Journal of Agricultural and Resource Economics 36(1):177–191

Modeling Stochastic Crop Yield Expectations with a Limiting Beta Distribution

David A. Hennessy

The use of plausible stochastic price processes in price risk analysis has allowed advances not seen in crop yield risk analysis. This study develops a stochastic process for yield modeling and risk management. The Pólya urn process is an internally consistent dynamic representation of yield expectations over a growing season that accommodates agronomic events such as growing degree days. The limiting distribution is the commonly used beta distribution. Binomial tree analysis of the process allows us to explore hedging decisions and crop valuation. The method is empirically flexible to accommodate alternative assumptions on the growing environment, such as intra-season input decisions.

Key Words: crop abandonment, crop insurance, derivative analysis, growing degree days, Pólya's urn, stochastic process

Introduction

Demand for information on crop yield expectations arises from producers who seek to understand and manage yield risks and from consumers concerned about future stock availabilities. To meet these needs, the U.S. federal government issues regular production forecasts, which are based on grower surveys and include official estimates of planting intentions and planted acres and yield forecasts through the growing season. State extension agents and private firms also develop estimates. This study explores a novel, flexible, and simple approach to organizing and using these forecasts. A crop yield expectation stochastic process is designed that represents how harvest yield expectations change during the course of a crop year. This approach can be viewed as an application of real options methods and is similar in spirit to how stock, commodity, and other financial prices are typically modeled. Our key innovation is to identify a reasonable yield expectation stochastic process, because processes commonly used in price analysis are not suitable.

Improvements in the modeling of crop yields are likely to promote public policy formation for two reasons. The first reason has to do with the extensive government intervention in crop insurance markets, especially in the United States. The reasons for public intervention have been discussed elsewhere (see Innes, 2003). Public assistance has cost taxpayers about \$4.25 billion per year over 2004–2010 (Shields, 2010). Private crop insurance companies sell and service policies through their agents, but the USDA's Risk Management Agency designs insurance products and establishes or approves premium rates. Many have remarked on geographic imbalances in insurance rates (e.g., Babcock, 2008) while crop insurance markets are also fraught with information deficiencies (Smith and Goodwin, 1996; Makki and Somwaru, 2001). If yield insurance rate-setting is to continue as a publicly provided service,

---

David A. Hennessy is professor, Department of Economics and Center for Agricultural and Rural Development, Iowa State University. Two anonymous reviewers and the *JARE* editorial staff provided very helpful comments. The author also thanks seminar participants at Iowa State and North Carolina State Universities.
Review coordinated by Gary W. Brester.

then it is unlikely there will be a strong private-sector motive to engage in research to improve rate-setting methodologies. As a result, most research has been conducted by academics (see, e.g., Just and Weninger, 1999; Atwood, Shaik, and Watts, 2003; Ramírez and McDonald, 2006; Hennessy, 2009; Harri et al., 2009).

The second reason that crop yield modeling should be of interest for public policy relates to international economic vulnerability and development assistance. Crop agriculture supports large vulnerable populations in Africa and South America. Crop yield risks can be large as a result of extreme weather conditions and pest hazards. Financial markets to cope with these risks are generally weak, leaving growers to choose low-productivity technologies, degrade soil resources, or use other costly risk management strategies (Dercon, 2004; Skees, Hartell, and Murphy, 2007). The World Bank, the Gates Foundation, and others have sought to strengthen risk management opportunities through assistance in the design, rate-setting, and operation of crop insurance markets.

To meet these challenges, we develop a plausible expected yield stochastic process of a maturing crop. In addition, we suggest how the model can be used to communicate about and manage yield risk and to make intra-season production decisions. We show that the model satisfies a set of properties that are desirable, if not essential, for the yield expectation stochastic process to possess. For example the process supports a yield distribution with bounded (non-infinite) positive realizations.[1] It is also internally consistent in the following sense. If crop prospects are assessed at several points (e.g., weekly) during the growing season, the present-time expectation of next week's expectation of harvest yield should equal the present-time expectation of harvest yield.

The model also possesses the property that when expected yield is extreme (high or low), then today's belief about the variability of next week's expected yield should be very low. In other words, when comparing a crop that has fared poorly during the early season with one that has had a moderate start, the expectation for the poor crop should be less sensitive to new information. The poor start may well lower a plant's capacity to take advantage of better late season weather, perhaps because of shallow early-season root penetration. The same should be true for a promising crop compared with one having moderate prospects. The crop with the more promising growth history should have more resources available, either stored internally or in the soil, to overcome future weather setbacks (Bloom, Chapin, and Mooney, 1985).

One further requirement is that expectations should tighten as harvest approaches—that is, the variance of next period expectations should decline as harvest nears. This is especially true when early conditions have been extreme, either good or bad. In late August there remains little opportunity for a string of good weather to improve a crop that has earlier experienced poor weather. Similarly, if good early-season weather has accelerated maturity, then the crop may be harvested early, which obviates vulnerability to unseasonable weather near the normal harvest date.

Our expected yield stochastic process is developed from the Pólya and Eggenberger urn model (Mahmoud, 2009). The stochastic process converges to the beta distribution (Mahmoud, p. 53), which is a popular yield distribution model (Nelson and Preckel, 1989; Borges and Thurman, 1994; Babcock and Hennessy, 1996; Goodwin, 2009; Claassen and

[1] Agronomists have long worked with the idea of potential yield (i.e., maximum yield for a cultivar under ideal fertility, cultivation, and weather; see Araus et al., 2008). Tollenaar and Lee (2002) report a variety of estimates for maize, including one of about 400 bushels per acre for the U.S. Corn Belt. Acknowledging bounded support is important. Distributions with tails, however small in probability weighting, can allow reasonable mean, variance, and skewness statistics, but such tails can distort expected crop insurance indemnity payouts.

Just, 2011). The distribution has finite support and accommodates both positive and negative skewness.

Our approach is flexible both as a model for abstract analysis and as an empirical model. The ability to use a reasonable stochastic model should allow analysts and commercial risk managers to better quantify, communicate, and characterize risk attributes. The model's flexibility affords risk managers the opportunity to overcome any deficiencies that time and use may reveal about our proposed basic model. Stated differently, if it becomes clear that the basic model fails to capture some important feature of yield distributions, then opportunities exist to tweak the model. We will identify and comment on some opportunities to better calibrate the model were the need to arise. The method's flexibility should also prove useful in production analysis. Many researchers have studied intra-season information use in crop management (Antle, 1983; Antle and Hatchett, 1986; Antle, Capalbo, and Crissman, 1994; Saphores, 2000). We will show how to adapt our basic model to address these production problems.

Marcus and Modest (1984) posit an alternative approach in which expected yields follow a geometric Brownian motion over the growing season. This is convenient because the Black and Scholes valuation approach can then be employed to arrive at production choices that maximize crop asset values. However, this class of model necessarily uses lognormal yield distributions that exhibit positive skewness and a slowly decaying right tail. Neither property is consistent with observed data regarding yield distributions, especially in the prime crop growing areas of the United States (Harri et al., 2009).

After presenting and explaining our model of how expected yield evolves over the growing season, some of its properties are demonstrated. We show how the model can be viewed as a binomial tree, i.e., where each piece of additional information allows for one of two outcomes regarding yield expectations. The binomial tree feature provides a way to implement the model to, e.g., characterize yield sensitivities to weather and to better understand how intra-season production decisions shape the yield distribution. When weather is assumed to be the origin of uncertainty, binomial tree analysis also allows for dynamic hedging against a weather derivative contract. We use this observation to value a crop and also to better understand the hedging of yield risks. The paper concludes with some observations on model adaptability, and on other possible applications of the method.

The Model

Our model assumes a growing year with $T + 1$ periods. The first period begins at time $t = 0$, the planting date, and continues to $t = T$, the harvest date. The step length could be a day, week, or month, depending on when new information becomes available, the desired level of accuracy, and model implementation costs. We are interested in yield expectations at each time point $t \in \{0, 1, \ldots, T\}$. With Y_T as the measure of actual harvest yield, Ω_t is the information set available at time t. The expected harvest yield given information available at t is written as $\mu_t = E[Y_T | \Omega_t]$.

If we assume that yield realizations are nonnegative and finite, then it can be assumed without further loss of generality that the yield distribution has support only on the interval $[0, 1]$. To demonstrate this, suppose that lower case y_T is true yield. Then appropriate choices of a and b for a linear location-and-scale transformation $Y_T = (y_T - a)/b$ ensure that $Y_T \in [0, 1]$. Thus, if we model the distribution of $Y_T \in [0, 1]$, we can translate the distribution to model $y_T \equiv a + bY_T \in [a, a + b]$.

It is also assumed that yield expectations have a logistic form. This specification has been widely used to model plant growth processes (Carrasco-Tauber and Moffitt, 1992; Tschirhart, 2000). The assumption allows for expected yields to possess standard regularity properties including monotonicity, concavity, and bounded output (Chambers, 1988, p. 9). Specifically, we define expected yield at time $t = 0$ as:

(1) $$\mu_0 = \frac{g(\mathbf{x})}{f(\mathbf{x}) + g(\mathbf{x})},$$

where input vector $\mathbf{x} = (x_1, \ldots, x_N) \in \overline{\mathbb{R}}_+^N$, the positive real numbers in dimension N, is a vector of input choices, while $f(\mathbf{x}): \overline{\mathbb{R}}_+^N \to \overline{\mathbb{R}}_+$ and $g(\mathbf{x}): \overline{\mathbb{R}}_+^N \to \overline{\mathbb{R}}_+$ are continuously differentiable functions.

As distinct from Antle's (1983) sequential decision model, our basic model assumes production choices are made only at planting time, $t = 0$. It is convenient to write $h(\mathbf{x}) = g(\mathbf{x})/f(\mathbf{x})$ so that $\mu_0 = h(\mathbf{x})/[1 + h(\mathbf{x})]$. Expected yield at time $t = 0$ is $\mu_0 = E[Y_T \mid \Omega_0]$. Expected yield is increasing in each input choice x_i, $i \in \{1, 2, \ldots, N\}$ if and only if $dh(\mathbf{x})/dx_i \geq 0$, while expected yield at time $t = 0$ is concave in any given argument whenever $d^2 h(\mathbf{x})/dx_i^2 \leq 0$. Writing the optimal input choices $\mathbf{x} = \mathbf{x}^*$ as given, we abbreviate $f^* = f(\mathbf{x}^*)$ and $g^* = g(\mathbf{x}^*)$.

The model is one of information-conditioned updating of yield expectations as relevant events and their yield consequences are processed. At $t = 1$, new information arrives and expected yield evolves as follows (Durrett, 1996, p. 241; Mahmoud, 2009, p. 53):

(2) $$\mu_1 = \begin{cases} \mu_1^+ \equiv \dfrac{g^* + c}{f^* + g^* + c} \equiv \dfrac{\mu_0 + \dfrac{c}{f^* + g^*}}{1 + \dfrac{c}{f^* + g^*}} & \text{with prob. } \mu_0, \\[2ex] \mu_1^- \equiv \dfrac{g^*}{f^* + g^* + c} \equiv \dfrac{\mu_0}{1 + \dfrac{c}{f^* + g^*}} & \text{with prob. } 1 - \mu_0, \end{cases}$$

for $c > 0$. Here $c > 0$ is a parameter that recognizes good weather over the first growing period. It can be viewed as the benefit from good weather in that $\partial\mu_1^+/\partial c = f^*/(f^* + g^* + c)^2 > 0 > -g^*/(f^* + g^* + c)^2 = \partial\mu_1^-/\partial c$. Parameter c may also be viewed as a measure of yield sensitivity to weather outcomes in that $\partial(\mu_1^+ - \mu_1^-)/\partial c = (f^* + g^*)/(f^* + g^* + c)^2 > 0$.

The general expression over t time periods is defined as:

(3) $$\mu_t = \begin{cases} \mu_t^+ \equiv \dfrac{\mu_{t-1} + m(f^*, g^*, t, c)}{1 + m(f^*, g^*, t, c)} & \text{with prob. } \mu_{t-1}, \\[2ex] \mu_t^- \equiv \dfrac{\mu_{t-1}}{1 + m(f^*, g^*, t, c)} & \text{with prob. } 1 - \mu_{t-1}; \end{cases}$$

$$m(f^*, g^*, t, c) = \frac{c}{f^* + g^* + (t-1)c},$$

where equation (2) arises when $t = 1$.[2] Expression (3) provides the expected yield stochastic process that we posit over $t \in \{0, 1, \ldots, T\}$. Notice that $1 > \mu_t^+ > \mu_{t-1} > \mu_t^- > 0$. Note also that c is time independent, while we have not precluded $T \to \infty$. Consequently, in order to ensure the process has bounded support, it is necessary for the favorable and unfavorable outcomes to depend on the prior probability, μ_{t-1}. By contrast, the stochastic processes most commonly encountered in economics—Brownian motion and geometric Brownian motion—do not even have bounded unconditional variance.[3]

The change in the random variable's value is state dependent, also unlike the stochastic processes most commonly encountered in economics. We believe state dependence to be reasonable because the change in expected yield should depend on events over the whole growing season. For example, rainfall early in the season provides a buffer against drought later on. Perhaps surprisingly, the process for μ_t possesses the Markov property. This requires that the distribution of μ_t given time $t-1$ information can be conditioned on μ_{t-1} alone among historical process realizations $\{\mu_0, \mu_1, \ldots, \mu_{t-1}\}$, i.e., $\text{Prob}(\mu_{t+1} \mid \mu_t, \mu_{t-1}, \ldots, \mu_0) = \text{Prob}(\mu_{t+1} \mid \mu_t)$.

Equation (3) illustrates that the Markov property applies in a nonhomogeneous manner.[4] Specifically, the distribution of μ_t at time $t-1$ depends on the value of $t-1$ as well as on the value of μ_{t-1}. Stating that the expectation at this time point is $\hat{\mu}$ is not enough to establish the next time point's conditioned distribution on the expectation. One also needs to identify that $\hat{\mu}$ was realized at time t because the distribution varies with t as well as with the expected value at t.

Martingale Property

Notice from (2) that $E[\mu_1 \mid \Omega_0] = \mu_1^+\mu_0 + \mu_1^-(1-\mu_0) = \mu_0$. In general form, (3) allows us to establish that $E[\mu_t \mid \Omega_{t-1}] = \mu_t^+\mu_{t-1} + \mu_t^-(1-\mu_{t-1}) = \mu_{t-1}$. Thus, μ_t satisfies the martingale property that the time $t-i$ expected value of the time t expectation of yield equals the time $t-i$ expectation of yield for all $i \in \{0, 1, \ldots, t\}$. The time $t-i$ expectation is consistent with later state expectations, a necessary condition for the rational formation of expectations.[5]

Binomial Tree Attribute

Viewing the stochastic process as a binomial tree is useful for a variety of reasons. In particular, it provides a way to use the process in Monte Carlo simulations where intra-season events are of interest. Binomial trees also allow for the process to be used as an ingredient in a dynamic programming model. To explain the binomial tree attribute, consider what equation (3) implies for the distribution of forecasts two steps after planting. Substitute in for μ_1 to obtain:

---

[2] Formally, the Pólya urn algorithm has an initial state of g green balls and f white balls. A ball is drawn randomly. If it is green, then c more green balls are added. If it is white, then c more white balls are added. At time $t-1$, a total of $(t-1)c$ new balls have been added to the urn. This explains why $(t-1)c$ appears in $m(\cdot)$ for the probability transition equations in (3). The probability of drawing a green ball at time t depends on the urn's composition at time $t-1$. This explains the role of μ_{t-1} in expression (3).

[3] The Ornstein-Uhlenbeck process and related mean-reverting processes are commonly used to preclude price drift over time toward implausibly extreme values and infinite unconditional variance (Hull, 2009, p. 751). Infinite unconditional variance is implausible because high prices will entice additional resources into producing the commodity, while substitutes will be found if the commodity's price does not eventually fall. These mean-reverting processes do allow for high price events, but with low probability, and so may be appropriate to account for commodity price spike events.

[4] On the Markov property, including instances of nonhomogeneity, see Bhattacharya and Waymire (1990, pp. 109–117, and especially pp. 116–117) or Durrett (1996, chapter 5).

[5] See Durrett (1996, chapter 4) for extensive detail on martingales.

$$(4)\qquad \mu_2 = \begin{cases} \dfrac{g^*}{f^*+g^*+2c} & \text{with prob. } \dfrac{f^*(f^*+c)}{(f^*+g^*+c)(f^*+g^*)}, \\ \dfrac{g^*+c}{f^*+g^*+2c} & \text{with prob. } \dfrac{2f^*g^*}{(f^*+g^*+c)(f^*+g^*)}, \\ \dfrac{g^*+2c}{f^*+g^*+2c} & \text{with prob. } \dfrac{(g^*+c)g^*}{(f^*+g^*+c)(f^*+g^*)}. \end{cases}$$

Due to the Markov property, yield expectations are the same under the good-then-bad weather outcome in expression (4) as under the bad-then-good outcome. This exchangeability attribute (Bhattacharya and Waymire, 1990, p. 116) simplifies the stochastic structure.[6]

The distribution given in expression (4) can be viewed as a two-period binomial tree (Hull, 2009).[7] Binomial trees are multi-step models of all possible paths of a random variable over time. Two outcomes are allowed at each step, and probabilities are provided for outcomes along each possible path. Thus, the probability distribution of the random variable at the time of interest can be established. As depicted in figure 1, $t = 2$ yield outcomes are right-most, and probabilities over each of the two time intervals are given under the appropriate arrow. There are three nodes established at points where information becomes available. These are (*a*) time point 0, NODE 0; (*b*) time point 1, NODE 1A that occurs when favorable information was obtained at node 0; and (*c*) time point 1, NODE 1B that occurs when unfavorable information was obtained at node 0. If the process proceeds T steps, then there are $T + 1$ distinct terminal outcomes in the interval $[0, 1]$.

The realization possibilities for time t expected yield are given by:[8]

$$(5)\qquad n(f^*, g^*, t, i, c) = \frac{g^*+ic}{f^*+g^*+tc}, \qquad i \in \{0, 1, ..., t\}.$$

The values are evenly spaced where least value $g^*/(f^*+g^*+tc)$ decreases and greatest value $(g^*+tc)/(f^*+g^*+tc)$ increases as time t increases. State probabilities are more involved, as they are products of terms such as $n(f^*, g^*, t-1, i, c)$ and its complement $1-n(f^*, g^*, t-1, i, c)$.

Variance Attributes

Using expression (3), we calculate the variance of yield expectations conditional on information available at the preceding time point. Write $V_{t-1}(\mu_t) = E[(\mu_t - \mu_{t-1})^2 \mid \Omega_{t-1}]$ so that:

---

[6] As a practical matter when hedging or assessing value, the process could be modified so that exchangeability is not required. This is often the case when modeling stock prices with binomial trees to account for dividend effects. Allowing for this is just a matter of introducing an additional line in the programming code that would be used to support asset valuation, where the assumed distribution can be established through Monte Carlo simulations.

[7] There are three time points: 0, 1, and 2. There are two periods: 0 through 1 and 1 through 2.

[8] The urn interpretation of general expression (3) also applies to equation (5). The total number of balls at time t is $f^* + g^* + c$. The fraction that are green can be $g^*/(f^* + g^* + tc)$, arising when all t weather draws were bad. The fraction can be $(g^* + tc)/(f^* + g^* + tc)$, arising when all t weather draws were good. Or it can be at $t - 1$ equally spaced points between these two values.

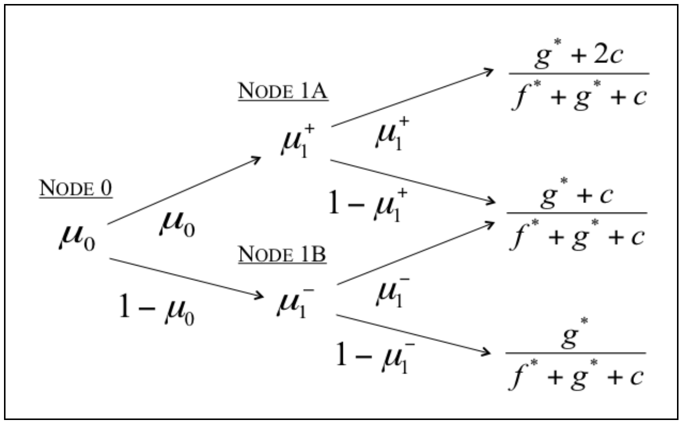

Figure 1. Binomial tree for the Pólya urn process, probabilities under arrows while {good then bad} and {bad then good} are equiprobable events

$$(6)\qquad V_{t-1}(\mu_t) = \left[\frac{\mu_{t-1} + m(f^*, g^*, t, c)}{1 + m(f^*, g^*, t, c)} - \mu_{t-1}\right]^2 \mu_{t-1} + \left[\frac{\mu_{t-1}}{1 + m(f^*, g^*, t, c)} - \mu_{t-1}\right]^2 (1 - \mu_{t-1})$$

$$= \frac{c^2\left[(1-\mu_{t-1})^2\mu_{t-1} + \mu_{t-1}^2(1-\mu_{t-1})\right]}{[1 + m(f^*, g^*, t, c)]^2[f^* + g^* + (t-1)c]^2} = \frac{c^2(1-\mu_{t-1})\mu_{t-1}}{[f^* + g^* + tc]^2}.$$

Variance is maximized when predicted yield is at the midpoint of the support, $\mu_{t-1} = 0.5$, and minimized at support extremes, $\mu_{t-1} \in \{0, 1\}$. For example, if the growing season started poorly, then expected yield is unlikely to move dramatically upward in light of one period's good weather. Similarly, if the growing season started well, then a future period's bad weather is unlikely to move expected yield markedly downward. These observations we refer to as resilience.

Equation (3) can be used to interpret equation (6) in terms of the magnitudes of forecast updates:

$$(7)\qquad \mu_{t+1}^+ - \mu_t = \frac{\mu_t + m(f^*, g^*, t, c)}{1 + m(f^*, g^*, t, c)} - \mu_t = (1-\mu_t)\frac{m(f^*, g^*, t, c)}{1 + m(f^*, g^*, t, c)},$$

$$\mu_{t+1}^- - \mu_t = \frac{\mu_t}{1 + m(f^*, g^*, t, c)} - \mu_t = -\mu_t \frac{m(f^*, g^*, t, c)}{1 + m(f^*, g^*, t, c)}.$$

If μ_t is close to 1, then any upward revision in the expectation is small. If μ_t is close to 0, then any downward revision in the expectation is small.

The sensitivity of $V_{t-1}(\mu_t)$ with respect to time is involved because t arises in two places. It appears in the denominator of the final expression in (6). This is referred to as the deterministic effect of time. Time also appears as a subscript on μ_{t-1}, and this appearance is referred to as the stochastic effect. The deterministic effect decreases step-ahead conditional variance over time and narrows the step-ahead conditional variance near harvest. The interpretation of the stochastic effect of time is less clear-cut in its consequence. If the process moves away from the middle of the distribution over time, then variance declines because of resilience. If the process moves toward the middle, then variance increases as the crop's prospects become more weather sensitive.

Marginal Weather Effects

Equations (2) and (3) indicate that the spread in expectations due to good weather is increasing and concave in c. To see this, differentiate $\mu_t^+ - \mu_t^- = c/(f^* + g^* + tc)$ to obtain:

(8) $$\frac{d^2(\mu_t^+ - \mu_t^-)}{dc^2} = -\frac{2(f^* + g^*)t}{(f^* + g^* + tc)^3} < 0.$$

This indicates there are diminishing returns to good weather. Also, the expression $d(\mu_t^+ - \mu_t^-)/dt < 0$ confirms that the step-ahead increase in variance decreases over time.

Beta Limiting Distribution

In binomial tree applications, constant proportional up and down movements approximate the lognormal distribution as step size decreases and the number of steps increase so that variance is unchanged. What happens to the discrete distribution in expression (3) under this limit operation? Pólya (1931) demonstrated that a limiting distribution for Y_T exists, and it has density:[9]

(9) $$\frac{\Gamma\left((f^* + g^*)/c\right)}{\Gamma\left(f^*/c\right)\Gamma\left(g^*/c\right)}\left(1 - Y_T\right)^{(g^*/c)-1} Y_T^{(f^*/c)-1},$$

where $\Gamma(\cdot)$ is the Gamma function. Thus, as the number of time periods to harvest increases, the realizations and associated probabilities for the discrete distribution given in (3) converge to a beta distribution.

Using standard formulae for beta distribution moments, the planting time information-conditioned mean and variance of this limiting distribution are

(10) $$\mu_0 = E[Y_T \mid \Omega_0] = \frac{(g^*/c)}{(f^*/c) + (g^*/c)} = \frac{g^*}{f^* + g^*}$$

and

(11) $$V_0(Y_T) = E[(Y_T - \mu_0)^2 \mid \Omega_0] = \frac{cf^*g^*}{(f^* + g^* + c)(f^* + g^*)^2}.$$

The mean is independent of the weather outcome parameter c, but the variance is not. Variance increases with the parameter c, which measures the sensitivity of yield expectations to weather realizations. Variance is 0 under $c = 0$, while $\lim_{c\to\infty} V_0(Y_T) = \mu_0(1 - \mu_0)$. The limiting distribution under $c \to \infty$ is the binomial, a special case of the beta distribution.

[9] See Mahmoud (2009, p. 53) for a proof.

Derivative Analysis

In this section, we develop the mechanics of hedging yield risk using weather derivatives, which have become popular offerings on formally organized exchanges.[10] Weather derivatives based on cooling and heating degree days are generally the most liquid (Morrison, 2009), and many researchers have explored their implications for crop insurance (Vedenov and Barnett, 2004; Woodard and Garcia, 2008; Xu, Odening, and Musshoff, 2008). Both the stochastic process described by expression (3) and the limiting distribution given in (9) identify crop yields to be dependent on stochastic weather realizations, as represented by the presence of c. Thus, our model provides an explicit link between weather stochastic processes and yield expectations. As such, yield has value that depends on weather, i.e., it is a derivative and its value can be hedged if a financial instrument exists that pays out on relevant weather events.

Without loss of generality, we assume a one-period process where yield is realized at $t = 1$ when weather derivatives are used to hedge yield losses.[11] An Arrow-Debreu financial instrument can be purchased in a liquid market at time 0, say planting, for $\$q$ where the pay-off at time $t = 1$ is \$1 whenever good weather occurs. The payoff is:

$$U = \begin{cases} 1 & \text{if good weather in period 1,} \\ 0 & \text{if bad weather in period 1.} \end{cases} \tag{12}$$

The crop price is assumed to be known as P at $t = 0$. Applying the standard modeling approach for hedging a derivative that can be represented as a binomial tree, assume the hedger takes a time $t = 0$ long position in v weather derivative instruments.

Equation (3) indicates that the time 1 state-contingent payoffs are:

$$P\mu_1 + vU = \begin{cases} \dfrac{P(g^* + c)}{f^* + g^* + c} + v & \text{with prob. } \dfrac{g^*}{f^* + g^*}, \\[2ex] \dfrac{Pg^*}{f^* + g^* + c} & \text{with prob. } \dfrac{f^*}{f^* + g^*}. \end{cases} \tag{13}$$

Set the derivative position at $v = v^{\#}$, where $v^{\#}$ is the value of v that equates payoffs across weather states. That is, solve $v + P(g^* + c)/(f^* + g^* + c) = Pg^*/(f^* + g^* + c)$ so that $v^{\#} = -Pc/(f^* + g^* + c)$. Hedge choice $v = v^{\#}$ secures payoff $P\mu_1 + vU = Pg^*/(f^* + g^* + c)$, which is invariant to weather outcome, i.e., it is risk-free. Thus, equation (13) represents a risk-free portfolio that should be a perfect substitute for a risk-free bond. Investors should be indifferent between owning a risk-free bond and owning a crop for which $v = v^{\#}$ weather derivatives have been purchased. Similar to equation (1), $v^{\#}$ bears a logistic relation with c.

For continuous-time interest rate r, the time $t = 0$ present value of the payout is $Pg^*e^{-r}/(f^* + g^* + c)$. With time $t = 0$ value of this crop defined as W_0, it follows that the $t = 0$ portfolio value is $W_0 + v^{\#}q = Pg^*e^{-r}/(f^* + g^* + c)$ so that

$$W_0 = \frac{P(g^*e^{-r} + cq)}{f^* + g^* + c}. \tag{14}$$

[10] Interested readers who are not familiar with financial binomial tree methods are directed to Hull (2009, chapter 11).
[11] For a more extensive tree, it is only necessary to apply the hedge to follow at each node.

This is a risk-neutral value for the crop, i.e., the fair market price when the opportunity to trade away risk is available. It can be achieved whenever the standard assumptions on transactions costs and market completeness apply.[12] Some reflection on how the valuation was arrived at shows that the method is essentially a dynamic program. The hedge identified in equation (13) would be applied at each node on a binomial tree and backward recursion would be used to arrive at risk-purged crop valuations during the growing season.

Some caution is warranted concerning the reliability of equation (14) when valuing crops. For example, hedging actions that support this valuation are sensitive to the model specification. If the yield expectation model is not representative of the true stochastic process, then the suggested hedging action will not completely remove risk and the present value of the resulting portfolio should not be calculated using the risk-free rate. In addition, basis risk will typically exist in that the weather affecting the crop being hedged will differ from the weather measurement underlying the weather derivative. To the extent unhedged risk would remain, it would not be appropriate to assume a risk-free rate when arriving at equation (14).

In equation (14), if $c = 0$, so that $V_0(Y_T) = 0$ by (6), then $W_0 = Pg^*e^{-r}/(f^* + g^*) = P\mu_0 e^{-r}$. This is the discounted expected crop value, known with certainty at the beginning of the growing season. In that case, $v^\# = 0$ and the good weather state price q does not enter the valuation because hedging serves no purpose. A model for the price of q, i.e., the pricing of weather risk, is not required when $c = 0$.

One approach to crop valuation is to use a risk-neutral measure (Hull, 2009), if it exists, is unique, and can be found. In our case, the risk-neutral measure is easy to obtain. If π is the risk-neutral probability of favorable weather, then equations (2) and (14) provide:

$$(15)\quad W_0 = \frac{P(g^*e^{-r} + cq)}{f^* + g^* + c} = Pe^{-r}\left[\frac{(g^* + c)}{f^* + g^* + c}\pi + \frac{g^*}{f^* + g^* + c}(1-\pi)\right] = \frac{Pe^{-r}(g^* + \pi c)}{f^* + g^* + c}.$$

Upon cancellation, it follows that $\pi = qe^r$. Thus, the risk-neutral, as distinct from true, probability of a favorable weather event over the period is taken as the forward discounted market valuation of an Arrow-Debreu derivative that pays off \$1 upon the event. Of course, risk preferences are not absent in (15); it is just that they have been corralled into a premium embedded in price q.[13]

This brings us to formally modeling the good weather state price $\$q$ where the payoff is \$1 in the event of good weather for time 1. Let the good weather state bear a capital asset pricing model (CAPM) beta coefficient of β_q with the capital market as a whole. Rate r_m is the overall market rate of return, where $r_m > r$ is assumed. Then the continuous-time equilibrium expected rate of return on the binary payout good weather event asset will be $r_q = r + \beta_q(r_m - r)$.

Since the expected payout is μ_0, the state price will be $q = \mu_0 \exp[-r - \beta_q(r_m - r)]$. Insert this value together with $\mu_0 = g^*/(f^* + g^*)$ into equation (14) to obtain:

$$(16)\qquad W_0 = P\mu_0 e^{-r}\frac{[f^* + g^* + ce^{-\beta_q(r_m - r)}]}{f^* + g^* + c},$$

[12] One needs a liquid market with payouts that are strictly monotonic in exactly the appropriate weather risk. Transactions costs should be negligible to support frequent portfolio rebalancing. Furthermore, risk-free interest rate r should be common to both lenders and borrowers.

[13] Readers interested in understanding the distinction between the risk-neutral and the true probability measure are referred to Hull (2009, pp. 241–242).

whereas $P\mu_0 e^{-r}$ is the crop value under risk neutrality. The risk premium is:

$$(17)\qquad P\mu_0 e^{-r} - W_0 = P\mu_0 e^{-r} c\,\frac{[1-e^{-\beta_q(r_m-r)}]}{f^* + g^* + c} = Pe^{-r}J(\cdot),\qquad J(\cdot) \equiv \mu_0 c\,\frac{[1-e^{-\beta_q(r_m-r)}]}{f^* + g^* + c}.$$

If good crop-growing weather is positively correlated with overall market returns, or $\beta_q > 0$, then $1 > e^{-\beta_q(r_m-r)}$ and $P\mu_0 e^{-r} > W_0$. Therefore, premium factor $J(\cdot)$ conveys that the crop's value at planting will be lower than its present value when discounted by the risk-free rate.

Finally, with $\mathbf{w} = (w_1, ..., w_N) \in \overline{\mathbb{R}}_+^N$ as the input price vector corresponding to input choice vector $\mathbf{x} \in \overline{\mathbb{R}}_+^N$, $\mathbf{x} = \mathbf{x}^*$ is chosen so that W_0 in expression (14) is its maximum possible value. So interior solutions satisfy

$$(18)\qquad \frac{\partial W_0}{\partial x_i} = Pe^{-r}\left(\frac{\partial \mu_0}{\partial x_i} - \frac{\partial J(\cdot)}{\partial x_i}\right) = w_i\,.$$

The marginal risk premium, reflected in $\partial J(\cdot)/\partial x_i$, could be positive or negative depending on whether the input is risk-increasing or risk-decreasing (Roosen and Hennessy, 2003).

Extensions

One limitation of the basic model is that stochastic realizations are not allowed to evolve over the growing season. Another is that its exchangeability property may fail in the sense that a good-then-bad realization does not lead to the same expectation as a bad-then-good realization. A third limitation is that the model does not allow for zero yield where zero yield may be due to crop abandonment or to an event that kills the growing crop. In this section, we discuss these matters and develop an approach that overcomes some of these limitations.

To address the first two concerns, one can allow parameter c to change with time steps (i.e., $c = c_t$) and/or with states. For example, in a two-period model, one can allow parameter c to change with time such that $c = c_1$ over step 1 and $c = c_2$ over step 2. In that case, expression (4) would become

$$(19)\qquad \mu_2 = \begin{cases} \dfrac{g^*}{f^* + g^* + c_1 + c_2} & \text{with prob. } \dfrac{f^*(f^* + c_1)}{(f^* + g^* + c_1)(f^* + g^*)} \quad \text{(bad then bad state)},\\[2ex] \dfrac{g^* + c_2}{f^* + g^* + c_1 + c_2} & \text{with prob. } \dfrac{f^* g^*}{(f^* + g^* + c_1)(f^* + g^*)} \quad \text{(bad then good state)},\\[2ex] \dfrac{g^* + c_1}{f^* + g^* + c_1 + c_2} & \text{with prob. } \dfrac{f^* g^*}{(f^* + g^* + c_1)(f^* + g^*)} \quad \text{(good then bad state)},\\[2ex] \dfrac{g^* + c_1 + c_2}{f^* + g^* + c_1 + c_2} & \text{with prob. } \dfrac{(g^* + c_1)g^*}{(f^* + g^* + c_1)(f^* + g^*)} \quad \text{(good then good state)}. \end{cases}$$

Allowing the values of c_1 and c_2 to differ relaxes the exchangeability assumption.[14] In particular, if $c_1 > c_2$, then it would be better to have good weather occur before bad weather rather than the reverse. This assumes plants with poor germination rates or weak early root development have limited capacity to take advantage of subsequent good weather.

[14] Equation (6) shows that letting the value of parameter c change with the time step allows the programmer to vary the variance in yield expectations from period to period as one moves through the growing season.

Alternatively, consider allowing parameter c to change with the state. In the two-step model, one could have just parameter value c over step 1, value $c + \varepsilon$ with $\varepsilon > 0$ over step 2 if the state is good in step 1, and value $c - \varepsilon$ over step 2 if the state is bad in step 1. Requirement $\varepsilon > 0$ allows for higher variance in yield expectations conditional on good weather in the first step when compared with a bad weather outcome in the first step. Cropping in the semi-arid Sahel region on the southern margin of the Sahara Desert illustrates an example of this situation. Millet and sorghum are planted only if spring rains occur. Both crops also need further rain after planting, as the soils are sandy and canopy cover is insufficient to allow for substantial soil water storage in the face of intense heat. Thus, the conditional expectation of yield and the conditional variance are both zero in the absence of spring rains. However, the conditional moments are positive when spring rains occur.

If parameter c is allowed to change with the state, distribution (4) becomes

$$(20) \quad \mu_2 = \begin{cases} \dfrac{g^*}{f^* + g^* + 2c - \varepsilon} & \text{with prob. } \dfrac{f^*(f^* + c)}{(f^* + g^* + c)(f^* + g^*)} \ \text{(bad then bad state)}, \\ \dfrac{g^* + c - \varepsilon}{f^* + g^* + 2c - \varepsilon} & \text{with prob. } \dfrac{f^* g^*}{(f^* + g^* + c)(f^* + g^*)} \ \text{(bad then good state)}, \\ \dfrac{g^* + c}{f^* + g^* + 2c + \varepsilon} & \text{with prob. } \dfrac{f^* g^*}{(f^* + g^* + c)(f^* + g^*)} \ \text{(good then bad state)}, \\ \dfrac{g^* + 2c + \varepsilon}{f^* + g^* + 2c + \varepsilon} & \text{with prob. } \dfrac{(g^* + c)g^*}{(f^* + g^* + c)(f^* + g^*)} \ \text{(good then good state)}. \end{cases}$$

The expression $(g^* + c - \varepsilon) / (f^* + g^* + 2c - \varepsilon) < (g^* + c) / (f^* + g^* + 2c + \varepsilon)$ indicates it is better to have good beginning weather followed by poor ending weather than the reverse. The expression reduces to $g^* < f^* + \varepsilon.$ If the probability of a good state at season outset is low, or $g^* < f^*,$ then it would be better to have a good start and a bad end than a good end and a bad start. The adaptations suggested above come at a cost in that the limiting distribution is very unlikely to be beta. Then Monte Carlo simulation analyses could be used to identify the shape of the harvest yield distribution.

One final point is that the model could be used to study production decisions. Suppose, as in Antle (1983), information is available about intra-season production decisions. If so, the dynamic program used to arrive at asset valuation (14) can be adapted to allow for optimum input choices at each node of the decision tree. In figure 1, it could be that Node 1A, where good early-season weather occurs, supports additional pesticide use relative to Node 1B so as to protect the promising crop. Accounting for these state-conditional decisions will affect valuation earlier in the season, and so input choices earlier in the season.

Even in the world's most productive cropping regions, it is not uncommon for a crop to fail completely (Mendelsohn, 2007). In order to understand how, as a practical matter, one might account for such risks when modeling a yield distribution, consider two scenarios. One is where some catastrophic event with probability η, perhaps a hail storm or flood, destroys an immature crop. In this case, suppose the probability of such an event over the growing season is independent of the crop's performance and the crop's performance follows the process we suggest. Then the yield distribution is a mixture of distributions. It is that given as the harvest outcome from algorithm (3) with probability $1 - \eta$, and is zero with probability η. Crop valuation along the lines of expression (14) above would involve identifying a premium for insuring against the catastrophic event.

There is an alternative way by which zero yield may arise, one that involves intra-season decisions. This is that the crop is abandoned (i.e., an endogenous failure) mid-season as an unprofitable bet for continued investment (Chen and Miranda, 2007). Consider figure 1 again and suppose, regardless of weather, costs amounting to $K > 0$ must be incurred mid-season or the crop dies. With output price P, expected profit is $P\mu_1^+ - K$ whenever weather is good early in the season. The grower does not abandon the crop since the profit situation is more favorable than it had been at planting. If weather is bad early in the season, then expected profit is $P\mu_1^- - K$, and this could be negative even though the grower saw fit to plant. If $P\mu_1^- < K$, then it would be privately logical, and socially efficient absent external issues, to abandon the crop.

This abandonment option has been studied widely in the economics of investment (Trigeorgis, 1996). The general assumption in the literature is that the stochastic process at issue follows geometric Brownian motion. For our purpose, abandonment would mean the yield distribution would no longer be characterized by the distribution in (4), but rather by

$$\text{(21)} \qquad \mu_2 = \begin{cases} 0 & \text{with prob.} \quad \dfrac{f^*}{f^* + g^*}, \\ \dfrac{g^* + c}{f^* + g^* + 2c} & \text{with prob.} \quad \dfrac{f^* g^*}{(f^* + g^* + c)(f^* + g^*)}, \\ \dfrac{g^* + 2c}{f^* + g^* + 2c} & \text{with prob.} \quad \dfrac{(g^* + c) g^*}{(f^* + g^* + c)(f^* + g^*)}. \end{cases}$$

One can readily program a multi-step algorithm to account for the possibility of mid-season abandonment when there may be multiple in-season decision points.[15] The yield distribution this program would construct depends on the timing and cost of decisions as well as the underlying yield expectation stochastic process.

Conclusion

This paper has developed a method for modeling expected yield as an information-conditioned stochastic process over the course of a crop growing season. The process has as its limiting distribution a commonly estimated crop yield distribution. The process allows for yield randomness similar to derivative valuation methods used for modeling commodity prices. In the presence of appropriate financial instruments such as weather derivatives, the process allows for the identification of optimal dynamic hedging strategies. In the absence of such instruments, risk managers can use the framework to sharpen real-time risk assessments.

The approach has two appealing features. First, it provides a model of stochastic crop production that is consistent with modern finance theory; and second, it is flexible. A binomial tree approach allows for calibration and adapts readily to discrete-time dynamic programming methods. For example, the c parameter can be allowed to vary as the growing season progresses to better comply with historical information on yield expectations and findings in the agronomy literature.

[15] There are no conceptual difficulties in risk-neutral crop valuation along the lines of valuation (14), but with the option to abandon.

Concerning further possible developments, a bivariate extension to account for price-yield correlations would be required to model revenue insurance contracts within this framework. Methods are available to model bivariate discrete time stochastic processes (Boyle, 1988; Ho, Stapleton, and Subrahmanyam, 1995). As far as we know, the use of these methods has been confined to approximating multivariate geometric Brownian motion. Therefore, modifications would be required if such models were to be used for rating crop revenue insurance contracts. Our approach also provides a framework for looking at interactions between how information becomes available and production decisions. One such interaction concerns the decision to abandon a crop. Our approach affords opportunities to understand how a grower's set of available information affects that choice. It may also be possible to introduce insurance indemnities into the model and then study how insurance contract design affects moral hazard.

[*Received October 2009; final revision received January 2011.*]

References

Antle, J. M. "Sequential Decision Making in Production Models." *Amer. J. Agr. Econ.* 65,2(1983):282–290.

Antle, J. M., S. M. Capalbo, and C. C. Crissman. "Econometric Production Models with Endogenous Input Timing: An Application to Ecuadorian Potato Production." *J. Agr. and Resour. Econ.* 19,1(1994):1–18.

Antle, J. M., and S. A. Hatchett. "Dynamic Input Decisions in Econometric Production Models." *Amer. J. Agr. Econ.* 68,4(1986):939–949.

Araus, J. L., G. A. Slafer, C. Royo, and M. D. Serret. "Breeding for Yield Potential and Stress Adaptation in Cereals." *Critical Reviews in Plant Sci.* 27,6(2008):377–412.

Atwood, J., S. Shaik, and M. Watts. "Are Crop Yields Normally Distributed? A Reexamination." *Amer. J. Agr. Econ.* 85,4(2003):888–901.

Babcock, B. A. "Corn Belt Contributions to the Crop Insurance Industry." *Iowa Ag Review* 14,2(2008):1–3 and 11.

Babcock, B. A., and D. A. Hennessy. "Input Demand Under Yield and Revenue Insurance." *Amer. J. Agr. Econ.* 78,2(1996):416–427.

Bhattacharya, R. N., and E. C. Waymire. *Stochastic Processes with Applications*. New York: Wiley, 1990.

Bloom, A. J., F. S. Chapin, and H. A. Mooney. "Resource Limitation in Plants—An Economic Analogy." *Annual Rev. Ecology and Systematics* 16,1(1985):363–392.

Borges, R. B., and W. N. Thurman. "Marketing Quotas and Random Yields—Marginal Effects of Inframarginal Subsidies on Peanut Supply." *Amer. J. Agr. Econ.* 76,4(1994):809–817.

Boyle, P. P. "A Lattice Framework for Option Pricing with Two State Variables." *J. Finan. and Quant. Analy.* 23,1(1988):1–12.

Carrasco-Tauber, C., and L. J. Moffitt. "Damage Control Econometrics: Functional Specification and Pesticide Productivity." *Amer. J. Agr. Econ.* 74,1(1992):158–162.

Chambers, R. G. *Applied Production Analysis: A Dual Approach*. New York: Cambridge University Press, 1988.

Chen, S.-L., and M. J. Miranda. "Effects of Insurance on Farmer Crop Abandonment." Paper presented at annual meetings of the American Agricultural Economics Association, Portland, OR, 29 July–1 August 2007.

Claassen, R., and R. E. Just. "Heterogeneity and Distributional Form of Farm-Level Yields." *Amer. J. Agr. Econ.* 93,1(January 2011): forthcoming.

Dercon, S. "Risk Insurance and Poverty: A Review." In *Insurance Against Poverty*, ed., S. Dercon, pp. 9–38. New York: Oxford University Press, 2004.

Durrett, R. *Probability: Theory and Examples*, 2nd ed. Belmont, CA: Duxbury, 1996.

Goodwin, B. K. "Payment Limits and Acreage Decisions Under Risk Aversion: A Simulation Approach." *Amer. J. Agr. Econ.* 91,1(2009):19–41.

Harri, A., C. Erdem, K. H. Coble, and T. O. Knight. "Crop Yield Distributions: A Reconciliation of Previous Research and Statistical Tests for Normality." *Rev. Agr. Econ.* 31,1(2009):163–182.

Hennessy, D. A. "Crop Yield Skewness and the Normal Distribution." *J. Agr. and Resour. Econ.* 34,1(2009): 34–52.

Ho, T.-S., R. C. Stapleton, and M. G. Subrahmanyam. "Multivariate Binomial Approximations for Asset Prices with Nonstationary Variance and Covariance Characteristics." *Rev. Finan. Stud.* 8,4(1995):1125–1152.

Hull, J. C. *Options, Futures, and Other Derivatives*, 7th ed. Upper Saddle River, NJ: Pearson Education, Inc., 2009.

Innes, R. "Crop Insurance in a Political Economy: An Alternative Perspective on Agricultural Policy." *Amer. J. Agr. Econ.* 85,2(2003):318–335.

Just, R. E., and Q. Weninger. "Are Crop Yields Normally Distributed?" *Amer. J. Agr. Econ.* 81,2(1999): 287–304.

Mahmoud, H. M. *Pólya Urn Models*. Boca Raton, FL: CRC Press, 2009.

Makki, S. S., and A. Somwaru. "Evidence of Adverse Selection in Crop Insurance Markets." *J. Risk and Insur.* 68,4(2001):685–708.

Marcus, A. J., and D. M. Modest. "Futures Markets and Production Decisions." *J. Polit. Econ.* 92,3(1984): 409–426.

Mendelsohn, R. "What Causes Crop Failure?" *Climate Change* 81,1(2007):1573–1580.

Morrison, J. "Managing Weather Risk: Will Derivatives Use Rise?" *Futures Industry* (January/February 2009):26–29.

Nelson, C. H., and P. V. Preckel. "The Conditional Beta Distribution as a Stochastic Production Function." *Amer. J. Agr. Econ.* 71,2(1989):370–378.

Pólya, G. "Sur Quelques Pointes de la Théorie des Probabilités." *Annals of the Institute of Henri Poincaré* 1(1931):117–161.

Ramírez, O. A., and T. McDonald. "Ranking Crop Yield Models: A Comment." *Amer. J. Agr. Econ.* 88,4(2006):1105–1110.

Roosen, J., and D. A. Hennessy. "Testing for the Role of Risk Aversion on Input Use." *Amer. J. Agr. Econ.* 85,1(2003):30–43.

Saphores, J.-D. M. "The Economic Threshold with a Stochastic Pest Population: A Real Options Approach." *Amer. J. Agr. Econ.* 82,3(2000):541–555.

Shields, D. A. *Federal Crop Insurance: Background and Issues.* Congressional Research Service, Washington, DC, 13 December 2010. Online. Available at http://www.nationalaglawcenter.org/assets/crs/R40532.pdf.

Skees, J. R., J. Hartell, and A. G. Murphy. "Using Index-Based Risk Transfer Products to Facilitate Micro Lending in Peru and Vietnam." *Amer. J. Agr. Econ.* 89,5(2007):1255–1261.

Smith, V. H., and B. K. Goodwin. "Crop Insurance, Moral Hazard, and Agricultural Chemical Use." *Amer. J. Agr. Econ.* 78,2(1996):428–438.

Tollenaar, M., and E. A. Lee. "Yield Potential, Yield Stability, and Stress Tolerance in Maize." *Field Crops Research* 75,2–3(2002):161–169.

Trigeorgis, L. *Real Options: Managerial Flexibility and Strategy in Resource Allocation*. Cambridge, MA: MIT Press, 1996.

Tschirhart, J. "General Equilibrium of an Ecosystem." *J. Theoretical Biology* 203,1(2000):13–32.

Vedenov, D. V., and B. J. Barnett. "Efficiency of Weather Derivatives as Primary Crop Insurance Instruments." *J. Agr. and Resour. Econ.* 29,3(2004):387–403.

Woodard, J. D., and P. Garcia. "Weather Derivatives, Spatial Aggregation, and Systemic Risk: Implications for Reinsurance Hedging." *J. Agr. and Resour. Econ.* 33,1(2008):34–51.

Xu, W., M. Odening, and O. Musshoff. "Indifference Pricing of Weather Derivatives." *Amer. J. Agr. Econ.* 90,4(2008):979–993.