Stochastic Evaluation of Commodity Support Program Alternatives

Introduction

Farmers are generally averse to risk – in particular, to uncertain and economically unfavorable outcomes (Hardaker et al., 2004). While many sources of uncertainty have been identified (Moschini and Hennessy, 2001), this report focuses on the exposure of the farmer and the government to production (specifically, yield) and price uncertainty, which together translate to revenue uncertainty. Given that the producer is unlikely to be indifferent (or neutral) to risk, the producer is concerned with more aspects of revenue than simply its mean value. In short, risk aversion means that a farmer would tend to prefer a commodity support program under which some yearly average level of revenue is forgone in return for lower variability in year-to-year revenue.

While the risk preferences of individuals have received extensive study in the academic literature, the risk preferences of government have not. In the case of support payments, risk preferences may be defined as the Government's desire to decrease the variation in payments from projected budget levels. But the Federal Government is a large and heterogeneous body, and anecdotal evidence suggests that it has no uniform risk preference. However, certain program rules suggest that government agencies have at least some risk aversion with respect to costs. For example, starting with the 1996 Farm Bill, the Congressional Budget Office has used probability, or stochastic, scoring to estimate farm program costs (Jagger and Hull, 1997; Gardner, 1996). In addition, the Office of Management and Budget requires agencies to use probability scoring for estimating program costs if costs are uncertain.

Regardless of government agencies' risk preferences with respect to variability in payment levels, evaluating program costs in a probabilistic framework can identify costs that might not be identified otherwise. Specifically, given the highly stochastic (random) nature of prices and yield (and the many other variables that may affect prices), estimating program costs based simply on the point estimates of variables may not capture full budgetary costs of program change (Jagger and Hull, 1997). For example, just because the expected season-average price for a crop is greater than the trigger price (loan rate) for a marketing loan program does not mean support payments will be nonexistent, given that the average can mask prices that fall below the loan rate during the loan availability period.

Probability scoring is a cost estimate procedure that uses different projection paths for the key variables that are likely to affect corresponding program costs, thereby generating a statistical distribution of program costs. Even if the probability scoring provides only the mean of the estimated distribution of program costs, as it usually does, some aspect of the budgetary risk can still be captured. For example, a proposed program may show no costs using point estimates but higher costs when the mean is based on a probabilistic analysis. Nonetheless, the estimated distribution of program costs (in particular, farm support payments) provides additional information that

may be of policy relevance to the government and of practical relevance to producers. For instance, if the government intended to reduce the likelihood that payments exceed a certain ceiling, then such an objective could be examined using the probabilistic approach.

To gain some insights into the policy implications of revenue support programs, this chapter compares the statistical distribution of payments from hypothetical revenue-based programs to those from a suite of programs similar to the traditional set of commodity support programs. While probability-based program analysis, as used in legally required government cost estimates, summarizes the distribution of program costs into mean estimates, other summary statistics – such as the variance and skewness (shape) of the distribution – are useful too. The estimated payment distributions have implications both for government policy and for farm-level benefits.
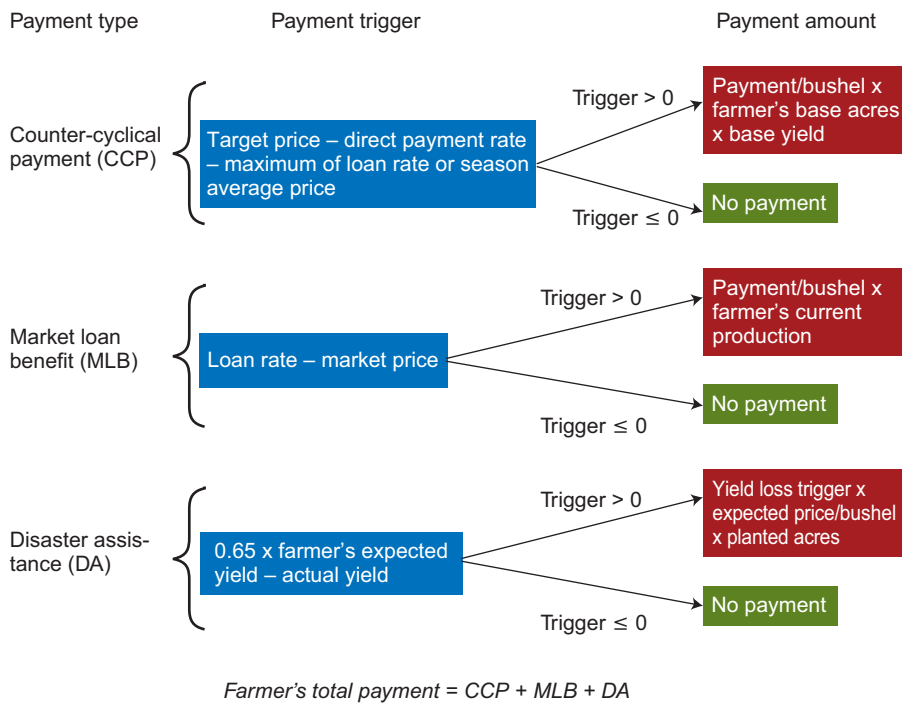
Commodity Support Program Scenarios

Actual program payments are sensitive to a broad array of program provisions, and seemingly small changes in these can cause large changes in payment levels. Hence, to make the support programs comparable, our program scenarios are designed to differ only in the fundamental program provisions. The goal is to investigate how payments are affected by using revenue targets rather than price or yield targets, and not how payments are affected by program parameters inherent to these targets. The traditional-style program scenario is compared with two revenue-based program scenarios, one based (in part) on revenue shortfalls with respect to a target revenue, and one based on revenue shortfalls with respect to an expected market revenue (see "Appendix B. Technical Details of the Stochastic Analysis").

Traditional-Style Domestic Program Scenario

Our scenario for a generic version of traditional commodity support has three components: countercyclical payments (CCP), marketing loan benefits (MLB), and disaster assistance (DA) payments. Disaster assistance payments are usually based on a shortfall in yield with respect to expected yield, where the lost production is valued at an "established" or expected price (see the three boxes in this section for representations of these program scenarios using flow diagrams). We assume that DA payments operate in this manner, but on a permanent rather than ad hoc basis, like a form of crop yield insurance that is free to the producer. As is frequently the case in actual practice (e.g., the 2001 and 2002 *ad hoc* disaster programs), we assume that payments are made when the producer's yield is reduced by more than 35 percent from the expected yield. Unlike the MLB, DA payments can be nonzero even if harvested yield is zero.

Schematic of payments under the traditional-style support program scenario

Payment type	Payment trigger		Payment amount

Counter-cyclical payment (CCP)

Target price – direct payment rate – maximum of loan rate or season average price

Trigger > 0 → Payment/bushel x farmer's base acres x base yield

Trigger ≤ 0 → No payment

Market loan benefit (MLB)

Loan rate – market price

Trigger > 0 → Payment/bushel x farmer's current production

Trigger ≤ 0 → No payment

Disaster assistance (DA)

0.65 x farmer's expected yield – actual yield

Trigger > 0 → Yield loss trigger x expected price/bushel x planted acres

Trigger ≤ 0 → No payment

Farmer's total payment = CCP + MLB + DA
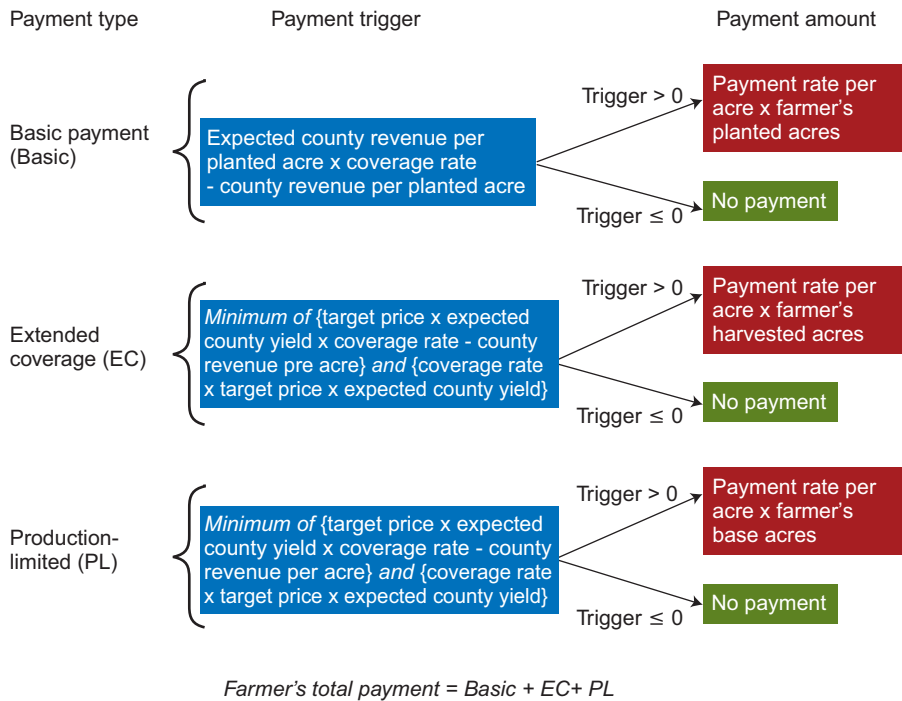
Target Revenue Program Scenario

We base the three components of this county-area revenue program on Babcock and Hart (2005) and NCGA (2006), with some minor differences (e.g., we use futures prices rather than cash prices). The "basic" component is a payment per planted acre to cover shortfalls with respect to expected revenue per acre, calculated at the county level. Expected county revenue is multiplied by a coverage rate between 0 and 1 such that, as with an insurance program, less than 100 percent of expected revenue is covered.

The "extended coverage" payment per harvested acre is based on a shortfall in revenue with respect to a target revenue based on a statutory price, and provides supplemental coverage over the basic payment. The revenue coverage rate for this component is greater than for the "basic" component, but still less than 1. As with the "basic" component, the payment rate for "extended coverage" is multiplied by the farmer's planted acreage for the current crop year.

The "production-limited" payment is similar to the extended coverage payment but applied to a fixed base acreage for the farmer, and provides supplemental coverage over the extended coverage payment. This payment is similar to the CCP in that payment does not require current production. The revenue coverage rate for this component is greater than for the "extended" component, but still less than 1.[12]

[12]The terms "basic," "production limited," and "extended coverage" substitute for the terms Babcock and Hart (2005) use, which are "green", "blue", and "amber," respectively. These colors ("boxes") are references to categorizations by the World Trade Organization's Agreement on Agriculture (AoA) of domestic subsidies according to their impacts on production. Since it is impractical to speculate on how a proposed program might be notified to the WTO and given the political controversy in multilateral negotiations over which support programs should be associated with each of these WTO "boxes," for the sake of avoiding the potential for confusion we avoid using the WTO terminology in our scenarios.

Schematic of payments under the target revenue program

| Payment type | Payment trigger | Payment amount |

| Payment type | Payment trigger | | Payment amount |

Basic payment (Basic)

Expected county revenue per planted acre x coverage rate - county revenue per planted acre

Trigger > 0 → Payment rate per acre x farmer's planted acres

Trigger ≤ 0 → No payment

Extended coverage (EC)

Minimum of {target price x expected county yield x coverage rate - county revenue pre acre} *and* {coverage rate x target price x expected county yield}

Trigger > 0 → Payment rate per acre x farmer's harvested acres

Trigger ≤ 0 → No payment

Production-limited (PL)

Minimum of {target price x expected county yield x coverage rate - county revenue per acre} *and* {coverage rate x target price x expected county yield}

Trigger > 0 → Payment rate per acre x farmer's base acres

Trigger ≤ 0 → No payment

Farmer's total payment = Basic + EC+ PL

Note: The coverage rate (value between 0 and 1) in each payment type are designed so that the farmer's total payment per acre does not exceed the target revenue per acre.

Market Revenue Program Scenario

The market revenue program proposal has two components: a national revenue payment (e.g., Zulauf, 2006; AFT, 2007a) and a supplemental county-area revenue payment. The national revenue payment (NRP) is calculated as a percentage decrease in national expected total revenue with respect to national average realized total revenue, times the farmer's expected revenue per planted acre times the farmer's planted acres.
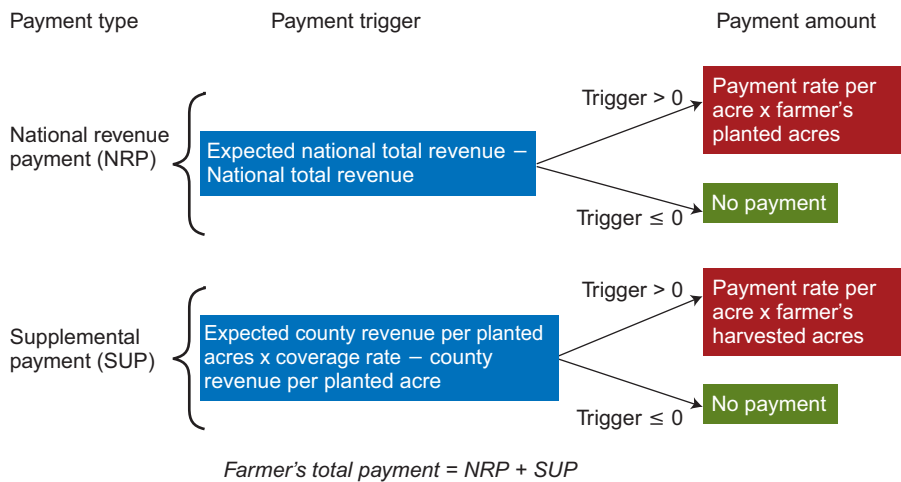
With the NRP triggered only by national shortfalls in revenue, Zulauf assumes that a Federal crop insurance program payment is used to ensure that the farmer is covered up to a guaranteed level. However, for the sake of comparability across scenarios, we instead use a supplemental county-area revenue payment to ensure that the farmer is covered up to a guaranteed level.

Comparability of the Payment Scenarios

Our target revenue program operates at the county level. To put each of the program scenarios on an equal footing for the simulation, all three are constructed to operate at the county level as well.

For the expected and harvest-time prices, we utilize futures prices, as discussed in more detail below. In the traditional-style and the target revenue programs, 2004 levels for acreage and yield serve as base acreage and yield. To calculate benefits in time t, we use the Olympic average of the prior 5 years' worth of yield data from USDA's National Agricultural Statistics Service (NASS), which is consistent with the approach used in various insurance products

Schematic of payments under the market revenue program

Payment type	Payment trigger		Payment amount

National revenue payment (NRP)

> Expected national total revenue − National total revenue

Trigger > 0 → Payment rate per acre x farmer's planted acres

Trigger ≤ 0 → No payment

Supplemental payment (SUP)

> Expected county revenue per planted acres x coverage rate − county revenue per planted acre

Trigger > 0 → Payment rate per acre x farmer's harvested acres

Trigger ≤ 0 → No payment

Farmer's total payment = NRP + SUP

administered by the USDA's Risk Management Agency (RMA), various disaster payments administered by the USDA's Farm Services Agency (FSA), and the revenue-based ACRE program passed into law in the 2008 Farm Act.

Programs can be compared against each other in many ways. Given limited information on the risk preferences of producers, it seems reasonable from a policy standpoint to assume that payment recipients would be reluctant to support a revised direct support program unless it provided at least the same support levels as the program it replaces. Hence, to narrow the range of possible program parameters, we calibrate the models by setting the program parameters so that the mean of total annual payments evaluated at each of the 31 price-yield points (over 1975 to 2005) is equal across the program scenarios. By doing so, we are not favoring one scenario over another with respect to the mean of the payment distribution. Given this calibration, other characteristics (for example, variance or skewness) of the distribution of payments can be compared, as can the program parameters necessary to achieve equality of mean total payments across programs. Details of the calibration procedure are presented in Appendix B, as is the methodology for estimating payments and the data sources.

Discussion of Results

Table 1 summarizes the results of the stochastic analysis, using 2005 data for planted acres and for the expected yield and price against which the price and yield deviations are applied. The first row under each scenario shows mean payments from the stochastic simulation and the next row the coefficient of variation of the payments. The coefficient of variation provides a measure of variability (the higher the value, the higher the variability) that allows for easier comparability across program scenarios than the standard deviation. The overall coefficients of variation for the two revenue approaches are roughly equal at 0.32 and 0.34. However, the coefficient of variation for the traditional program scenario is twice as high (0.68), with most of the contribution to this value coming from the fully production-coupled MLBs (the disaster payments have a higher coefficient of variation but account for a smaller portion of total payments).

Table 1

Stochastic analysis of the distribution of corn program payments under alternative U.S. programs (2005 expected prices and yields)

	Payment type			
Target Revenue Program	Total	Extended Coverage	Production Limited	Basic[1]
Mean payment ($ billion)	3.03	1.16	1.64	0.22
Coefficient of variation[2]	0.32	0.52	0.24	1.06
90% confidence interval (lower, upper)	1.62, 4.80	0.39, 2.28	1.06, 2.37	0.02, 0.73
Market Revenue Program	Total	National[3]	Supplmental	
Mean payment ($ billion)	3.17	2.33	0.85	
Coefficient of variation	0.34	0.430	0.59	
90% Confidence interval	1.55, 5.09	0.76, 4.06	0.37, 1.97	
Traditional-Style Program	Total	P-MLB	P-CCP	Disaster
Mean payment ($ billion)	3.11	1.26	1.67	0.19
Coefficient of variation	0.68	1.35	0.53	1.46
90% confidence interval	0.38, 7.10	0.00, 4.78	0.00, 2.28	0.02, 0.83

[1]The "basic" payment covers shortfalls in county revenue per acre with respect to expected county revenue per acre. The "extended coverage" payment is based on a target revenue using a statutory price, and provides supplemental coverage over the basic payment. The "production-limited" payment is similar to the extended coverage payment but applied to a fixed base acreage for the farmer, and provides supplemental coverage over the extended coverage payment.

[2]The coefficient of variation in this application is a measure of the dispersion of the probability distribution of revenue per acre that allows comparisons across populations with different means, and is the standard deviation of revenue per acre divided by the mean revenue per acre. The smaller the coefficient of variation, the lower the dispersion relative to the mean value of the distribution.

[3]The "national" revenue payment rate is based on the difference between national expected and actual revenue per acre, and the "supplemental" revenue payment provides additional coverage based on a county- level payment rate.

Among the three traditional-style program payment types, the price-based CCP has the lowest coefficient of variation (0.53), which is not surprising given the hard ceiling on the CCP payment rate. In fact, the coefficient of variation for the price-based CCP is lower than for the "basic" component (1.06) of the target revenue approach, but more than twice the value (0.24) of the "production limited" component. This difference is attributed to the formula for the "production limited" revenue payment rate (equation B.8 in Appendix B) versus the price-CCP payment rate (equation B.1) – the former has a more explicit limit on the payment rate than the latter.

The third row in table 1 presents the 90-percent confidence intervals calculated from the same bootstrap output. The lower bound of the 90-percent confidence band for the current-style scenarios includes zero or near-zero payment levels in all three payment types, but also several billion dollars at the upper end.[13] The traditional-style program scenario has a 90-percent lower bound that is more than $1 billion lower than for either of the two revenue-based programs, but an upper bound that is over $2 billion higher. This indicates that both farmers and the Government would face less uncertainty in budgeting for expected payments under the revenue-based alternatives examined here.

The Government is concerned with more than just the mean and variance of the empirical payment distribution. For example, in comparing program alternatives, it would be useful to have information on the probability that

[13]Actual production-coupled corn payments vary greatly from year to year. For instance, over 1996-2006, actual LDPs for the crop year were $0 in each of 4 years, but as high as $4.3 billion in the 2005 crop year (payment variation is less extreme on a fiscal-year basis).

Economic Aspects of Revenue-Based Commodity Support / ERR-72
Economic Research Service/USDA

commodity support levels will exceed those agreed to under a multilateral agreement on domestic support. The right-hand tail of the frequency distribution (a graph of how may times the bootstrapped payments fall within each billion-dollar interval) provides this information. Figures 5a-5c show both the frequency of total payments and the subset of payments most likely to face payment ceilings in future multilateral agreements on agricultural support. For example, the traditional-style scenario shows payments net of disaster payments given that disaster payments can under certain conditions be exempt from support ceilings (fig. 5a). Likewise, under the target revenue program, the basic portion of payments could be exempt from support ceilings, and so figure 5b shows payments net of basic payments as well as total payments. Figure 5c shows the market revenue payment net of the supplemental payment, although this breakdown is not intended to suggest that any portion of the market revenue payment be exempt from payment ceilings.

Given the premise of achieving the same mean annual payment level across the program scenarios, figures 5a-5c clearly show that the traditional-style support scenario has a fatter right-hand tail – or higher probability of exceeding a support ceiling – than the two revenue-based programs. For example, excluding the portion of payments that may not be subject to limits, the two revenue-based programs would exceed $6.5 billion less than 1 percent of the time, while the traditional-style program would exceed $6.5 billion in payments 12 percent of the time.

Budgetary Impacts Under Alternative Scenarios

This section presents an approach to empirically demonstrating how the within-season probability distribution of U.S. domestic commodity support for corn differs between traditional-style approaches to support and revenue-based support. In general, official government assessments of the costs of a program that use a probabilistic setting (known as "probability scoring") present only the mean of the probability distribution of program costs. However, other summary statistics, such as variance or skewness (shape) of the distribution of payments, may provide useful information as well, especially when comparing across program alternatives. For the revenue-based support scenarios evaluated here, variability around total expected annual payments and the probability of high payments are both lower than for the traditional-style approach. These results suggest less budgetary uncertainty for the Federal Government and easier adherence to multilateral commitments regarding limits to domestic commodity support. Of course, the empirical results in this section showing the benefits of revenue-based support with respect to the Federal budget pertain to the specific program scenarios examined here, and may not necessarily hold for program scenarios not examined here.[14]

Regional Implications of Revenue-Based Versus Price-Based Direct Commodity Support

The previous section examined the implications for Federal budgetary planning of the three support proposals by summing up the county-level payments from the stochastic simulation to the national level. This section examines how payments vary by region, focusing (for brevity's sake) on

[14]For instance, a price-based support program that is production-limited (that is, not coupled to current production) and has a hard ceiling on the effective farm price could have a lower coefficient of variation than a revenue-based support program that is production-limited but does not have a hard ceiling on the payment rate.

Figure 5a

Frequency of commodity payments for corn – traditional-style program
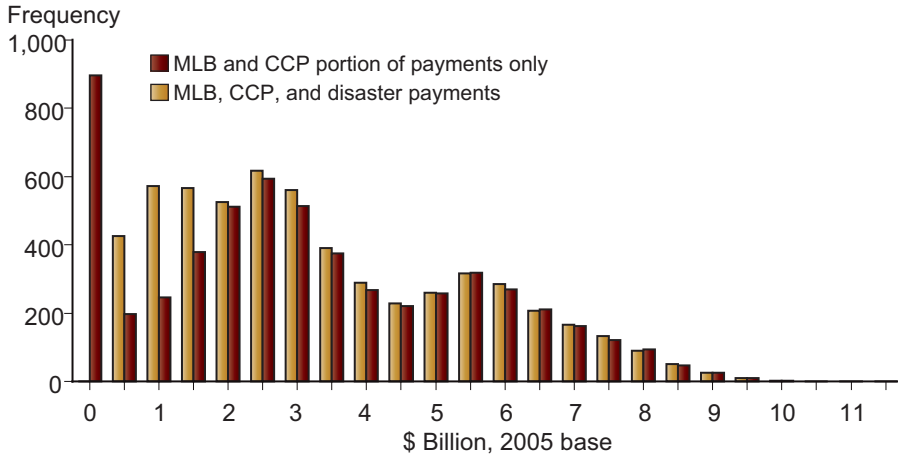*The traditional style programs more frequently have high payment*

Frequency

Figure 5b

Frequency of commodity payments for corn – target revenue program
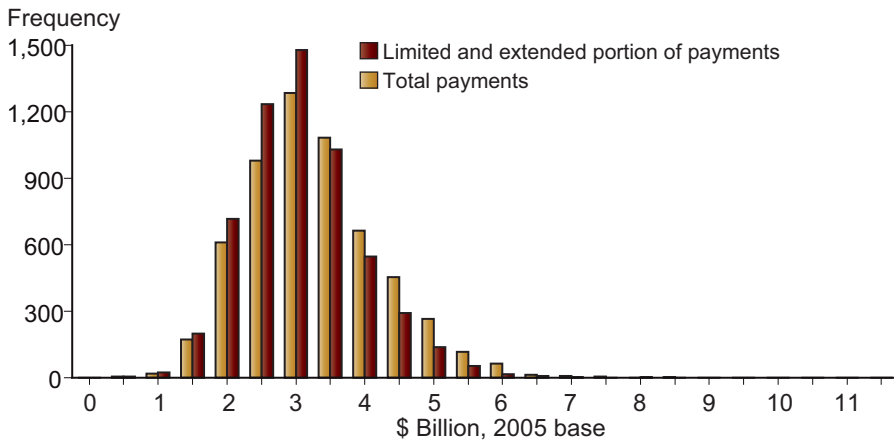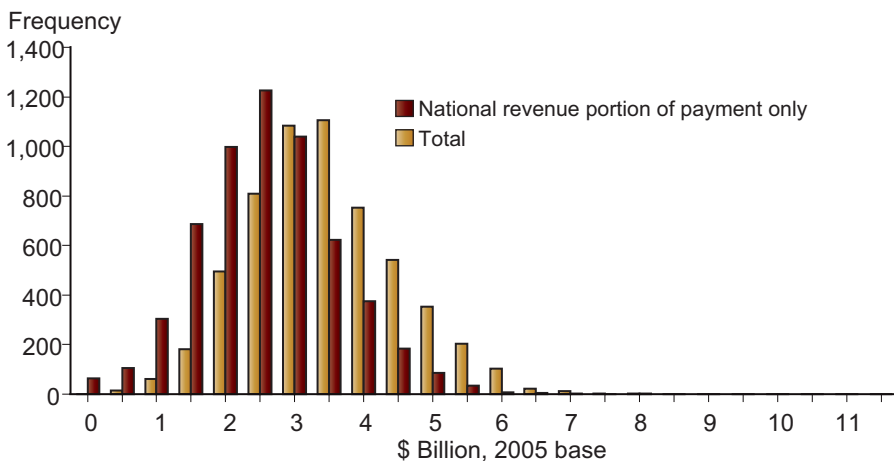*The target revenue programs produces a tighter range of payments.*

Frequency

Figure 5c

**Frequency of commodity payments for corn – market revenue program**

Frequency

Note: Each bar covers a $500 million range of payments. The taller the bar, the greater the number of payments falling in the associated range.

the traditional-style program versus the target revenue program. The results for the market revenue program are similar to those for the target revenue program, however, and can be found in Cooper (2007, 2009b).

Figure 6 shows the coefficient of variation for gross corn revenue by county. The smaller the coefficient, the lower the variation in average county revenue per acre relative to its mean. The pattern of groupings in the map suggests that the coefficient of variation has a significant regional component. Table 2 presents average county returns per acre and the associated coefficient of variation for corn, as summarized by ERS Farm Resource Regions (Heimlich, 2000). The table lists both the gross returns per acre (price times yield per acre) as well as total gross returns (gross returns plus the per-acre government payment) under both the current-style and target revenue programs.

Perhaps not surprisingly, the Heartland region has the lowest coefficient of variation for gross corn returns, indicating its comparative advantage in corn production. The coefficient of variation for total gross returns is lower under the target revenue than traditional-style programs for each region except the Fruitful Rim, where it is the same across programs (table 2). For the Heartland region, it is almost three times lower. Since the mean returns are roughly the same (by design) under either approach, a safety net intended to reduce variability in total gross income might benefit from a revenue-based approach, for corn at least.

Figure 6
Coefficient of variation of gross corn revenue per acre
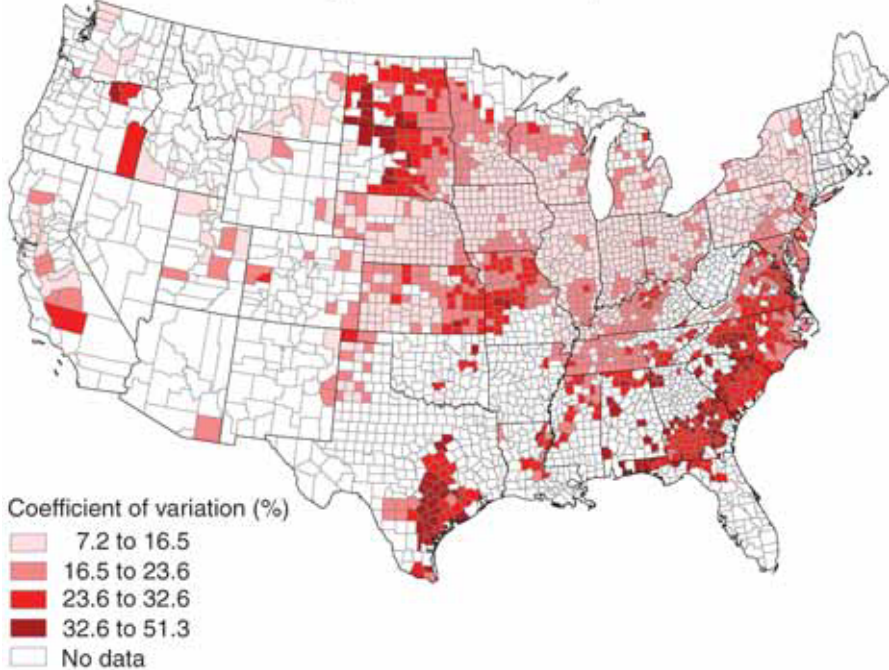
Coefficient of variation (%)
- 7.2 to 16.5
- 16.5 to 23.6
- 23.6 to 32.6
- 32.6 to 51.3
- No data

Note: The coefficient of variation in this application is a measure of the dispersion of the probability distribution of revenue per acre that allows comparisons across populations with different means, and is the standard deviation of revenue per acre divided by the mean revenue per acre. The smaller the coefficient of variation, the lower the dispersion relative to the mean value of the distribution.

Source USDA, Economic Research Service.

Table 2
County-level revenue per acre with and without program payments, by farm resource region

Farm resource region	Share of total corn acres (percent)	County gross revenue		County traditional-style plus gross revenue		County target revenue payment plus gross revenue	
		Mean ($/acre)	Coefficient of variation (percent)	Mean ($/acre)	Coefficient of variation (percent)	Mean ($/acre)	Coefficient of variation (percent)
Heartland	61.6	286	15	326	14	330	5
Northern Crescent	13.5	246	16	277	13	277	7
Northern Great Plains	5.8	220	24	248	19	248	11
Prairie Gateway	12.4	241	20	272	16	272	10
Eastern Uplands	1.3	188	22	227	14	236	12
Southern Seaboard	2.7	214	28	251	20	269	11
Fruitful Rim	1.4	226	34	268	15	269	15
Basin and Range	0.1	339	17	380	12	379	9
Mississippi Portal	1.2	241	22	277	17	283	11

Note: The coefficient of variation in this application is a measure of the dispersion of the probability distribution of revenue per acre that allows comparisons across populations with different means, and is the standard deviation of revenue per acre divided by the mean revenue per acre. The smaller the coefficient of variation, the lower the dispersion relative to the mean value of the distribution.

Figure 7 maps the percentage change in the coefficient of variation for total gross revenue under the target revenue program versus the traditional-style program. The lighter the color, the greater the decrease in variation offered by the target revenue program. Areas with high levels of correlation between national average yield and county average yield (e.g., the Heartland) tend to show a greater decrease in the coefficient of variation of the target revenue program with respect to the current-style program. In only a few randomly occurring counties does the coefficient of variation in the target revenue program increase over that in the current-style program.

Producer Preferences for Mean Versus Variability of Gross Revenue

If gross revenue plus support payments are a proxy for the annual contribution to a grower's wealth (defined as total gross revenue) and if the only information available on estimated payments under various program alternatives is the mean level of payments, one would expect the eligible producer to prefer the program that offers the greatest mean total gross revenue.[15] But what if the decision criteria involved variability in payments and gross revenue? While the coefficient of variation for total gross revenue may help in determining a preference for mean versus variance, the coefficient is only a measure of dispersion. By itself, it cannot indicate whether a farmer would prefer a program that results in lower mean total gross revenue and lower variability in revenue to one that results in higher mean revenue with higher variability.

Economic theory suggests that producers may balance the mean level of total gross revenue against the variability in the total gross revenue in deciding which support program they would prefer. In particular, almost any individual would view an increase in their mean level of total gross revenue as desirable,

[15]As costs of production do not figure in the calculation of the support payments, we simplify the analysis by using total gross revenue rather than total net revenue.

Figure 7

Difference in coefficient of variation for total gross revenue under target revenue versus traditional-style program payments

In most counties, the volatility of revenue is lower under the target revenue program.

Percent difference
- -89 to -60
- -60 to -30
- -30 to 0
- 0 to 176
- No data
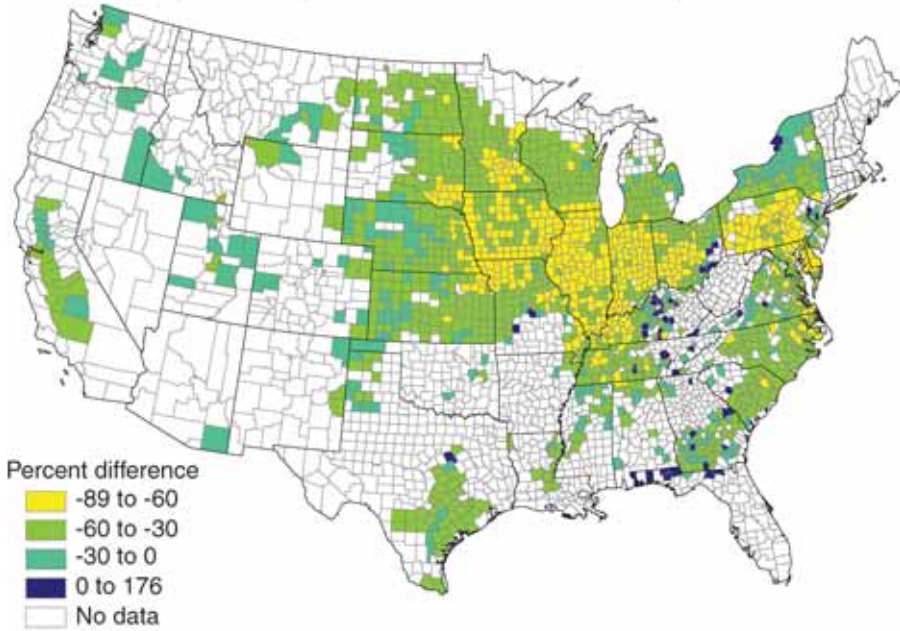
Note: The lighter the color, the lower the coefficient of variation of total gross revenue under the revenue program relative to the traditional style programs.

Source USDA, Economic Research Service.

whereas farmers are typically risk averse and would view increasing variability in total gross revenue as undesirable (Serra et al., 2006).

Serra, Zilberman, and Goodwin (2006) present parameter estimates of the preferences of Kansas farmers for mean level of returns versus variability in returns. To assess whether farmers would prefer the target revenue program over the traditional-style program scenario, we apply that preference structure to the estimated means and variances in county-level total gross revenue from the target revenue-based and current-style payment scenarios. More specifically, for a generic corn farmer in each county (that is, on a corn farm with a yield the same as the county's mean), we calculate the farmer's preference level for expected total gross revenue and variability of total gross revenue. The farmer's preference levels are dictated from an equation in which benefits to the farmer increase as mean revenue increases and decrease as variability of revenue increases.[16] If the estimated preference level is higher under the target revenue program than under the current style program, then a typical farmer in the county is assumed to prefer the former program to the latter. Details of this approach to comparing payment programs are presented in Cooper (2008).

[16]The last section of Appendix B provides technical details of this approach.

The results of the simulation suggest that the target revenue program is preferred over the current-style program by representative corn farmers in 60 percent of counties. While the main purpose of this simulation is to demonstrate that program preferences depend on tradeoffs between mean payments and the variance of payments, results do indicate that farmer preferences for

type of support program have a geographic component (fig. 8). Comparing this pattern with that in figure 3 suggests that farmer preference for program type is more complex than a mere a function of the "natural hedge" between price and yield.

There is a pronounced preference for the target revenue program over the current-style program in the Southern Seaboard, a region where the natural hedge between price and yield is relatively low (figs. 3 and 8). Recall that the national price/national average yield correlation for corn is significantly negative, and that the correlation of corn yields in the Southern Seaboard with national average yield tends to be fairly low (fig. 1). For farmers such as these, the potential benefits of a revenue-based versus price-based program are higher than for farmers whose yields correlate more closely with national aggregate yields, generating more negative correlation between price and farm-level yield.

While the representative farmer shows a preference for the target revenue program in most Illinois, Indiana, and Ohio counties, for Iowa this tends to be the case only in the eastern portion of the State. This suggests that for some Heartland counties, less variable revenue under the target revenue program does not fully compensate for a reduction in mean revenue from the traditional-style program, which tends to over-compensate for revenue losses in areas with more negative price-yield correlations.

Comparing the mean level of returns to variability in returns ignores farmer preferences regarding skewness (shape) of the distribution of revenue. For

Figure 8
Simulation of producer preference for a target revenue versus traditional-style program

Producer preference
■ Prefer pseudo traditional payment
■ Prefer revenue-based payment
□ No data
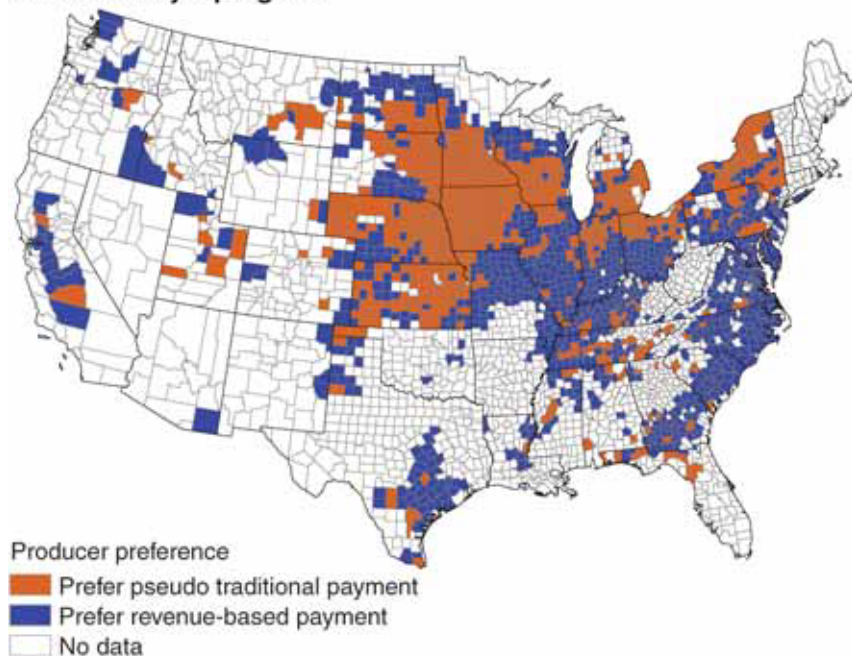
Note: For the simulation, support program parameters are chosen such that average total national payment levels are equal between the two programs when evaluated using historical price and yield data.

Source USDA, Economic Research Service.

example, the entrepreneur may prefer positively skewed revenue distributions because the likelihood of extremely low earnings is smaller (Fisher and Hall, 1969). However, preferences of U.S. farmers for attributes other than the mean and variance of income have received little empirical examination to date.

The stochastic analysis in this report has attempted to outline the implications of the statistical distribution of payments for both the government and producers. Still uncertain is the extent to which government and producer attitudes to risk differ, and whether those differences can be balanced in the final policy outcomes. For example, a support program that reduces the chance of total payments exceeding some ceiling may not be the same program that provides the greatest benefit to producers.