial correlation in residuals from estimating implied coefficients, the data frequency was reduced to not less than 15 minutes between transactions for the same options contract. This resulted in data sets of between 200 to 800 recorded transactions for a particular trading day for a total of around 1.1 million observations used in estimation. For each data point we separately estimate implied volatility using CRR binomial trees with 500 steps. Then, for each data point, the price of a European option using Black's formula is calculated using the same parameters (futures price, interest rate, time to maturity) as that recorded for the American option. In addition, volatility is set equal to the one implied for American options. These 'artificial' European options are then used in fitting parameters of GLD option pricing model for each trading day separately.

As stated in the previous section, the implied skewness used in the econometric analysis is calculated as a simple average over 10 business days following the June WASDE report. Due to the high incidence of limit-move days and days with high intraday price changes the year 2008 is excluded from the sample. Including 2008 would render the calculation of higher moments unreliable. Descriptive statistics of the variables used in econometric analysis are given in Table 1, and corn supply/demand balance sheets are in Table 2.

Figure 9 presents a scatter diagram of expected ending stocks-to-use vs. implied skewness. Note the inverse relationship between these variables and the beneficial impact of the acreage response. For example, in the summer of 1996, carryout stocks-to-use were only 4.03%, two standard deviations below the average for 1995-2009. However, skewness was below the mean, due to a 12.2% increase in expected acreage, which is 2.2 standard deviations above the average increase of 1.4%. Similarly, in 2007 carryout stocks were only 8.56% of demand, but a massive acreage increase of 15.5%, by far the largest in this sample, reduced the skewness below the

24

mean. It is instructive to look at 2006 as well. Although ending stocks were bountiful at 19.67% of demand, a reduction in acreage of 4.6% made for the third largest skewness in the sample.

5. Estimation procedure and results

5.1. Estimating parameters of GLD distribution and implied higher moments

As stated in section 3.1., for each contract, for each trading day, we separately estimate the parameters $\sigma, \lambda_3,$ and λ_4 in the GLD option pricing formula. In particular, we minimize the squared difference in option premiums calculated with the GLD formula, and prices of European options as implied by Black's model. To do so, we first need a starting value for the implied volatility of an option with a strike price closest to the underlying futures price. The starting values for the λ_3 and λ_4 parameters were chosen to correspond to the skewness and kurtosis of the terminal futures price as they would be under the restriction that the logarithm of the terminal price is normally distributed with variance equal to $\sigma^2 t$, where σ^2 is the square of the starting value for the implied sigma parameter. Excel Solver is used to run the minimization problem, utilizing a FORTRAN compiled library (.dll file) created by Corrado (2001) that estimates GLD European Call prices. A formula for the GLD European put option was then programmed in Visual Basic for Applications.

Estimated lambda parameters are employed to calculate implied skewness and kurtosis. GLD option prices seem to work rather well, with an average absolute pricing error about 3/8 of a cent per bushel, and a maximum pricing error usually reaching not more than 2 cents (this occurs for the least liquid and most away from the money options). While there may be issues regarding the robustness of implied parameters with respect to starting values, the implied parameters seem to

be rather stable from one day to the next. For December 2007 corn, for example, the skewness estimated between June 11 and June 25, 2007 varies between 1.15 and 1.26. For that year, the average absolute pricing error was 7/8 of a cent per bushel, with a maximum pricing error of 7.9 cents.

For all years in the sample, the implied skewness is 1.2 to 3 times higher than it would be if the logarithm of the terminal futures price was really expected to be normal. Implied kurtosis is 1.2 to 1.6 times higher than that predicted by Black's model. We thus see that deviations from Black's model are particularly pronounced in implied skewness.

5.2. Dynamics of intra-year implied skewness: results

The results of intra-year models for dynamics of implied skewness are presented in Table 3, and predicted implied skewness for each contract month is plotted in Figure 10. For all five contracts, the quadratic model improves fit dramatically over the linear model. To perform a formal test whether a model with one switching time and quadratic segments fits the data better than the quadratic model, we have used the bootstrapping procedure described by Cox, Hansen and Jimenez (2004). In doing this test, we shall refer to the quadratic model (16) as the restricted model and model (18) as the unrestricted model. These models are nested, i.e. model (16) is obtained by imposing restrictions $\alpha_1 = \alpha_2, \beta_1 = \beta_2, \gamma_1 = \gamma_2$. Under the null hypothesis that these restrictions hold the switching time τ_1 is not identified. To test the null hypothesis, we first make 2000 bootstrap samples using the fixed-regressors residual bootstrapping method. In particular, for each simulation values of implied skewness are calculated by adding a draw from the empirical distribution of residuals to predicted value of the dependent variable. Fitting is done using the estimated coefficients from the restricted model, in this case model (16). Then, for each

26

bootstrapped sample, parameters of the unrestricted model, including switching time, are calculated by the same method as before, i.e. combining a grid search and concentrated restricted least squares. A Wald statistic $W_n = n\dfrac{\left(SSE_0 - SSE_1\right)}{SSE_1}$ is then calculated for that particular replication, where n is the number of observations in the sample, SSE_0 is the sum of square errors in the restricted model (zero switching points) using bootstrapped data and SSE_1 is the sum of square errors of model (18) using bootstrapped data. The entire process is repeated 2000 times to obtain a finite sample distribution of the Wald statistic. The null hypothesis is rejected if the Wald statistic obtained using the original data is higher than the 95th percentile of the simulated distribution. We see from table 3 that model (16) is strongly rejected in favor of model (18) for all five contract months.

We also estimate a model with two regimes changes. For the May contract, the optimal first switching time solves to a corner solution, i.e. 20 days less than the maximum time-to-maturity used in estimation. We interpret this as evidence that for the May contract, a model with two regime switching times does not explain the data any better than models with one change in regime, and is in fact a misspecification, i.e. number of break points is stipulated to be higher than actually exist. For other contract months, the optimal switching time solves out to the interior of the allowable set of times, and we need to perform a formal test to investigate if models with two switching times are indeed better representations of the data. Bootstrapping is again employed. In particular, the null hypothesis now is that true model is model (18), and the unrestricted model is model (19). Model (18) can be obtained from model (19) by restricting it, such that coefficients satisfy: $\alpha_1 = \alpha_2, \beta_1 = \beta_2, \gamma_1 = \gamma_2$.

We again use fixed-regressors residual-bootstrap technique and add draws from the empirical

distribution of residuals obtained from model (18) to the implied skewness measures predicted

using the estimated coefficients of model (18). For each replication, a Wald statistic

$$W_n = n \frac{\left(SSE_1 - SSE_2\right)}{SSE_2}$$ is calculated, where SSE_1 is the sum of square errors obtained by

estimating model (18) on bootstrapped data, and SSE_2 is calculated by estimating the model with

two switching times on bootstrapped data. As before, the entire process is repeated 2000 times to

obtain a finite sample distribution of the Wald statistic. The null hypothesis is rejected if the

Wald statistic obtained using original data is higher than the 95[th] percentile of the simulated

distribution. We find that the Wald statistics obtained using the original data are low enough that

the null hypothesis cannot be rejected for any contract month, and p-values are exceptionally

large. In conclusion, statistical tests show that a model with 1 regime change is superior. To test

if model (18) explains the data any better than model (17) with two linear segments we can use

standard critical values in Wald test, as both models have the same number of regimes. We find

that the null is rejected for all contract months.

The next issue to investigate and explain concerns evaluated knot times and their confidence

intervals. Point estimates are found using the already explained estimation procedure. Residual-

based bootstrap is then used to obtain confidence intervals. For a particular contract month,

simulated data is created by adding draws from the empirical distribution of residuals to the

predicted implied skewness using the same model for which we evaluate confidence intervals of

the knot. The model is then re-estimated on simulated data, and a new optimal knot value is

noted. The procedure is repeated 2000 times, with the confidence interval obtained using the

2.5[th] and 97.5[th] quantile as the lower and upper bounds, respectively. For model (18) with

quadratic segments, we find that the confidence intervals for switching times are substantial for May and July contracts. A possible reason is that the first segment in the model is convex, and the second concave, creating a rather smooth transition. In such a setting, changing the knot value can be very easily compensated for by changes in the slopes parameters. For the September, December and March contracts, both segments are estimated with concave curves, and exhibit much tighter confidence intervals of the switching times. Results are presented in Table 3. As a robustness check, we also calculate asymptotic confidence intervals using a method developed by Hansen (2000) that involves inverting a likelihood ratio statistic. We find that our bootstrapping method matches closely the results obtained using asymptotics for all contract months except July. For that contract month, the curve for the likelihood ratio statistic is rather flat and close to the asymptotic critical value for time-to-maturity values included in the bootstrapped confidence interval. In that sense, we perhaps could say that the bootstrap produces more conservative estimates for the confidence intervals. Another likely reason for observed differences could be that we estimate our model with the additional restriction of continuity in predicted variable, whereas asymptotic distribution is developed for unrestricted least squares estimation.

In the model with 1 regime change and quadratic segments, optimal switching time for September contract is 69 days to maturity, and for December it is 160 days. It will help us to be able to map time-to-maturity measures to a particular date in a year. Option contract specifications state that last trading day is *"The last Friday preceding the first notice day of the corresponding corn futures contract month by at least two business days."* The first notice day is the first day of the delivery month. For simplicity, we approximate the last option trading day to be 25th of the month preceding the delivery month. Under such an approximation, regime

switching times for new-crop contracts correspond to June 18[th] for the September contract and

June 19[th] for the December contract. To test if regime switching times for these two contracts

really fall on the same calendar date, we perform a Wald test. This is a non-standard test and we

use residual-based bootstrapping to generate data under the null hypothesis that calendar

switching dates are the same, which is equivalent to restriction that $\tilde{\tau}^{DEC} = \tilde{\tau}^{SEP} + 90$. The null

hypothesis is not rejected, and p-value is 0.9995, with the original Wald statistic is higher than

only one out of 2000 Wald statistics simulated under the null hypothesis. The *Crop progress*

report[4] published in last week of June is normally the first such report to list corn silking

progress. These reports suggest that on average about 5% of the U.S. corn crop is silking by June

26[th]. Thus, dynamics of implied skewness for new-crop contracts appear to change right at the

start of the corn silking period.

For carry contracts we find quite a different dynamic. Regime changes for the March, May and

July contracts occur respectively at 130, 125 and 112 days to maturity. Suspecting that these

days account for similar patterns across contracts, we tested whether regime changes occur at

statistically significantly different time-to-maturity horizons. Similar to a previous hypothesis

test, we express that hypothesis as restrictions on switching times: $\tilde{\tau}^{MAR} = \tilde{\tau}^{MAY} = \tilde{\tau}^{JUN}$. The

optimal switching time under the null hypothesis (that the restriction holds) is 128 days to

maturity, and the null is not rejected. This common switch is manifested in Figure 8 as concave

and substantial decrease in implied skewness close to contract expiry. This likely reflects the

decline in overall uncertainty concerning terminal prices as maturity approaches. It is more

---

[4] Crop Progress report is a government report produced weekly from April through November of each year by the National Agricultural Statistics Service, an agency of the United States Department of Agriculture. Historical Prospective Plantings reports can be accessed at
 http://usda.mannlib.cornell.edu/MannUsda/viewDocumentInfo.do?documentID=1048

interesting to note that for carry contracts, the first segment (i.e. before the switch) is convex for the May and July contracts, but concave for March. This could reflect that fact that for March contract, corn growth-sensitive period falls in the middle of the March contract lifetime, while for May and July this growth period is at the beginning of the contract life.

Table 4. showcases the relative contribution of the corn growth period (silking through harvest, approximated by the dates June 20 to October 20) to skewness reduction during a contract's lifetime. For carry contracts the days spent in second regime at the end of the contract life are excluded. For example, for the July contract, the maximum time to maturity was 350 calendar days. The contract traded for 230 calendar days prior to entering the "finish-line" period, i.e. the last 120 days in which we find a strong reduction of skewness. Out of those 230 days, 113 days, or 49.1% of time, falls in the growth sensitive period. At the maximum time-to-maturity horizon predicted implied skewness is 1.419, and at switch time it is 1.218. Although the growth sensitive period constitutes only one half of that time, it accounts for 76.5% of the difference between the maximum time-to-maturity horizon and switch time implied skewness. For the March contract, we see a situation that is even more extreme – the sensitive growth period constitutes 59.5% of the pre-switch life, but accounts for 94.3% of the difference between skewness at maximum time-to-maturity and at the switch-time.

5.3. Intra-year variation in implied skewness: results

Results of the previous section further justify using implied skewness for December contract over 10 days after June WASDE report in investigating effect of expected stocks-to-use at the end of a marketing year (Aug. 31) on implied skewness. To test this, a simple linear regression is estimated for the period 1995-2009 using implied skewness as the dependent variable. The independent variables include a constant, the expected ending stocks-to-use, the expected

planned change in planted acreage and changes in the unemployment rate. Regression statistics are reported in Table 5. Due to very low degrees of freedom (10), we have to rely on t-table for critical values, and use a one-tail test for the stocks-to-use coefficient.

An 1 % increase in stocks-to-use reduces skewness by 0.015. This coefficient is statistically significant at the 95% confidence level. To put this number in perspective, the difference between the lowest and the highest ending stocks-to-use recorded in the sample reduces skewness from 1.47 to 1.24, which is 47% of the difference between the highest and the lowest recorded skewness in the sample. Coefficients for demand uncertainty and acreage response are also statistically significant and have the expected sign.

6. Conclusions and further research

An option pricing model based on a generalized lambda distribution provides a useful heuristic in thinking about determinants of the shape of terminal futures price conditional distributions. Results indicate that crop inventories and plant physiology play a significant role in determining the expected asymmetry of the terminal distribution. In particular, results reveal that implied skewness is much more persistent than implied by Black's model. In years with low implied volatility implied skewness remains much higher than would be the case under the lognormality restriction, and dynamics are dominated not by time to maturity, but by temporal patterns in the resolution of uncertainty regarding crop yields.

Further research will focus on extending this analysis to soybeans and wheat. The U.S. is a major world player in corn, with 55.6% of world exports. That is higher than 45.3% of world exports of soybeans, and much higher than 17.7% percent in wheat. Extending the analysis to other crops will identify the effect of trade and non-overlapping growing seasons in different countries on

the magnitude, inter-year differences and intra-year dynamics on implied higher moments of the terminal price distribution.

Thus far the literature has focused on evaluating the impacts of government reports on implied volatility coefficients. The model presented here allows us to extend this to higher moments and examine how reports (i.e., information) influence the entire distribution of prices, not just the second moment. For example, we could use weekly crop progress reports to explain inter-year differences in the evolution of skewness through the summer months.

In the absence of high frequency data, many researchers use end of day reported prices for futures and options to evaluate implied higher moments. By re-estimating this model using only end of day data it is possible to examine the amount of noise and possible direction of bias such an approach brings to estimates of implied higher moments.

What happens when storage is not available to partially absorb the shocks to supply? It would be interesting to use the GLD option pricing model to examine the evolution and determinants of higher moments of non-storable commodities. Further research is needed to examine the impact of durability of production factors for commodities that are themselves not storable.

Finally, impacts of market liquidity and trader composition on the levels and stability of implied higher moments is a promising new area for research. With careful design of the analysis, we may be able to find a way to separate the part of the option price that is due to implied terminal price distributions from additional premium influences incurred due to hedging pressure or lack of market liquidity.

References

Adam-Müller, A.F.A., and A. Panaretou. 2009. "Risk management with options and futures under liquidity risk." Journal of Futures Markets 29(4):297-318.

Berra, Y. 2010. *"The Yogi Book: I didn't really say everything I said"*, New York: Workman Publishing.

Black, F. 1976. "The pricing of commodity contracts." *Journal of Financial Economics* 3(1-2):167-179.

Black, F. 1976b. "Studies of Stock Price Volatility Changes." Proceedings of the 1976 Meetings of the American Statistical Association, pg. 177-181.

Buschena, D., and K. McNew. 2008. "Subsidized Options in a Thin Market: A Case Study of the Dairy Options Pilot Program Introduction." *Review of Agricultural Economics* 30(1):103-119.

Corrado, C.J. 2001. "Option pricing based on the generalized lambda distribution." *Journal of Futures Markets* 21(3):213-236.

Cox, D., Hansen, B.E. and Jimenez, E. 2004. "How responsive are private transfers to income?" *Journal of the Public Economics*, 88:2193-2219.

Cox, J.C., S.A. Ross, and M. Rubinstein. 1979. "Option pricing: A simplified approach." *Journal of Financial Economics* 7(3):229-263.

Deaton, A., & Laroque, G. 1992. "On the Behaviour of Commodity Prices." *The Review of Economic Studies*, 59(1):1-23.

Dumas, B., J. Fleming, and R.E. Whaley. 1998. "Implied Volatility Functions: Empirical Tests." *The Journal of Finance* 53(6):2059-2106.

Egelkraut, T. M., P. Garcia, and B.J. Sherrick. 2007. "The Term Structure of Implied Forward Volatility: Recovery and Informational Content in the Corn Options Market." *American Journal of Agricultural Economics*, 89(1):1-11.

Fackler, P.L., and R.P. King. 1990. "Calibration of Option-Based Probability Assessments in Agricultural Commodity Markets." *American Journal of Agricultural Economics* 72(1):73-83.

Fortenbery, T.R., and D.A. Sumner. 1993. "The effects of USDA reports in futures and options markets." *Journal of Futures Markets* 13(2):157-173.

Gardner, B. L. (1977). "Commodity Options for Agriculture." *American Journal of Agricultural Economics*, 59(5):986-992.

Geman, H. 2005. Commodities and commodity derivatives. West Sussex: John Wiley & Sons Ltd.

Glauber, J. W., and M. J. Miranda. 1989. "Subsidized Put Options as Alternatives to Price Supports." Proceedings of the NCR-134 Conference on Applied Commodity Price Analysis, Forecasting, and Market Risk Management. Chicago, IL. [http://www.farmdoc.uiuc.edu/nccc134].

Hansen, B.E. 1996. "Inference when a nuisance parameter is not identified under the null hypothesis" *Econometrica* 64:413-430.

Hansen, B.E. 1999. "Testing for Linearity" *Journal of Economic Surveys* 13, 551-576.

Hansen, B.E. 2000. "Sample splitting and threshold estimation," *Econometrica*, 68:575-603.

Hilliard, J.E., and J. Reis. 1998. "Valuation of Commodity Futures and Options Under Stochastic Convenience Yields, Interest Rates, and Jump Diffusions in the Spot." *The Journal of Financial and Quantitative Analysis* 33(1):61-86.

Isengildina-Massa, O., S.H. Irwin, D.L. Good, and J.K. Gomez. 2008. "Impact of WASDE reports on implied volatility in corn and soybean markets." *Agribusiness*, 24(4):473-490.

Jarrow, R., and A. Rudd. 1982. "Approximate option valuation for arbitrary stochastic processes." *Journal of Financial Economics* 10(3):347-369.

Ji, D., and W. Brorsen. 2009. "A relaxed lattice option pricing model: implied skewness and kurtosis." *Agricultural Finance Review* 69(3):268-283.

Kang, T., and B.W. Brorsen. 1995. "Conditional heteroskedasticity, asymmetry, and option pricing." *Journal of Futures Markets* 15(8):901-928.

Lence, S.H., Y. Sakong, and D.J. Hayes. 1994. "Multiperiod Production with Forward and Option Markets." *American Journal of Agricultural Economics* 76(2):286-295.

Lien, D., and K.P. Wong. 2002. "Delivery risk and the hedging role of options." *Journal of Futures Markets* 22(4):339-354.

Ramberg, J.S., and B.W. Schmeiser. 1974. "An approximate method for generating asymmetric random variables." *Commun. ACM* 17(2):78-82.

Ramberg, J.S., P.R. Tadikamalla, E.J. Dudewicz, and E.F. Mykytka. 1979. "A Probability Distribution and Its Uses in Fitting Data." *Technometrics* 21(2):201-214.

Rubinstein, M. 1994. "Implied Binomial Trees." *The Journal of Finance* 49(3):771-818.

Rubinstein, M. 1998. "Edgeworth Binomial Trees." *The Journal of Derivatives* 5(3):20-27.

Rui, X., and M.J. Miranda. 1995. "Commodity Storage and Interseasonal Price Dynamics." Paper presented at the Conference on Applied Commodity Price Analysis, Forecasting, and Market Risk Management., Chicago, IL.

Shaw, R. H., J. E. Newman, D. G. Baker, R. V. Crow, R. F. Dale, Rich Feltes et al. 1988. "Weather Stress in the Corn Crop." OSU current report - Oklahoma State University, Cooperative Extension Service, Vol. 2099.

Sherrick, B.J., P. Garcia, and V. Tirupattur. 1996. "Recovering probabilistic information from option markets: Tests of distributional assumptions." *Journal of Futures Markets* 16(5):545-560.

Thomsen, M.R., A. McKenzie, and G.J. Power. 2009. "Volatility Surface and Skewness in Live Cattle Futures Price Distributions with Application to North American BSE Announcements." Paper presented at 2009 AAEA & ACCI Joint Annual Meeting. Milwaukee, WI.

World Agricultural Outlook Board. 2010. "World Agricultural Supply and Demand Estimates." United States Department of Agriculture. Retrieved October 4, 2010 from http://usda.mannlib.cornell.edu/MannUsda/viewDocumentInfo.do?documentID=1194

National Agricultural Statistics Service. 2010. "Prospective Plantings." United States Department of Agriculture. Retrieved October 4, 2010 from http://usda.mannlib.cornell.edu/MannUsda/viewDocumentInfo.do?documentID=1194

Vercammen, J. 1995. "Hedging with Commodity Options When Price Distributions Are Skewed." *American Journal of Agricultural Economics* 77(4):935-945.

Williams, J.C., and B. Wright. 1991. *Storage and commodity markets*. Cambridge ; New York: Cambridge University Press.

Figure 1. Typical pattern for implied volatility coefficients for options on agricultural futures

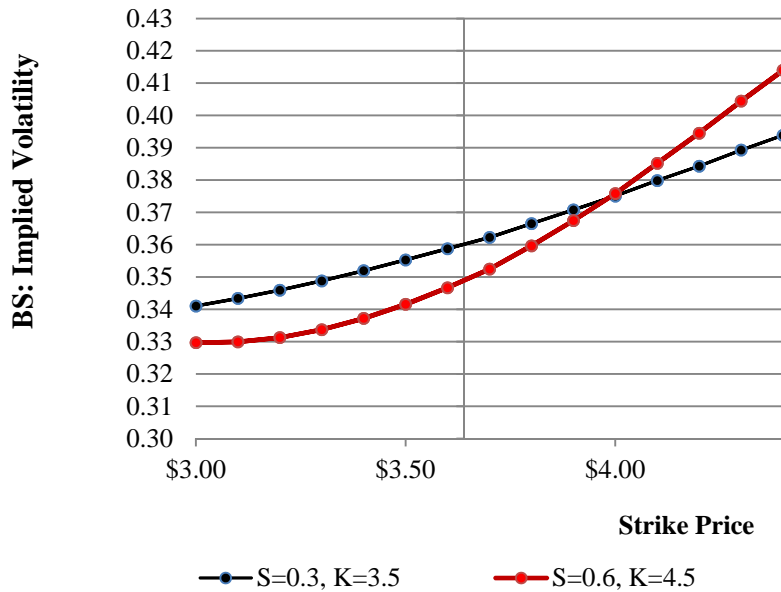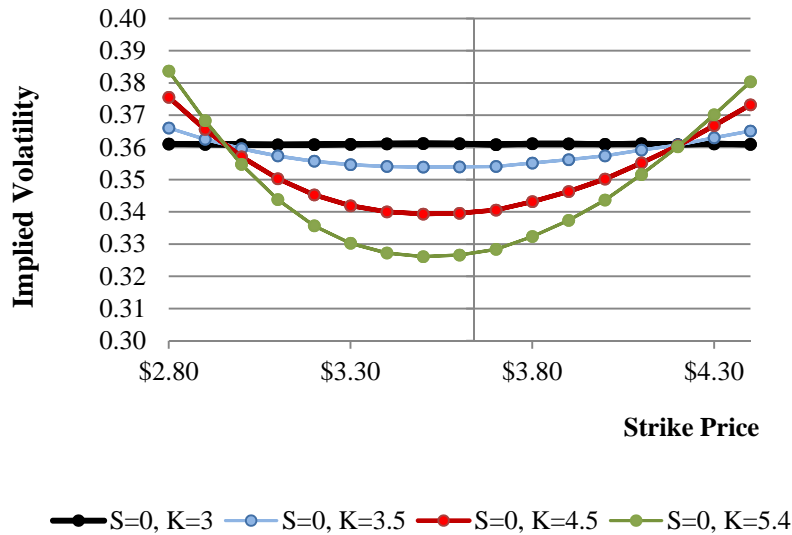Implied Volatility

Log(Strike/Underlying Futures Price)
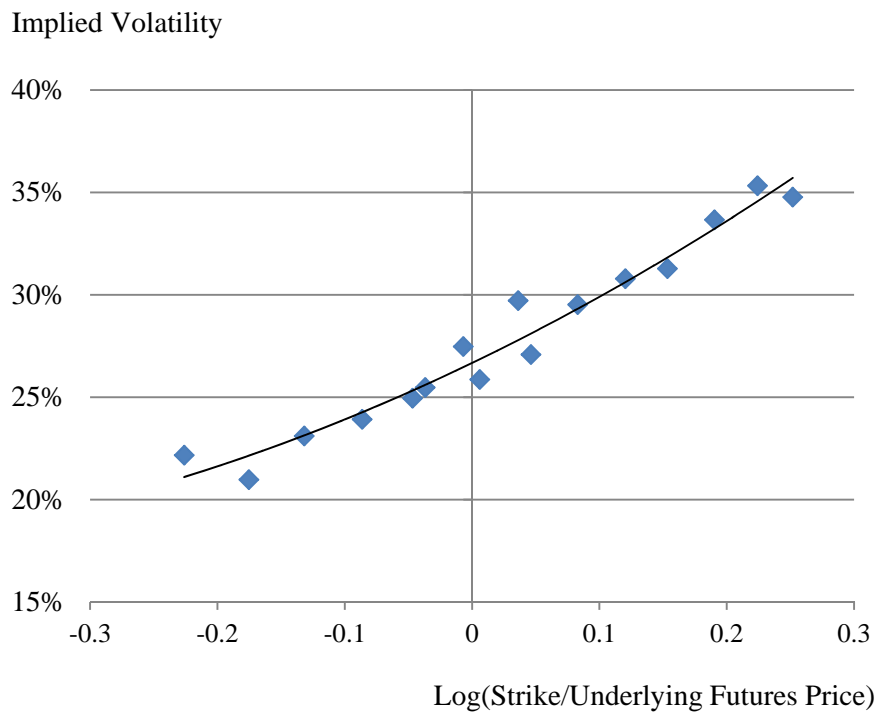
Notes: Implied volatility coefficients are estimated for options on the December 2006 corn futures contract, on 6/26/2006 using Cox, Ross and Rubinstein's binomial tree with 500 steps. The underlying futures price was $2.49/bu. Dots represent implied volatility coefficients for each strike, and the smooth line is a fitted quadratic trend.

Figure 2. Evolution of implied volatility curve for options on Dec' 04 and Dec '06 corn futures.
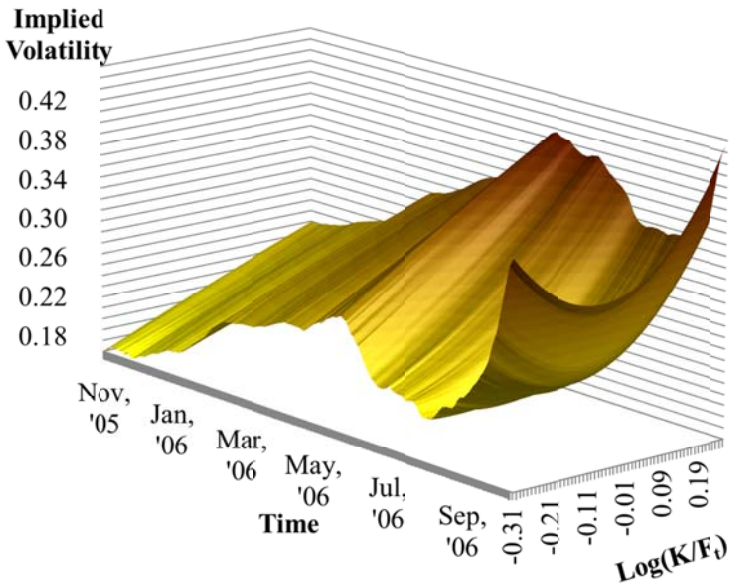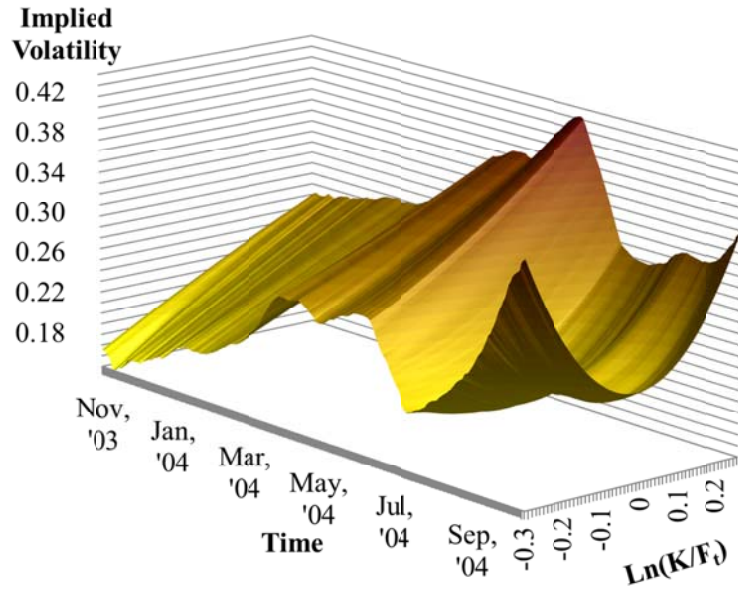
Notes: For each day, implied volatility is estimated for each traded option using 15 minute interval data. A quadratic trend curve is fitted to produce the implied volatility curve. A 30 day moving average is calculated to increase smoothness of the volatility surface and make it easier to see principal characteristic of the IV curve evolution. The Z-axis shows option moneyness calculated as the logarithm of the ratio between option strike (K) and underlying futures price F_t. When the option strike price is equal to the current futures price moneyness is zero.

Figure 3. Effects of Excess Kurtosis and Positive Skewness on Implied Volatility

Notes: S stands for skewness, and K for kurtosis of terminal futures log-prices. Option premiums are calculated via Rubinstein's Edgeworth binomial trees that allow for non-normal skewness and kurtosis, and implied volatility is inferred using Cox, Ross and Rubinstein's binomial tree which assumes normality in terminal futures prices. The black line in the above diagram with S=0 and K=3 corresponds to assumptions of Black's model, where implied volatility curve is flat across all strikes. Excess kurtosis (K>3) creates convex and nearly symmetric "smiles", and positive skewness produces an upward sloping implied volatility curve.

Figure 4. New crop price distributions conditional on storage adequacy
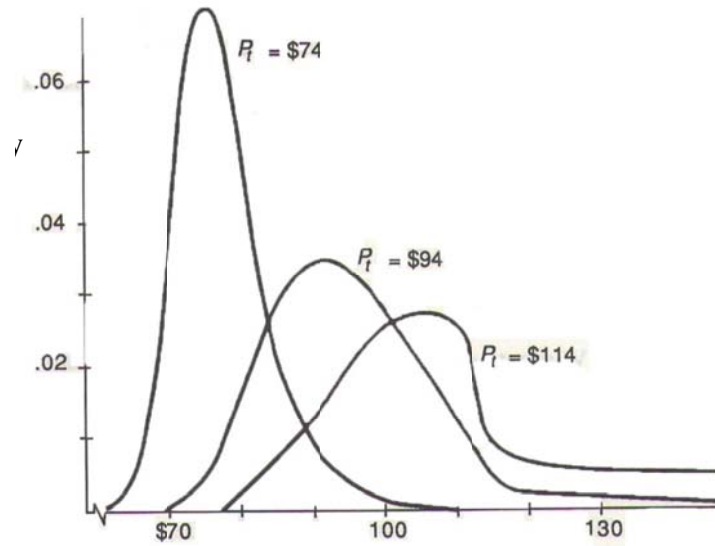
Note: Reproduced from Williams and Wright (1991). Conditional price distributions obtained from a rational expectations competitive equilibrium model with storage. The time frequency is one year, i.e. t+1 represents the next harvest. New crop price distributions are conditional on information known after old crop carryout stocks have been determined, but before weather shocks are revealed. Prices and quantity are standardized to make non-stochastic equilibrium at $100 and 100 units. Higher prices at time t reflect lower carryout stocks, and correspond to higher skewness of new crop price distribution.

Figure 5. Weather Stress in the Corn Crop
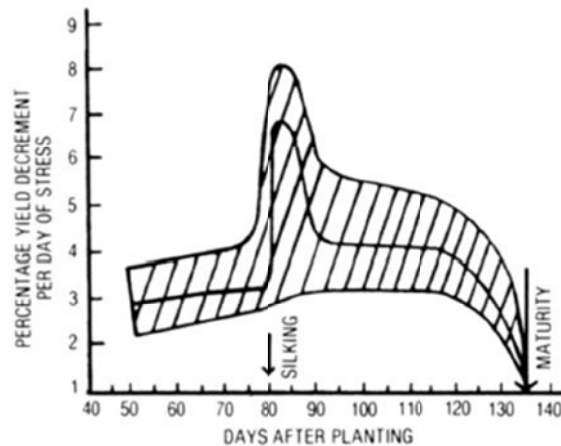
Note: Reproduced from Shaw et al. (1988). This figure shows the relationship between the age of the corn crop and the percentage yield reduction due to one day of moisture stress. Outer lines show boundaries of experimental results, while the middle line shows the average.

Figure 6. Monthly projected corn yield 1980-2008 - deviation from trend
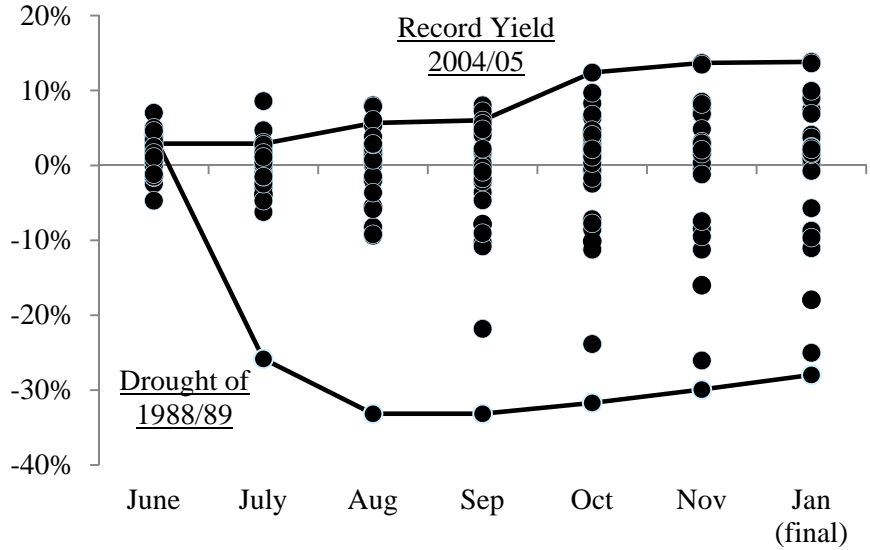
Note: For each year, trend yield was calculated as a simple linear regression over previous years, starting in 1960. Monthly projected yields were obtained from the WASDE report either directly or by calculations based on projected planted area and expected production size.

Figure 7. Monthly projected corn yield 1980-2008 - deviation from the final estimate (January)
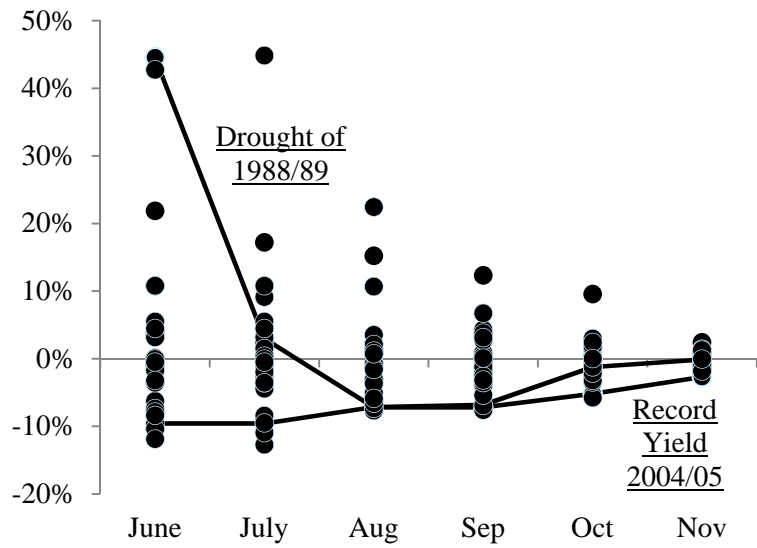
Note: For each month, projected yield was obtained from WASDE reports. Final estimates are taken from WASDE reports published in January.

Figure 8. Options on corn futures: Dynamics of Implied Skewness
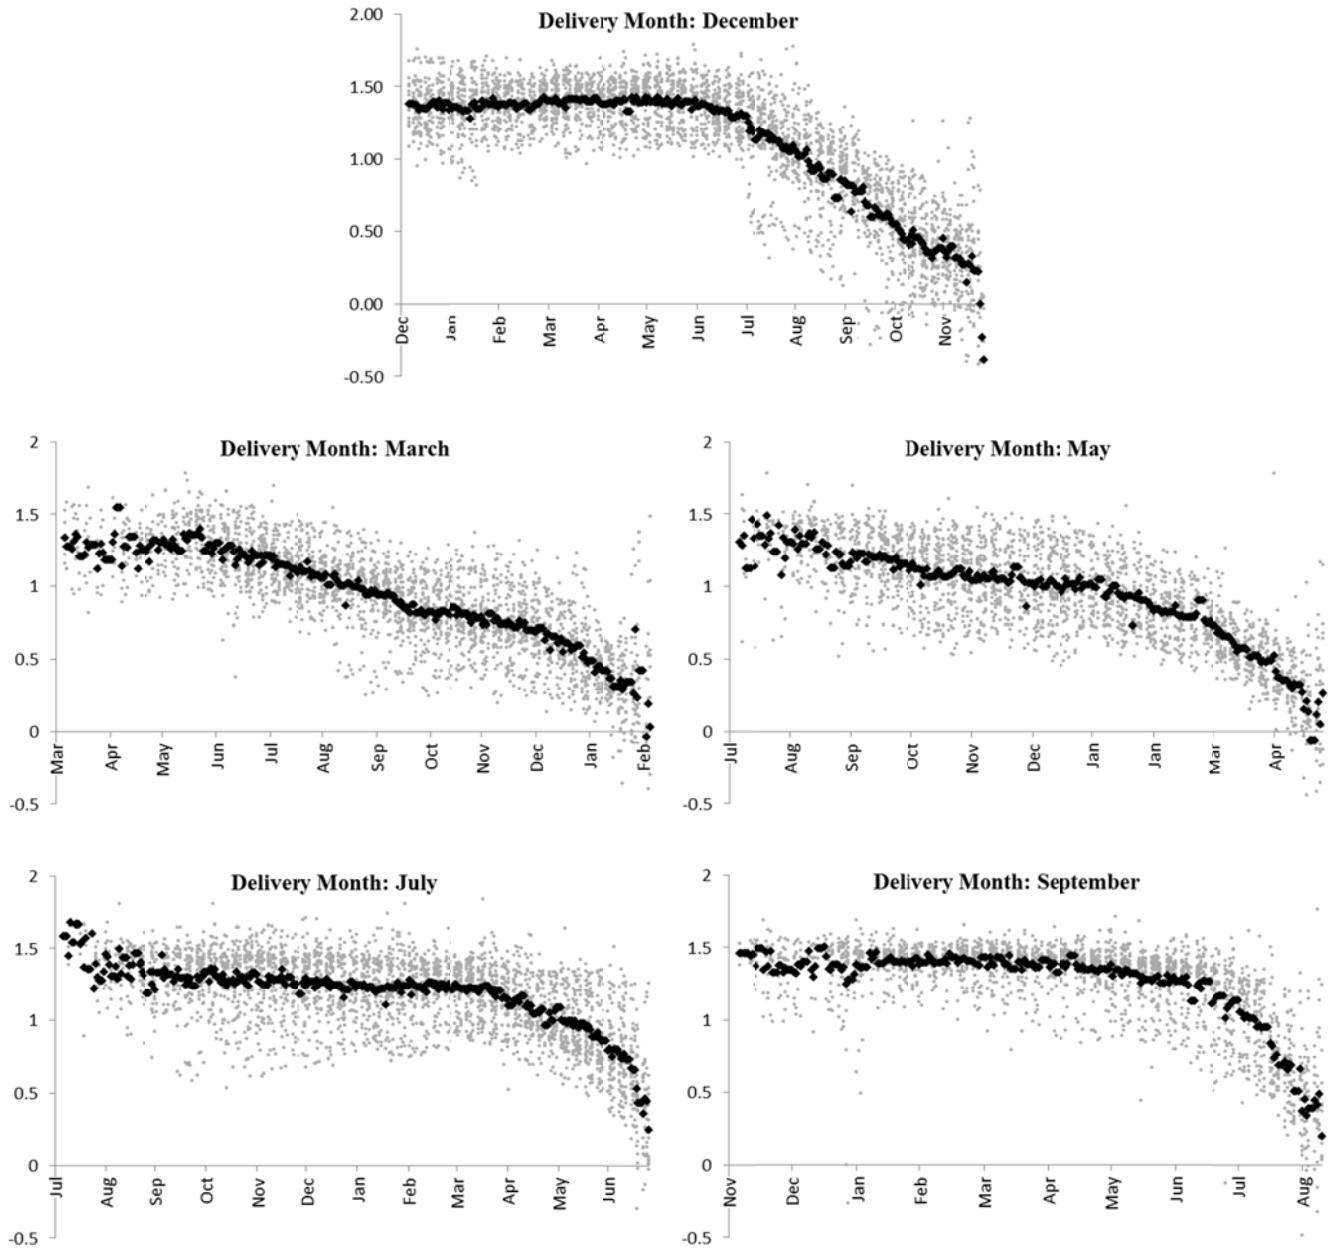
Note: Implied skewness is estimated for each contract and for each trading day separately using the generalized lambda distribution pricing model for options on commodity futures. Graphs show individual estimates over 1995-2009 for each contract and each time to maturity as gray scatter diagrams with the average for a particular time-to-maturity bolded.

Figure 9. Relationship between implied skewness and expected ending stocks-to-use
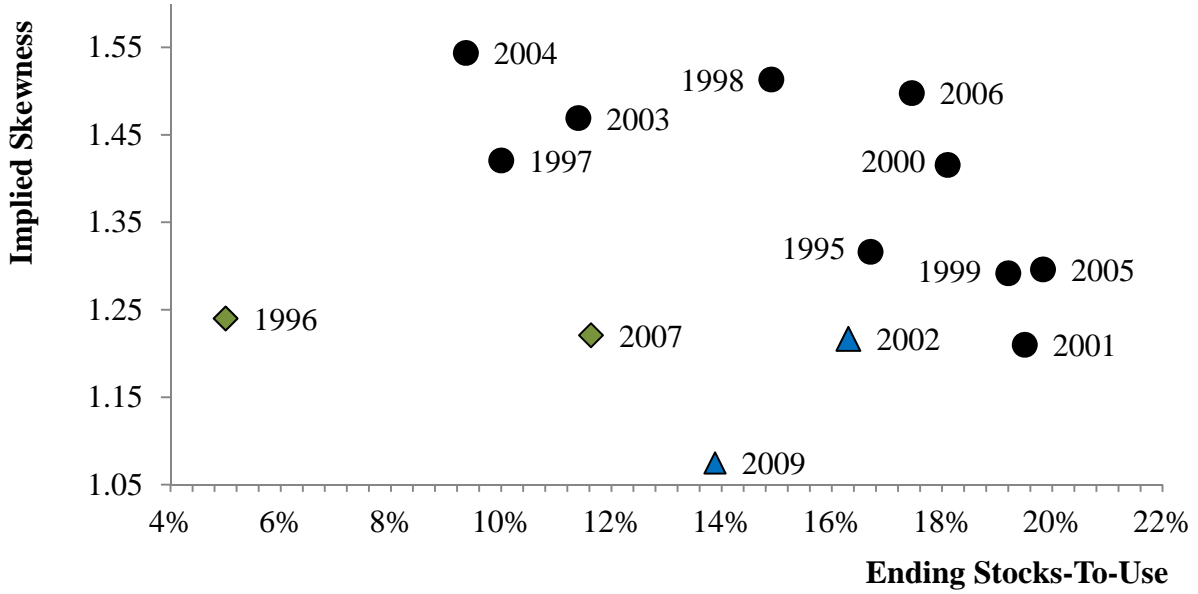
Note: Years with increase in intended cultivated acreage of 5 or more percent are drawn using green rhombs. Years with the June-to-June increases in the unemployment rate of 1 percent or more are drawn using blue triangles.

Figure 10. Predicted intra-year dynamics of implied skewness for options on corn futures contracts
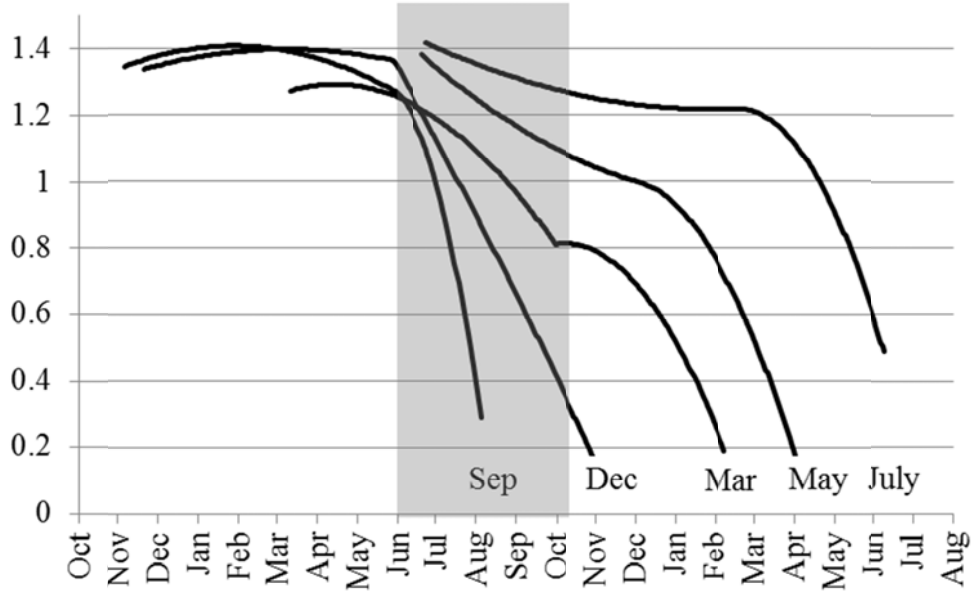
Note: Implied skewness is modeled with quadratic segments with one regime switch. The grayed area covers the corn growth sensitive period, i.e. silking through harvest time (approximately Jun 20 – Oct 20).

Table 1. Determinants of implied skewness: descriptive statistics

Descriptive Statistics

Variable	Mean	Standard Deviation	Min	Max
Implied Skewness	1.33	0.14	1.07	1.54
Ending Stocks-to-Use (%) WASDE June Projection	14.4	5.36	4.03	21.23
Intended Acreage Planted – Percentage Change	1.37	5.89	-4.84	15.48
Unemployment Rate Change	0.17	0.23	-0.7	4.00

Note: Implied skewness was calculated for December corn contracts as the average for implied parameters over 10 trading days following the June WASDE report. On average, 100-150 data points were used in estimating implied parameters for each trading day in the stated periods.

Table 2. Corn supply/demand balance sheet 1995-2009.

		Marketing year	1995	1996	1997	1998	1999	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009
SUPPLY	Exp.	Exp. acres planted	73.3	79.0	81.4	80.8	78.2	77.9	76.7	78.0	79.0	79.0	81.4	78.0	90.5	86.0	85.0
		Exp. acreage change	-7.4%	11.0%	2.4%	0.7%	-2.5%	0.6%	-3.5%	2.9%	-0.1%	0.4%	0.6%	-4.6%	15.6%	-8.1%	-1.2%
		Exp. yield	119.7	126.0	131.0	129.6	131.8	137.0	137.0	135.8	139.7	145.0	148.0	149.0	150.3	148.9	153.4
	Realized	Acres planted	71.2	79.5	80.2	80.2	77.4	79.5	75.8	79.1	78.7	80.9	81.8	78.3	93.6	86.0	86.5
		Acres harvested	65.0	73.1	72.7	72.6	70.5	72.4	68.8	69.3	70.9	73.6	75.1	70.6	86.5	78.6	79.6
		%Harvested	91.3%	91.9%	90.6%	90.5%	91.1%	91.1%	90.8%	87.6%	90.1%	91.0%	91.8%	90.2%	92.4%	91.4%	92.0%
		Yield	113.5	127.1	127.0	134.4	133.8	137.1	138.2	130.2	142.2	160.4	147.9	149.1	151.1	153.9	164.7
		Production	7,374	9,293	9,366	9,761	9,437	9,968	9,507	9,008	10,114	11,807	11,112	10,535	13,074	12,101	13,110
		Beginning stocks	1,558	426	883	1,308	1,787	1,718	1,899	1,596	1,087	958	2,114	1,967	1,304	1,624	1,673
		Imports	16	13	9	19	15	7	10	14	14	11	9	12	20	15	8
		Total supply	8,948	9,732	10,258	11,088	11,239	11,693	11,416	10,618	11,215	12,776	13,235	12,514	14,398	13,740	14,792
DEMAND	Exp.	Exp. total demand	8,600	8,820	9,000	9,360	9,480	9,645	9,725	9,535	10,405	10,560	11,060	11,525	12,960	12,140	13,190
		Exp. ending stocks	347	909	1,259	1,727	1,759	2,048	1,621	1,084	806	2,215	2,176	987	1,433	1,600	1,603
		Exp. stocks-to-use	4.03%	10.3%	13.9%	18.4%	18.5%	21.2%	16.7%	11.4%	7.7%	20.9%	19.6%	8.5%	11.1%	13.2%	12.2%
	Realized	Feed & residual	4,696	5,360	5,505	5,472	5,664	5,838	5,877	5,558	5,798	6,162	6,141	5,598	5,938	5,205	5,159
		Food/Seed/Ind.	1,598	1,692	1,782	1,846	1,913	1,967	2,054	2,340	2,537	2,686	2,981	3,488	4,363	4,993	5.938
		Ethanol	N/A	N/A	N/A	N/A	N/A	N/A	N/A	996	1,168	1,323	1,603	2,117	3,026	3,677	4,568
		Exports	2,228	1,797	1,504	1,981	1,937	1,935	1,889	1,592	1,897	1,814	2,147	2,125	2,436	1,858	1,987
		Total demand	8,522	8,849	8,791	9,299	9,514	9,740	9,820	9,490	10,232	10,662	11,269	11,211	12,737	12,056	13,084
		Ending stocks	426	883	1,467	1,789	1,725	1,953	1,596	1,128	983	2,114	1,966	1,303	1,661	1,684	1,708
		Stocks-to-use	5.0%	10.0%	16.7%	19.2%	18.1%	20.1%	16.3%	11.9%	9.61%	19.8%	17.5%	11.6%	13.0%	14.0%	13.1%
		Avg. farm price	3.24	2.71	2.43	1.94	1.82	1.85	1.97	2.32	2.42	2.06	2.00	3.04	4.20	4.06	3.55

Note: Acres planted and harvested are measured in million acres, yield in bushels per acre, beginning and ending stocks, imports, exports and other demand categories are measured in million bushels. Average farm price measured in U.S. dollars per bushel. Corn marketing year starts on September 1 of the current calendar year, and ends on August 31 the following calendar year. Expected acres planted based on "Prospective Plantings" report published at the end of March preceding the marketing year. Expected total demand, ending stocks, and stocks-to-use are taken from June WASDE report. For example, marketing year 2001/02 (denoted in table simply as 2001) started on 09/01/2001, and ended on 08/31/2002. For that year, expected acres planted was published on 03/31/2001 and expected total demand, ending stocks and stocks-to-use were taken from WASDE report published in 06/12/2002. Variables used in econometric analysis are bolded.

Table 3. Models of intra-year skewness dynamics: regression results

	March	May	July	September	December
Maximum time to maturity	330	290	350	275	350
Number of observations	2736	2558	3058	2321	3448
1. Linear model with no regime change					
SSE	178.49	156.39	180.62	166.28	295.67
R^2	0.55	0.56	0.30	0.34	0.58
2. Quadratic model with no regime change					
SSE	171.32	135.61	159.60	114.36	163.98
R^2	0.57	0.62	0.38	0.55	0.77
3. Linear model with one change in regime (timing is estimated endogenously)					
SSE	170.94	127.50	143.83	98.90	158.80
R^2	0.57	0.65	0.45	0.61	0.77
Switch time	63	64	62	56	162
4. Quadratic model with one change in regime (timing is estimated endogenously)					
SSE	167.36	126.36	141.17	96.10	158.30
R^2	0.58	0.65	0.46	0.62	0.77
Switch time	130	125	112	69	160
Switch date	Oct, 19	Dec, 22	Mar, 6	Jun, 18	Jun, 19
Switch time 95% CI (boot.)	(119-140)	(86-167)	(80-149)	(49-94)	(151-169)
Switch time 95% CI (asy.)[§]	(113-144)	(86-151)	(70-95)	(51-95)	(142-183)
Wald test: (2) vs. (4)					
Critical val. (simul., 95%)	10.89	10.94	10.57	10.73	10.77
Wald-statistic	64.72	187.14	399.30	441.11	123.76
(p-value)	0.0000	0.0000	0.0000	0.0000	0.0000
Wald test: (3) vs. (4)					
Critical val: $\chi^2(0.95,2)$	5.99	5.99	5.99	5.99	5.99
Wald-statistic	58.56	23.06	57.65	67.64	10.98
(p-value)	<0.001	<0.001	<0.001	<0.001	0.004

Wald tests for knots:

Null: $\tilde{\tau}^{DEC} = \tilde{\tau}^{SEP} + 90$ — Wald: 4.16, Crit. val.: 24.74 (p-value: 0.001)

Null: $\tilde{\tau}^{MAR} = \tilde{\tau}^{MAY} = \tilde{\tau}^{JUL}$ — Wald: 1.35, Crit. val.: 6.83 (p-value: 0.662)

	March	May	July	September	December
5. Quadratic model with two changes in regime (timing is estimated endogenously)					
SSE	167.28	126.19	140.96	95.59	158.16
R^2	0.58	0.65	0.46	0.62	0.77
Switch time 1	120	268[†]	270	220	270
Switch time 2	75	137	110	77	159
Wald test: (4) vs. (5)					
Critcal val. (simul., 95%)	13.82		13.79	11.57	12.15
Wald-statistic	1.30	N/A	4.55	1.47	3.05
p-value	0.9265		0.9094	0.9690	0.9575

[†] Model estimates to a corner solution, and is therefore treated as misspecified.

[§] Estimated using asymptotic likelihood ratio test developed by Hansen (2000), which does not impose restriction of continuity in predicted variable.

Table 4. Relative contribution of corn growth-sensitive period to reduction in implied skewness

Contract		
March	% of contract life	% skewness reduction
non-growth period	40.5%	5.7%
in growth period	59.5%	94.3%
May	% of contract life	% skewness reduction
non-growth period	39.1%	19.6%
in growth period	69.1%	80.4%
July	% of contract life	% skewness reduction
non-growth period	50.9%	23.5%
in growth period	49.1%	76.5%
September	% of contract life	% skewness reduction
non-growth period	76.0%	7.5%
in growth period	24.0%	92.5%
December	% of contract life	% skewness reduction
non-growth period	64.6%	33.47%
in growth period	35.4%	76.53%

Note: For carry contracts (March, May, and July), percentages reported refer to contract life before a regime switch, i.e. excluding the last four months of contracts life. For new-crop contracts (September and December), percentage reported are over the entire contract life.

Table 5. Determinants of implied skewness: regression results

Explanatory variables	Dependent Variable: GLD Implied Skewness
Constant	1.55 (0.10)
Ending Stocks-to-Use (%)	-1.28 (0.64)
Intended Acreage Planted – Percentage Change	-1.52 (0.58)
Unemployment Percentage Change	-0.07 0.02
Degrees of Freedom	10
Mean Root Square Error	0.075
R^2	0.66

Note: The critical t-statistic for 10 d.f. at 95% confidence is 1.81 for one-tail tests and 2.22 for two-tail tests. All coefficients are statistically significant at the 95% confidence level (Ending stocks-to-use coefficient is significant at the 95% using one-tailed test, or 90% using two-tailed test).