Do Decoupled Payments Stimulate Production?

Estimating the Effect on Program Crop Acreage Using Matching

Abstract. The extent to which decoupled government payments influence agricultural production has important implications for international trade policy. This study applies propensity score matching to farm and ZIP code-level panel data from the U.S. Agricultural Census to evaluate the effect of decoupled payments on crop acreage. The study compares the change between 1997 and 2002 in the program crop acreage of farmers receiving high levels of decoupled payments in 1997 to the change in acreage of similar farmers receiving low levels of payments. For continuing operations, results are consistent with earlier empirical estimates and suggest a small but statistically significant farm-level response to payments. However, for ZIP codes, the estimated supply elasticity is substantially larger than previous estimates and potentially economically significant in scale.

Key words: decoupled payments, supply response, government payments, program crops, trade policy

1. Introduction

Agricultural support payments that cause no or minimal production distortions can be categorized as “green-box”, and therefore except from World Trade Organization restrictions. In the United States, production flexibility contract (PFC) payments under the 1996 Federal Agricultural Improvement Reform Act and direct payments under the 2002 Farm Security and Rural Investment Act are largely “decoupled” because they are based on historical yields and acreages, not current production or prices. The extent to which such decoupled payments significantly affect production and distort trade has emerged as a point of dispute in recent World Trade Organization negotiations (FAO, 2005; Sumner, 2005; USTR, 2004).

In recent years, economists have sought to understand the mechanisms through which decoupled payments could influence agricultural supply and to estimate the extent of this production effect (for reviews see OECD, 2005; Bhaskar and Beghin, 2007). In an influential study, Hennessy (1998) showed that decoupled payments could stimulate production by reducing income variability (an insurance effect) or by reducing the absolute level of risk aversion (a wealth effect) of farmers with decreasing absolute risk aversion. The link between decoupled payments and risk has been the basis of several calibrated models used to estimate of the effect of payments on supply (Young and Westcott, 2000; Anton and Le Mouel, 2004; Sckokai and Moro, 2006). This approach has also been extended to allow for both input and output price risk (Serra, et al., 2006).

Other model-based empirical research has considered the consequences for production of payment-induced increases in land values (Dewbre, Anton, and Thompson, 2001; Gohin, 2006; Roberts, Kirwan and Hopkins, 2003; Kirwan, 2008). It has been hypothesized that decoupled payments increase land values and therefore collateral for obtaining loans, which in turn permits greater investment and output (Roe, Somwaru and Diao, 2003).

Other studies have examined how decoupled payments indirectly affect production through farm household labor market decisions. Ahearn, El-Osta and Dewbre (2006) and El-Osta, Mishra and Ahearn (2004) estimated econometrically the extent to which decoupled payments, which raise household wealth and thereby could influence labor-leisure tradeoffs, affect the allocation of household labor on and off the farm. Key and Roberts (forthcoming) show that farmers with preferences for farm (versus off-farm) work could respond to higher decoupled payments by decreasing off-farm labor and increasing farm labor. [Should we cite our meetings paper?]

Even when current payments are decoupled from production, farmers may have an incentive to respond to anticipated but uncertain policy changes by altering their acreage or input decisions to maximize their expected future stream of payments (Lagerkvist, 2005; Sumner, 2003). The 2002 Farm Bill, which extended and increased the fixed decoupled payments of the 1996 Act, gave producers an opportunity to update their base acreage and yields and allowed them to include soybean acreage in their base. Hence, prior to 2002, farmers may have altered their acreage decisions in anticipation of the base updating, even though current payments were decoupled from current production.

The soundness of predictions derived from simulations using calibrated models rests on the validity of the coupling mechanisms underpinning the models and the accuracy of the coefficients used for calibration. An alternative empirical approach is to estimate the effect of decoupled payments on agricultural production econometrically using survey data. However, this approach brings multiple conceptual and practical challenges. First, decoupled payments generally originate from agricultural programs that are open to all program crop farmers and program participation is often nearly universal. This makes it difficult or impossible to distinguish between a treatment and control group – a prerequisite for a standard program evaluation. Second, changes in agricultural policies that cause changes in farm-level payments generally occur simultaneously across the nation. This makes event studies infeasible and can make it difficult to distinguish between the effects of changes in policy versus changes in prices, technology or other time-varying factors. Third, it is usually not clear what drives variation in payments across observationally similar farms. This leaves open the possibility that unobservable factors could be associated with both program participation (or payment levels) and supply response, resulting in sample selection bias. Consequently, analyses require a high degree of confidence that there are no missing factors correlated with both agricultural supply and government payments.

Goodwin and Mishra (2006) made an important contribution to the econometric literature by estimating the effect of coupled and decoupled payments on acreage decisions using a cross-sectional survey that asked how much land farmers planted in the current and previous periods. The researchers estimated a linear relationship between current payments per acre and current crop acreage conditioning on past acreage and other factors. As the authors recognized, a potential problem with this approach arises when there are unobservable factors that influence both current plantings and current payment levels. For example, a problem could arise if farmers with particularly high yields and therefore high revenues in the previous period rent or buy more land of relatively high quality in the current period. Since land quality is correlated with payments per acre, this would cause current payments per acre and current crop acreage to be spuriously correlated. If land quality is unobservable, estimates of the effect of payments on acreage would be biased upward.

The current study uses a large panel data set and matching econometric technique to address some of the empirical challenges associated with estimating the effect of decoupled payments on crop acreage. Data are from the U.S. Agricultural Census which include information on the amount of land allocated to particular crops in 1997 and 2002 and decoupled payments in 1997 – these payments are comprised mainly of decoupled PFC payments, as discussed later in the paper. A matching estimation based on the propensity score is used to compare the change between 1997 and 2002 in program crop acreage of farms with a high level of 1997 payments per acre to similar farms with a low level of payments. The analysis is conducted for both individual continuing farms, and for a representative farm at the ZIP-code level. The ZIP-code level analysis is useful because it accounts for entering and exiting farms in addition to continuing farms.

The paper offers several contributions to the empirical literature on decoupled payments. First, because the data are drawn from a farm-level panel, it is possible to estimate how variations in payments per acre affect *subsequent* plantings. While unobservable factors could cause current payments to be correlated with contemporaneous plantings, it is much less likely that unobservable factors affecting current payments are correlated future plantings. Second, the comprehensive sample of the Census data – which includes essentially all U.S. producers of program crops – minimizes measurement errors associated with sample design and response rates and allows for regional fixed effects. The near-population sample of farmers allows for comparisons across very similar farms and ZIP codes, which increases the precision of the estimators and reduces potential sample selection bias. Third, the matching methodology provides a way to estimate the effect of payments on acreage that does not require distributional and functional form assumptions about the relationship between the treatment and outcome. The findings should help inform debates on future agricultural policy reforms and trade negotiations.

2. Empirical Model

Farm operators can be assumed to make land allocation decisions to maximize their expected utility. As discussed in the review of the literature in the previous section, attitudes towards risk, expectations about future policy changes, or preferences for farm versus non-farm work could all enter the agent's optimization problem. In addition, acreage decisions could be constrained by imperfections in credit, labor, land or other markets. The complexity of the optimization problem rules out deriving a feasible structural acreage response equation. Instead, we posit a general reduced-form equation describing the change in program crop acreage over five years:

$$(1) \qquad A_t = f\left(A_{t-5}, G_{t-5}, \mathbf{P}_{t-5}, w_{t-5}, \mathbf{Z}_{t-5}\right),$$

where A_t is program crop acreage in year *t* (2002), A_{t-5} is program crop acreage in the previous census year (1997), and G_{t-5} is decoupled payments per acre of program crops. The vector $\mathbf{P}_{t-5}$ includes base-year factors affecting expected future profits (e.g., prices and idiosyncratic factors affecting the production technology, such as land quality), w_{t-5} is household wealth (which allows for the possibility that risk preferences vary with wealth, which in turn could affect production decisions, and allows for the fact that credit constraints could vary with wealth), and $\mathbf{Z}_{t-5}$ includes other factors that could influence land use decisions, such as the operator's age.

The estimated equation based on (1) is:

$$(2) \qquad A_t = f\left(A_{t-5}, D_{t-5}, C_k, AC_{j,t-5}, Y_{t-5,} S_{t-5}, Age, Age^2\right) + \varepsilon\,,$$

where D_{t-5} is a discrete variable indicating whether the operator received high or low payments per acre of program crops in 1997 (defined by being above or below the sample median). A discrete treatment variable is used because it facilitates implementation of the matching estimation approach and it is less likely to be influenced by outliers than a continuous variable would be. Fixed effects C_k for each ZIP code *k* are included to capture price variation across regions. Price variation across farms within each ZIP code is likely to be small. In addition, the response of the dependent variable, total program crop acreage, to within-ZIP code price variation is likely to small. ZIP-code fixed effects also control for time-invariant unobserved factors that affect all farms within each ZIP-code. Within-ZIP-code land quality controls include acreage of crop *j* [what is crop *j* ?] $AC_{j,t-5}$ and the farm's corn yield Y_{t-5}.[1] Total agricultural sales S_{t-5} is used as a proxy for wealth.[2] The operator's age and age-squared are included to capture factors associated with lifecycle decisions affecting land use, including wealth and experience.

Matching

[1] As discussed in the Data section, all farms in the sample grew corn in 1997.

[2] The census of agriculture does not provide a good measure of wealth. One possible proxy is the value of land and building on the operation, but this is only available on the "long form" census questionnaire, which was distributed to only about one-third of all operators.

The testable null hypothesis is that mostly-decoupled 1997 agricultural payments per acre had no influence on subsequent program crop acreage. We first test this hypothesis by estimating (2) assuming a linear functional relationship between the independent variables. A problem with the linear regression model is that it imposes strict distributional and functional form assumptions about the relationship between the treatment and outcome. Matching provides a non-parametric approach that does not require functional form assumptions in the outcome equation. Matching techniques have been shown to outperform standard nonexperimental estimators such as regression and latent variable selection models, relative to an experimental benchmark (Dehejia and Wahba, 1999, 2002; Smith and Todd, 2005).

The empirical approach is to estimate the effect of the binary treatment (high/low government payments per acre) on the continuous outcome, program crop acreage harvested. Dropping the time subscript to simplify notation, let A_{0i} and A_{1i} denote two potential outcomes: A_{0i} is the acreage of individual *i* not exposed to the treatment and A_{1i} is the acreage if exposed to the treatment. The sample-average treatment effect for the treated (the additional crop acreage harvested for those receiving high payments compared to what it would have been had they received low payments) is given:

(3) $$\tau_T^S = \frac{1}{N_T} \sum_{i:D_i=1} \left(A_{1i} - A_{0i} \right),$$

where N_T is the number of treated individuals and the where $D_i = \{0,1\}$ is the treatment indicator.

As with most policy analyses, the sample-average treatment effects cannot be computed because we only observe one of the two possible outcomes for each individual. For example, if an individual received the treatment, then we observe A_{1i}, but we do not observe what the outcome would have been had the individual not received the treatment (A_{0i}). The basic idea behind the matching estimator is to estimate A_{0i} using the average outcomes of similar individuals who were not exposed to the treatment. Analogously, if we observe the outcome for an individual who did not receive a treatment, then we can estimate A_{1i} using the average of outcome of similar individuals who were exposed.

Matching estimators compare outcomes across pairs of "similar" treated and control units. Ideally, pairs would be matched on all relevant observable variables. In practice, matching subjects on a vector of characteristics is not computationally feasible with a large sample when the number of characteristics is large. Propensity score matching is a method that summarizes the characteristics of each observation into a single index (the propensity score) to make matching feasible. Since Rosenbaum and Rubin (1983) proposed matching individuals based on their propensity score – that is, on their probability of receiving the treatment – these methods have been widely used in the evaluation of economic policies and medical trials.

The propensity score is defined as the conditional probability of receiving a treatment given pre-treatment characteristics *X*: $P(X) \equiv \Pr(D = 1 | X)$. A probit (or other standard probability model) can be used to estimate the propensity score:

(4) $$\Pr(D_i = 1 | X_i) = \Phi(h(X_i)),$$

where Φ denotes the normal c.d.f. and $h(X_i)$ is a function of the covariates.

Rosenbaum and Rubin (1983) showed that if exposure to treatment is random within cells defined by *X* then it is also random within cells defined by the values of $P(X)$. Hence, if the propensity score is known, then the average treatment effect can be estimated as the expected difference in outcomes between individuals receiving the treatment and not receiving the treatment, conditional on the having the same propensity score. The average effect of treatment on the treated (ATT) can therefore be estimated as:

$$\tau^T = E\left[E\left[A_{1i} \middle| D_i = 1, P(X_i)\right] - E\left[A_{0i} \middle| D_i = 0, P(X_i)\right] \middle| D_i = 1\right] \tag{5}$$

Since the propensity score is a continuous variable, the probability of having two observations with the same value of $P(X)$ is zero. Consequently, an estimate of the propensity score is not sufficient to estimate (5). We use two alternative methods to overcome this problem: stratification matching (or blocking) and nearest neighbor matching. With stratification matching the range of the propensity score is divided into intervals such that within each interval the treated and control observations have the same average propensity score. Then, within each interval the difference between the average outcomes of the treated and controls is computed. The ATT is weighted average of the average differences in each block, with the weights given by the frequency of the treated observations.

The nearest neighbor method matches each treated observation with the control unit with the closest propensity score. Once each treated unit is matched, the difference between the outcome of the treated and untreated units is computed, and the ATT is the average of these N_T differences.

For both methods, the ATT estimator is essentially the difference between two sample means, so the variance calculation for the estimator is straightforward. Details on the implementation of the stratification and nearest neighbor methods and the variance calculations are given in Becker and Ichino (2002).

Exogeneity of the Treatment

A key assumption underlying this and most econometric analyses of the effect of payments on acreage is that after controlling for observable differences between farms, differences in payments per acre across farms are not associated with unobservable factors that influence the change in program crop acreage. We maintain that this assumption is reasonable in the current study for two main reasons. First, a substantial portion of the variation in payments across observationally similar farms likely results from exogenous ex-ante variation in government payments caused by random variations in base yields. Base yields were determined by averaging realized yields between 1981 and 1985. While yields are tied to land quality, they also vary widely from year-to-year and across space, due to weather outcomes. Indeed, summary statistics reported by Roberts, Key and O'Donoghue (2006) indicate that field-level yields are typically 30% to 50% above or below their mean, and a county-level yield shock accounts for only about half the field-level variability. Thus, some variation in base yields is likely random.[3]

Second, as discussed in the next section, the sample is large and, on average, there are relatively small differences between the characteristics of operations with high and low payments per acres even before matching. Large samples having substantial overlap in characteristics between the treatment and control are particularly well suited for matching techniques (Shandish, Cook, and Campbill, 2002). Matching allows for direct comparisons between very similar treatment and control groups, which further reduces the likelihood that unobservable factors could result in sample selection bias (Altonji, Elder, and Taber, 2005).

3. Data

Farm-level data are from the Census of Agriculture maintained by the USDA National Agricultural Statistics Service (NASS).[4] The Census collects data on farm and operator characteristics every five years from essentially all farms in the country. NASS assigns each respondent a unique Census File Number (CFN) to track the farm, ranch, or other agricultural entity controlled or operated by the individual filing the census. Since every farm operator is required (by law) to respond to the Census, it is possible to track operations across time as long as they remain in business.

To select a homogenous group of operations, this study examines operations in the "Heartland" – a relatively homogenous geographical region defined at the county-level by the USDA.[5] The Heartland includes all counties in Illinois, Indiana, and Iowa, and some bordering counties in Kentucky, Minnesota, Missouri, Nebraska, Ohio, and South Dakota. Data are organized in two ways: 1) individual continuing operations, and 2) all operations aggregated at the ZIP-code level.

Continuing operations

[3] We do not observe base yields or base acres in the Census data so we cannot use these as instrumental variables.
[4] More information about the Census of Agriculture can be found at: http://www.agcensus.usda.gov/.
[5] See www.**ers**.usda.gov/publications/aib760/aib-760.pdf for a map of the region and more details.

Continuing operations are defined as those that remained in business (with any level of cropland) in 2002. To increase the homogeneity of the sample, we consider only on farms harvested at least 50 acres of "program crops" (corn, wheat, barley, oats, rice, cotton, and sorghum) in 1997. Soybeans acres were not included in program-crop acres because prior to the 2002 Act soybeans and other oilseeds were not eligible for direct payments. In 1997, a total of 475,486 farmers responded to the Census in 1997 in the Heartland. Of these, 156,884 harvested at least 50 acres of program crops in 1997, and 106,431 of these remained in business for at least five years. To further reduce sample heterogeneity and to permit the use of corn yields as a measure of land quality, we dropped the 1,978 farms (about 1.9%) that did not grow corn for grain, leaving a final sample of 104,453 operations.

For this study, government payments are defined as total payments received for participation in Federal farm programs (not including Commodity Credit Corporation loans or crop insurance payments) net of payments received for participation in the Conservation Reserve Program and the Wetlands Reserve Program.[6] In 1997, these payments were derived almost entirely from Production Flexibility Contracts, which were tied to historically enrolled contract acreage, not current plantings (USDA, 2008). Because the Census data does disaggregate the source of payments by individual crop, it is necessary to use total program crop acreage as the dependent variable, rather than individual program crop acreage.

ZIP code regions

[6] In the 1997 census respondents were asked for the "total amount received for participation in Federal farm programs (not including CCC loans)." Respondents were also asked to provide "how much was received for participation in the Conservation reserve program and Wetlands Reserve Program (CRP and WRP)?" The latter was subtracted from the former to obtain the measure of payments used in this study.

An analysis of continuing farm operations cannot account for the acreage decisions of farms entering or exiting production between Census periods so it cannot provide an estimate of how payments influence total program crop acreage. To explore the effect of decoupled payments on aggregate program crop acreage we conduct an analysis at a regional level and include all respondents to the 1997 or 2002 Censuses who had operations located in the Heartland.

A regional analysis could be conducted at the ZIP code, county, or state level. ZIP codes are the smallest geographic unit where farms can be located with the data, so they provide the maximum number of observations and cross-sectional variation in the dependent and independent variables. In the Heartland, there are 9,718 ZIP codes compared to 544 counties. Although a very small fraction of ZIP codes change over time, most changes have occurred in relatively urban areas undergoing rapid population growth and where agriculture is less prevalent, which mitigates this potential problem in this analysis.

In the Heartland, there were 5,134 ZIP codes in which some operations reported harvesting any of the program crops in 1997 and 2002. To increase estimate precision, we kept only those 4,197 ZIP codes having at least 15 operations. To allow for the use of the average corn yield as a control for land quality, we dropped the 13 ZIP codes where no corn for grain was harvested, resulting in a final sample of 4,186 ZIP codes.

Table 1 provides a comparison of the mean values of variables used in the analysis for the payments per program-crop acre categories. For the farm-level program acreage analysis, the treatment group consists of those that received payments that were above $26.82 per acre of program crops harvested in 1997. For the ZIP-code analysis, the cutoff is $28.19.

For most operator and operation characteristics, differences between the averages of the treatment and control groups are not large, and in most cases would not be considered economically significant (though many differences are statistically significant). However, the differences in crop acreage illustrated in the first six rows of the table suggest that payments could play an economically significant role in acreage response. Between 1997 and 2002, the treatment (high payments per acre) farms increased land harvested in program crops by an average of 14.8 acres compared to a decline of 9.0 acres for the control group. For ZIP codes, the treatment group harvested 10.5 additional acres in 2002 while the control group harvested 0.4 additional acres.

4. Results

Table 2 presents the results of the linear regression models with 2002 program crop acres as the dependent variable. Since we control for 1997 program crop acres, the positive (negative) coefficients can be interpreted as increasing (decreasing) program crop acreage between 1997 and 2002. For continuing operations, operators who had higher sales, more farmland, and harvested more acres of corn, soybeans, and oats in 1997 were significantly more likely to increase their program crop acreage between 1997 and 2002. Program crop acreage expanded with age, but at a decreasing rate. Corn grain yields, a proxy for land quality, was positively correlated with an expansion of program crop acreage. The results for ZIP codes are similar to the individual farm results except that acres of corn and soybean harvested are not significant. Of greatest relevance for this analysis, the regression results indicate that being in the high payments-per-acre category was associated with an increase of 14.4 acres in program crops between 1997 and 2002 for continuing operations, and an increase of 6.8 acres in aggregate.

Table 3 presents the coefficients of the probit models of the propensity scores. The results indicate many statistically significant variables, but the model explains only a small share of the variation in the dependent variable (pseudo R-squares are 0.07 and 0.10 for the farm and ZIP code analyses, respectively). This is consistent with our maintained assumption that after controlling for observables, differences in payments per acre in 1997 are caused largely by random factors that determined base yields in 1981-1985. In other words, if random events caused variation in payments per acre in observationally similar farms, then we could not expect to explain a farm's placement in the high or low payments category with observables.

The propensity score algorithm splits the sample into blocks until the average propensity score of the treated and control units do not differ (for details, see Becker and Ichino (2002)). For the continuing operations, the same average propensity scores across blocks are achieved with 21 blocks; for the regional analysis with 10 blocks. Table 4 illustrates the final distribution of the treated and control observations across these blocks.

Table 5 presents the estimated average treatment effect for the stratification and nearest neighbor matching models, for continuing operations and ZIP codes. The coefficients can be interpreted in the same way as the indicator variable in the linear model. For continuing operations, the estimates imply that the treatment (high payments per acre) caused farmers to increase their program crop acreage between 1997 and 2002 by 11.9 – 15.7 acres (4.0 – 5.2 percentage points) more than they would have had they received low payments per acre. In terms of the aggregate effect, the ZIP code estimates indicate that the treatment increased program crop acreage by 10.6 – 15.9 acres (7.3 – 10.9 percentage points).

For continuing operations, the treatment represents an average 307% increase in payments per harvested acre of program crops (from $11.6 to $47.2 per acre), which implies a response elasticity of 0.013 – 0.017. These farm-level results are generally consistent with earlier studies. Goodwin and Mishra (2006), who pooled several cross-sectional surveys of farms in the Heartland from 1998-2001, found acreage elasticities for corn and soybeans of about 0.01 to 0.03 (p.87) for decoupled PFC payments.

The estimated aggregate effect is much larger than the farm-level results. For ZIP codes, the treatment represents a 94% increase in payments per acre, which implies an aggregate elasticity estimates of 0.077 – 0.116. While the farm-level analysis included only those farms that survived for 5 years, the aggregate analysis compared the change in acreage for all farms in each region – this also includes farms that were in business in 1997 that exited before 2002, and those that entered production between census periods. There is empirical evidence that payments have a positive effect on the rate of farm survival and that there is strong positive relationship between the duration of farm survival and farm size (Key and Roberts, 2006, 2007). It follows that farmers who remain in business longer because of high payments are likely to operate larger farms and harvest more acreage after five years than would entering farmers who would have replaced them. Hence, the results suggest that the positive effect of payments on program crop acreage occurs, in part, because payments increase the duration of farm survival, not simply because payments allow continuing farms to expand production.

5. Conclusion

Understanding how and to what extent decoupled government payments influence agricultural production has been the focus of recent economic research. However, there have been relatively few econometric examinations of the supply response to decoupled payments because of associated methodological challenges. This paper uses a large panel data set and matching technique to address these challenges. The approach compares the change between 1997 and 2002 in acreage allocated to program crops by farmers receiving high levels of decoupled payments in 1997 to the change in acreage of similar farmers receiving low levels of payments. The panel aspect of the data allows for conditioning current acreage decisions on past payments, rather than on current payments, which reduces the likelihood that unobservable factors correlated with payments and acreage response bias the estimates. The large sample of homogenous farms and the matching methodology allow for comparisons across farms that are observationally very similar, which also reduces the potential for sample selection bias. Finally, the matching technique provides a flexible empirical methodology that does not require distributional and functional form assumptions about the relationship between the treatment and outcome.

For continuing program crop producers in the Heartland, the results imply a payment-acreage elasticity between 0.013 and 0.017, which is consistent in magnitude with prior empirical estimates and suggests a small but statistically significant farm-level effect. However, for the ZIP codes analysis, which includes all farms, the estimated aggregate elasticity of 0.077 to 0.116 is substantially larger than previous estimates and is potentially economically significant. The aggregate results imply that a 50% change in decoupled payments would change the supply of all program crops by about 5% over five years. It is feasible that a supply response of this magnitude could result in meaningful changes in U.S. exports and international commodity prices.

The finding that the acreage response elasticity was larger in aggregate compared than for continuing operations suggests that decoupled payments affect total production partly by allowing farmers to remain in business longer. Farms that remain in business are likely to be larger and harvest more acreage than would new entrants. One possible explanation for this pattern is that decoupled payments, which are not tied to production or prices, provide additional liqudity and buffer-stock of land-based wealth, which serve as insurance in bad years and thereby reduce the likelihood of financial insolvency.

References

Ahearn, M., El-Osta, H., and Dewbre, J. (2006). “The impact of coupled and decoupled government subsidies on off-farm labor participation of U.S. farm operators. *American Journal of Agricultural Economics* 88(2): 393–408.

Altonji, J., T.E. Elder, and C.R. Taber, 2005. "Selection on Observed and Unobserved Variables: Assessing the Effectiveness of Catholic Schools," *Journal of Political Economy*, University of Chicago Press, vol. 113(1): 151-184.

Anton, J. and Le Mouel, C. (2004). “Do counter-cyclical payments in the 2002 US Farm Act create incentives to produce?” *Agricultural Economics* 31(2-3): 277–284.

Bhaskar, Arathi and Beghin, John C., “How Coupled are Decoupled Farm Payments? A Review of Coupling Mechanisms and the Evidence” Working paper. (September 26, 2008). Available at: http://ssrn.com/abstract=1274060

Becker, S. and A. Ichino, “Estimation of Average Treatment Effects Based on Propensity Scores,” *The Stata Journal* 2:4 (2002), 358-377.

Dehejia, R. and Wahba, S. (1999) “Causal effects in nonexperimental studies: reevaluating the evaluation of training programs.” *Journal of the American Statistical Association* 94: 1053-1062.

Dehejia, R., Wahba, S. (2002), "Propensity score matching methods for non-experimental causal studies", *The Review of Economics and Statistics*, Vol. 84 No.1, pp.151-61.

Dewbre, J., Anton, J., and Thompson, W. (2001). The transfer efficiency and trade effects of direct payments. *American Journal of Agricultural Economics* 83(5): 1204–1214.

El-Osta, H., Mishra, A., and Ahearn, M. (2004). “Labor supply by farm operators under ‘decoupled’ farm program payments.” *Review of Economics of the Household* 2: 367–385.

Food and Agriculture Organization. “The effect of direct payments of the OECD countries in world agricultural markets. Evidence from partial and general equilibrium frameworks*” FAO Commodity and Trade Policy Research Working Paper No. 12. August, 2005.*

Hennessy, D. A. “The Production Effects of Agricultural Income Support Policies under Uncertainty.” *American Journal of Agricultural Economics* 80(1998): 46-57.

Gohin, A. (2006). Assessing CAP Reform: Sensitivity of modeling decoupled policies. *Journal of Agricultural Economics* 57(3): 415–440.

Goodwin, B. K. and A. K. Mishra. “Are ‘Decoupled’ Farm Payments Really Decoupled? An Empirical Evaluation.” *American Journal of Agricultural Economics* 88(1) (February 2006): 73-89.

Key, N. and M.J. Roberts. “Nonpecuniary Benefits to Farming: Implications for Supply Response to Decoupled Payments.” *American Journal of Agricultural Economics*. (forthcoming).

Key, N. and M.J. Roberts. "Do Government Payments Influence Farm Size and Survival?" *Journal of Agricultural and Resource Economics*, Vol. 32, No. 2 (August 2007): 330-348.

Key, N. and M.J. Roberts. “Government Payments and Farm Business Survival.” *American Journal of Agricultural Economics*. 88(2) (May 2006): 382-392.

Lagerkvist, C. (2005). “Agricultural policy uncertainty and farm level adjustments - the case of direct payments and incentives for farmland investment.” *European Review of Agricultural Economics* 32(1): 1–23.

OECD, “A Review of Empirical Studies of the Acreage and Production Response to US Production Flexibility Contract Payments under The FAIR Act and Related Payments under Supplementary Legislation”, Directorate For Food, Agriculture And Fisheries Committee For Agriculture, March 2005. Available at: http://www.oecd.org/dataoecd/15/15/34997377.pdf

Roe, T., A. Somwaru and X. Diao (2003), “Do Direct Payments Have Intertemporal Effects on US Agriculture?”, in C.B. Moss and A. Schmitz (eds.), *Government Policy and Farmland Markets*, pp. 115-139, Ames, Iowa State Press.

Roberts, M.J., N. Key, and E. O’Donoghue. “Estimating the Extent of Moral Hazard in Crop Insurance Using Administrative Data” *Review of Agricultural Economics*. 28(3) (2006): 381-390.

Rosenbaum, P. and D. Rubin, “The Central Role of Propensity Score in Observational Studies for Causal Effects,” *Biometricka* 70 (1983), 41-55.

Sckokai, P. and Moro, D. (2006). “Modelling the reforms of the Common Agricultural Policy for arable crops under uncertainty.” *American Journal of Agricultural Economics* 88(1): 43–56.

Serra, T., Zilberman, D., Goodwin, B., and Featherstone, A. (2006). “Effects of decoupling on the mean and variability of output.” *European Review of Agricultural Economics* 33(3): 269–288.

Shadish, W.R., T.D.Cook, and D.T.Campbell, (2002). *Experimental and Quasi-Experimental Designs for Generalized Causal Inference*. Boston: Houghton-Mifflin.

Smith, J. A. and Todd, P. E. (2005): Does matching overcome LaLonde's critique of nonexperimental estimators? *Journal of Econometrics* 125: 305-353.

Sumner, D. (2003). Implications of the US Farm Bill of 2002 for agricultural trade and trade negotiations. *Australian Journal of Agricultural and Resource Economics* 47(1): 99–122.

Sumner, D. "Boxed In: Conflicts Between U.S. Farm Policies and WTO Obligations" *Trade Policy Analysis* No. 32. CATO Institute. December 2005.

U.S. Department of Agriculture, Economic Research Service, Farm and Commodity Policy: Government Payments and the Farm Sector, http://www.ers.usda.gov/Briefing/FarmPolicy/gov-pay.htm, Accessed April 2008.

U.S. Trade Representative. "Panel Issues Mixed Verdict in Cotton Case" Available at: http://www.ustr.gov/Document_Library/Press_Releases/2004/September/WTO_Panel_Issues_Mixed_Verdict_in_Cotton_Case.html Press Release. September 2004.

Young, C.E. and P.C. Westcott. "How Decoupled Is U.S. Agricultural Support for Major Crops?" *American Journal of Agricultural Economics* 82(2000): 762-767.

Table 1. Descriptive Statistics and Test of Equality of Means

	Continuing operations		ZIP codes	
	Low payments per acre (control)	High payments per acre (treatment)	Low payments per acre (control)	High payments per acre (treatment)
1997 program crops (acres harvested)	300.7	273.8***	145.9	124.5***
2002 program crops (acres harvested)	291.7	288.6	146.3	135.1***
Change in program crops (acres harvested)	-9.0	14.8***	0.4	10.5***
Government payments ($)	3,783	12,061***	3,242	4,433***
Gov. payments per program crop acre ($/acre)	11.6	47.2***	21.8	42.2***
Sales ($1000)	204.9	227.4***	113.1	124.9**
Operator age (years)	49.7	50.9***	53.2	53.1
Land in farm (acres)	665.5	642.1***	380.1	342.4**
Corn acres (field corn for grain, harvested)	261.6	254.4***	122.2	116.1***
Soybean acres (harvested)	233.5	237.8***	124.6	108.7***
Oat acres (harvested)	3.5	2.5***	1.2	1.4***
Hay acres (all types, harvested)	24.2	21.8***	17.2	16.4
Silage acres (corn or sorghum, harvested)	4.5	5.7***	2.4	2.5
Hogs (hog and pig inventory, head)	132.7	168.7***	65.7	117.5***
Cattle (cattle and calf inventory, head)	53.0	56.0**	30.9	31.8
Corn grain yield (bushels per acre harvested)	117.1	126.6***	109.4	116.8***
Observations	52,227	52,226	2,093	2,093

Notes: *** p-value < 0.001, ** p-value <0.01, * p-value <0.05, where p-values correspond to the two-tailed significance probability under the null hypothesis of equal means from a pooled sample. All variables are for 1997 except where noted. Program crops include corn, wheat, barley, oats, rice, cotton, and sorghum. Government payments are total payments received for participation in Federal farm programs (not including Commodity Credit Corporation loans or crop insurance payments) net of payments received for participation in the Conservation Reserve Program and the Wetlands Reserve Program. Land in farm is total land owned + rented in – rented out.
Source: 1997 and 2002 Census of Agriculture.

Table 2. Linear Regression Estimates (Dependent Variable: 2002 Program Crop Acres)

	Continuing operations		ZIP codes	
	Coefficient	Std. Err.	Coefficient	Std. Err.
Intercept	218.912	17.702***	259.507	65.647***
High payments per acre (treatment)	14.375	1.399***	6.826	1.973***
1997 program crop acres	0.532	0.010***	0.353	0.040***
Sales	0.037	0.003***	0.041	0.003***
Age of operator	-8.940	0.347***	-9.074	2.527***
Age-squared	0.064	0.003***	0.060	0.023***
Land in Farm	0.038	0.003***	0.346	0.018***
Corn acres	0.197	0.010***	0.028	0.031
Soybean acres	0.200	0.006***	0.004	0.029
Oat acres	0.096	0.059***	1.329	0.471***
Hay acres	-0.097	0.016	-1.412	0.101***
Silage acres	-0.037	0.034	-1.429	0.243***
Hog inventory	0.002	0.001	-0.013	0.003***
Cattle inventory	-0.036	0.005***	-0.063	0.033*
Corn yield	0.319	0.029***	0.230	0.057***
ZIP code fixed effects	yes		no	
R-square	0.65		0.75	
Observations	104,453		4,186	

Note: *** p-value < 0.001, ** p-value <0.01, * p-value <0.05.

Table 3. Probit Model Estimates (Dependent Variable: High Payments per Acre Program Crops)

	Continuing operations		ZIP codes	
	Coeff.	Std. Err	Coeff.	Std. Err.
Intercept	-1.981	.0906***	-6.908	1.434673***
1997 program crop acres	-0.00174	7.03E-05***	-0.01729	0.001201***
Sales	0.000133	1.79E-05***	0.000789	2.78E-04***
Age of operator	0.0387	0.00210***	0.235	0.0550***
Age-squared	-0.00031	2.01E-05***	-0.00232	0.000505***
Land in Farm	0.000173	1.98E-05***	0.00208	0.000431***
Corn acres	0.000718	6.97E-05***	0.0109	0.000986***
Soybean acres	0.00061	3.54E-05***	-0.00076	0.000686
Oat acres	-0.00505	0.000409***	0.0265	0.0104**
Hay acres	0.000391	9.72E-05***	-0.0048	0.00228**
Silage acres	0.00220	0.000201***	0.00309	0.00519
Hog inventory	1.44E-05	6.86E-06**	0.000363	0.000145**
Cattle inventory	-0.00014	2.86E-05***	-0.00132	0.000744*
Corn yield	0.00514	0.00017***	0.0152	0.00124***
ZIP code fixed effects	yes		no	
Pseudo R-square	0.07		0.10	
Observations	104,453		4,186	

Note: *** p-value < 0.001, ** p-value <0.01, * p-value <0.05.

Table 4. Distribution of Treated and Control for Each Propensity Score Block

Lower bound of propensity score block	Control	Treatment	Total
Continuing operations			
0	598	40	638
0.1	635	99	734
0.15	1,339	258	1,597
0.2	2,472	629	3,101
0.25	3,783	1,314	5,097
0.3	4,906	2,244	7,150
0.35	5,329	3,261	8,590
0.4	5,911	4,424	10,335
0.45	6,302	5,852	12,154
0.5	3,245	3,359	6,604
0.525	2,984	3,676	6,660
0.55	2,960	3,731	6,691
0.575	2,612	3,900	6,512
0.6	2,318	3,725	6,043
0.625	1,910	3,362	5,272
0.65	2,635	5,693	8,328
0.7	1,436	3,838	5,274
0.75	607	2,023	2,630
0.8	207	745	952
0.9	20	39	59
0.95	18	14	32
Total	52,227	52,226	104,453
ZIP codes			
0	222	34	256
0.2	211	53	264
0.3	326	163	489
0.4	508	337	845
0.5	294	261	555
0.55	212	277	489
0.6	157	345	502
0.65	85	316	401
0.7	69	268	337
0.8	9	39	48
Total	2,093	2,093	4,186

Table 5. Matching Model Estimates

Matching method	Continuing operations		ZIP Codes	
	ATT	Std. Err.	ATT	Std. Err
Stratification matching	11.86	2.42***	10.65	3.69**
Nearest neighbor matching	15.72	3.25***	15.86	5.12**

Note: *** p-value < 0.001, ** p-value <0.01, * p-value <0.05. The ATT is the average treatment effect of treated (high payments per acre for those with high payments).