Agricultural Subsidy Incidence: Evidence from Commodity Favoritism

Nathan P. Hendricks* Krishna P. Pokharel†

March 24, 2016

Abstract

We use county-level data in the United States to estimate the incidence of direct payments on cash rental rates. Direct payments were fixed subsidies not tied to price or production—thus, standard theory suggests direct payments should be fully reflected in rents. Our econometric model exploits variability in direct payments due to variation in the proportion of cropland with cotton or rice base acres while controlling for expected market returns. Cotton and rice base acres received substantially larger direct payments, arguably because cotton and rice—historically produced in the South—are politically favored compared to commodities produced in other regions. Estimates from two-stage least squares indicate that roughly \$0.81 of every dollar of direct payments accrues to landlords through higher rental rates in the long run. We also construct revised standard errors that account for potential violations of the exclusion restriction. Most previous literature exploits changes in subsidies over time or differences in subsidies across areas producing the same set of commodities. Our estimate of the incidence of direct payments on rental rates is larger than most previous literature because we exploit large, persistent differences in subsidies.

Keywords: Incidence, agricultural subsidies, decoupled payments, rental rates.
JEL codes: Q18, H22.

*Hendricks is an assistant professor in the Department of Agricultural Economics at Kansas State University. Department of Agricultural Economics, Kansas State University, Manhattan, KS 66506. nph@ksu.edu.

†Pokharel is a PhD candidate in the Department of Agricultural Economics at Kansas State University.

Political support for government interventions in the market often depends as much on the distribution of benefits and costs as the overall change in social welfare. In recent years, the beneficiaries of agricultural subsidies in the United States have come under increased scrutiny due to the pressure to reduce budgetary expenditures in the Farm Bill. The United States spent roughly $7.6 billion annually between 2000 and 2013 on agricultural commodity subsidies (U.S. Department of Agriculture 2016).[1] One concern is that non-operator landowners may benefit from these agricultural subsidies—even though the subsidies are generally paid directly to farm operators. Non-operator landowners may capture a portion of the subsidies by adjusting rental rates.

Economists have long recognized that the economic incidence of government subsidies differs from the initial recipient of such subsidies. Standard economic theory predicts that non-operator landowners capture all of a purely decoupled subsidy but only capture a portion of a subsidy directly tied to production (Floyd 1965; Alston and James 2002). Direct payments in the United States (2002–2014) were one example of a fixed subsidy that was not tied to current production or price.[2] There are, however, several reasons why landowners may not capture the entire direct payment. First, tenants are often related to the landowner (Schlegel and Tsoodle 2008), so some rental rates may not reflect the competitive rate (Perry and Robison 2001; Tsoodle, Golden, and Featherstone 2006).[3] Second, direct payments are not purely decoupled (e.g., Hennessy 1998; Just and Kropp 2013; Hendricks and Sumner 2014). Third, tenants may exercise market power in the rental market (Kirwan 2009; Kirwan and Roberts Forthcoming).

Most studies examining the impact of government payments on rental rates find that less than $0.50 of every dollar of subsidies is captured by changes in the rental rate (Kirwan 2009;

[1] In this calculation, we only include production flexibility contract, fixed direct, ACRE, counter-cyclical, and loan deficiency payments. Expenditures are much larger after accounting for crop insurance subsidies, ad hoc disaster assistance, and conservation programs.

[2] Note that we refer to direct payments in this paper as the specific type of subsidy implemented in the U.S. between 2002 and 2014, rather than referring to direct payments more broadly as any payment made directly to farmers.

[3] However, Bryan, James Deaton, and Weersink (2015) do not find a strong impact of family relations on rental rates.

Breustedt and Habermann 2011; Hendricks, Janzen, and Dhuyvetter 2012; Ciaian and Kancs 2012; Kilian et al. 2012; Herck, Swinnen, and Vranken 2013; Michalek, Ciaian, and Kancs 2014; Kirwan and Roberts Forthcoming). There are a few exceptions in the literature that find larger impacts on rental rates (Lence and Mishra 2003; Patton et al. 2008; Goodwin, Mishra, and Ortalo-Magné 2011), but these studies are subject to concerns that unmeasured variability in productivity inflates their coefficient estimates.

One unresolved puzzle is that previous literature usually finds a large impact of government payments on land values (Latruffe and Le Mouël 2009) even though the estimated impact on rental rates is usually small. For example, Ifft, Kuethe, and Morehart (2015) find that an additional dollar of direct payments increases land value by about $18. Given that rents are a major determinant of land values (Alston 1986; Burt 1986), it seems odd that non-operators would be willing to pay a premium for land with greater government payments but not extract the government payments through higher rental rates. The most plausible explanation of the puzzle is that either the land value or the rental rate literature exploits variability in the data that over or underestimates the true effect.

Intuitively, our empirical strategy compares cash rental rates in counties that have similar market returns, but that have different direct payments due to the favoritism shown to areas that historically produced cotton or rice. Our econometric model uses county-level data and regresses cash rental rates on direct payments, expected market returns, and the proportion of cropland enrolled in the Average Crop Revenue Election (ACRE) program. We instrument direct payments with the share of cropland with cotton or rice base acres. We argue that the favoritism shown to cotton and rice is primarily due to political favoritism which should have no direct impact on rental rates except through government payments. Since cotton and rice production is concentrated in a particular region, there could be concerns that our instrument is correlated with differences in unmeasured expected market returns or differences in the rental market for this region. We use the framework of Conley, Hansen, and Rossi (2012) to construct revised standard errors that allow for a potential violation of the exclusion restriction.

According to the OECD Producer Support Estimates, the 2000–2014 average commodity-specific government transfers as a percent of total gross commodity receipts was only 5% for corn and soybeans and 7% for wheat while it was 20% for cotton and 12% for rice. Data that we construct for this paper also indicate that counties with cotton or rice base acres received substantially larger direct payments than counties with similar market returns but no cotton or rice base acres. There are several potential explanations for political favoritism towards cotton and rice. Gardner (1987) argues that farm programs are primarily a means of income redistribution and a commodity receives greater support if income can be redistributed more efficiently for that commodity. Thus, government support depends on supply and demand elasticities and the cost of political lobbying specific to each commodity (Gardner 1987). Another explanation for cotton and rice favoritism is that one-party rule in the Southern U.S. up to 1960 resulted in Southern lawmakers holding powerful positions (Gardner 1987).[4]

Exploiting this large, persistent difference in direct payments gives a more plausible estimate of the long-run incidence on rental rates compared to other articles that exploit changes in government payments between time periods (e.g., Kirwan 2009; Hendricks, Janzen, and Dhuyvetter 2012; Michalek, Ciaian, and Kancs 2014) or between fields with the same crop planted (Kirwan and Roberts Forthcoming). Rental rates within a particular geographic region may not fully reflect differences in direct payments if rates are established by the customary arrangements in the region (see Young and Burke 2001). However, rental rates between different regions may fully reflect direct payments as the customary arrangements in each region reflect the typical direct payments of that region. Similarly, small changes in direct payments over time may have a negligible impact on rental rates if rents tend to

[4] From 1931 to 1995, the chairman of the House Committee on Agriculture was from a Southern state for all but 10 years. From 1933 to 1995, the chairman of the Senate Committee on Agriculture was from a Southern state for all but 12 years.

be established at round numbers.[5] The most relevant parameter for understanding the ultimate beneficiaries of agricultural subsidies is to understand how rental rates would differ if subsidies were eliminated—a large, persistent shock.

We estimate that roughly \$0.81 of every dollar of direct payments accrues to non-operator landlords, but we cannot reject the null hypothesis of full incidence. Exploiting the variation in payments due to cotton and rice favoritism is critical to our results. If we restrict our analysis to only counties that have negligible cotton or rice base acres, then our estimate of the incidence has severe upward bias because we cannot perfectly control for expected market returns between counties in the same region. However, our two-stage least squares empirical strategy only requires that our estimates of expected market returns are not systematically over or underestimated for counties with cotton or rice base acres and we also allow for potential violations of the exclusion restriction.

Even though direct payments were eliminated in the 2014 Farm Bill, our estimate of the incidence is relevant to current and future farm programs for two reasons. First, understanding the incidence of fixed payments not tied to production in real world rental markets provides an important baseline for understanding the incidence of more complex programs. If direct payments are not fully reflected in rental rates, then economic theory under perfectly competitive rental markets may not provide realistic estimates of the long-run incidence of other types of programs. Second, Agriculture Risk Coverage (ARC) and Price Loss Coverage (PLC) payments, which were introduced in the 2014 Farm Bill, are both tied to base acres and base yields rather than current production.[6] Therefore, the incidence of ARC and PLC payments is likely similar to the incidence of direct payments although the incidence could be smaller for ARC and PLC due to uncertainty about the payments.

[5] For example, if rent is \$100/acre and direct payments decrease by \$2.27/acre, then rent may not change in order to keep the rental rate at a round number. However, if direct payments decrease by \$10/acre, then rent may decrease to \$90/acre.

[6] ARC provides payments when county-level revenue falls below a trigger and PLC provides payments when price falls below a trigger.

Identification Challenges

In this section, we review the main challenges in identifying the incidence of agricultural subsidies. We also describe approaches of previous literature and compare them to our approach in this paper.

Measuring the Rental Rate

The first challenge is to obtain data on the cash rental rate for the dependent variable. Several previous studies estimate the relationship between government payments and land values (Goodwin and Ortalo-Magné 1992; Just and Miranowski 1993; Weersink et al. 1999; Barnard et al. 1997; Ifft, Kuethe, and Morehart 2015). Translating these results into estimates of the proportion of subsidies reflected in land values, however, requires assumptions about the discount rate and expected stream of government payments (Kirwan 2009; Hendricks, Janzen, and Dhuyvetter 2012). Identifying the impact on rental rates provides a cleaner identification strategy since rental rates presumably depend on the current expected returns from agricultural production.

However, data on rental rates have not been as widely available as land value data. Some studies use cash rent calculated as total rent divided by total rented acres (Kirwan 2009; Hendricks, Janzen, and Dhuyvetter 2012), but this underestimates the true cash rental rate since total rented acres include acres rented by cash and crop-share agreements.[7] Hendricks, Janzen, and Dhuyvetter (2012) show how this measurement error biases the coefficient on government payments downward with their data and use secondary data to correct for the bias.

In this paper, we use data on the average cash rental rate for cropland at the county level. These data are obtained from NASS surveys of the cash rental rate for irrigated and

[7] Furthermore, the Census and Kansas Farm Management Association data include rent for pasture which does not receive government payments. The Farm Accountancy Data Network (FADN) used by Michalek, Ciaian, and Kancs (2014) and Ciaian and Kancs (2012) also only reports total rent and total rented acres but it is not clear to us how crop-share acreage is treated in their data.

nonirrigated cropland, rather than constructing the rental rate from total rent divided by rented acres. Other studies that use data on actual cash rental rates include Kirwan and Roberts (Forthcoming) and Goodwin, Mishra, and Ortalo-Magné (2011).

Expectation Error

The second challenge is to accurately measure *expected* government payments. Farm subsidy programs often depend on the harvest price—and more recently yield. Cash rental rates are negotiated before harvest, and thus government payments are uncertain. The econometrician, however, only observes data on the realized government payments. Regressing rent on realized government payments results in classical measurement error since the observed variable has a larger variance than the true variable. Therefore, the coefficient on government payments is likely to be biased towards zero, *ceteris paribus*.

Kirwan (2009) provides a creative solution to the measurement error problem. He argues that government payments in 1997 were known with certainty due to the introduction of production flexibility contracts that did not depend on price or current production. Therefore, Kirwan (2009) uses the 1997 government payments as an instrument for the difference in 1997 and 1992 government payments. Several other studies use lagged or future government payments as an instrument for current government payments (Lence and Mishra 2003; Hendricks, Janzen, and Dhuyvetter 2012; Kilian et al. 2012). Goodwin, Mishra, and Ortalo-Magné (2011) consider different specifications where they use the previous 5-year average of government payments to approximate expected payments or various instruments. Kirwan and Roberts (Forthcoming) include direct payments—which were known with certainty—in their regression and also include a dummy variable for whether or not the farmer expected to receive a counter-cyclical payment. Kirwan and Roberts (Forthcoming) use data from 2006 and 2007 when counter-cyclical and loan deficiency payments comprised a significant portion of total government payments but the amount of payments was uncertain at the time rents were established.[8]

We use rent data from 2012 when prices were so high above the triggers that farmers arguably perceived a negligible probability of receiving counter-cyclical and loan deficiency payments.[9] Direct payments, on the other hand, provided a fixed per acre payment for the life of the Farm Bill that did not depend on price or current production. One potential concern with our analysis, however, is that the 2008 Farm Bill also introduced the Average Crop Revenue Election (ACRE) Program. ACRE was a voluntary program that provided farmers with payments when state-level revenues fell below a trigger. Farmers that enrolled in ACRE lost 20% of their direct payments. Therefore, direct payments decreased in counties with greater ACRE enrollment. Farmers, however, did not likely anticipate receiving less government payments in these counties, or else they would not have enrolled in the ACRE program. We include the proportion of cropland enrolled in ACRE as a control.

Omitted Variable Bias

The third challenge is to control for expected returns other than direct payments. Not completely controlling for market returns biases the coefficient on direct payments upwards, *ceteris paribus*, since the unobserved variability in market returns is likely positively correlated with cash rent and direct payments. Another potential omitted variable is the expected payments from ACRE since the proportion of cropland enrolled in ACRE is not likely to completely control for expected ACRE payments. Expected ACRE payments are positively correlated with rent but negatively correlated with direct payments since farmers sacrificed

[8] Counter-cyclical and loan deficiency payments totaled $1.2 billion for production in 2006 and $0.8 billion for production in 2007 compared to $5.1 billion of direct payments (U.S. Department of Agriculture 2016). And for production in 2005, counter-cyclical and loan deficiency payments totaled $4.8 billion. Counter-cyclical and loan deficiency payments are usually paid in the year following production so we use data from government payments in the following year.

[9] Counter-cyclical and loan deficiency payments were essentially zero for 2012 crop production. Furthermore, counter-cyclical and loan deficiency payments were less than $22 million from production in the previous two years (U.S. Department of Agriculture 2016).

direct payments to enroll in ACRE. Therefore, the bias from omitting expected ACRE payments is likely downward.

Several articles exploit panel data and include fixed effects to control for time-invariant productivity (Kirwan 2009; Hendricks, Janzen, and Dhuyvetter 2012; Ciaian and Kancs 2012; Herck, Swinnen, and Vranken 2013; Michalek, Ciaian, and Kancs 2014). Patton et al. (2008) include fixed effects but Kirwan and Roberts (Forthcoming) argue that unobserved heterogeneity still biases their results since payments not tied to production were implemented in the last year of their sample so Patton et al. (2008) effectively include the level of payments as the explanatory variable. Lence and Mishra (2003) and Patton et al. (2008) use lagged returns as an instrument for current market returns to reduce attenuation bias of the effect of market returns. Goodwin, Mishra, and Ortalo-Magné (2011) use an historical average of agricultural sales minus production costs at the county-level as a control, but this includes returns from livestock production.

Kirwan and Roberts (Forthcoming) argue that they control for differences in expected market returns across fields by including farmers' "yield goal" as a control. The yield goal represents an expectation of yields rather than actual yields. One disadvantage of their approach is that the data are crop specific. Kirwan and Roberts (Forthcoming) have data on the rent of land planted to soybeans, for example, and the yield goal for soybeans but the yield goal for other crops planted in the rotation may have an even larger impact on the rental rate. Kirwan and Roberts (Forthcoming) argue that after controlling for the yield goal, the variation remaining in subsidies is due to random variability in historical yields used to calculate base. Such random variability is likely small since the base yield is calculated from a multi-year average and farmers had the option to update base yield in 2002 if yields from a recent period represented an improvement.

We take great effort to construct a control for market returns that accounts for variation in returns across space and across crops. However, we recognize that we are unlikely to perfectly control for expected market returns and expected ACRE payments so we propose an instrumental variable approach. Our approach and assumptions are described in detail in the next section.

Long-Run Incidence

The fourth challenge is to estimate the long-run incidence, allowing for adjustments in rental rates. Rental rates are likely to have substantial inertia to changes in government payments and market returns due to multi-year contractual agreements and customary rates may not adjust to small changes in expected benefits.

Using panel data with fixed effects exploits year-to-year changes which only capture short-run rental rate adjustments (Ciaian and Kancs 2012; Herck, Swinnen, and Vranken 2013; Michalek, Ciaian, and Kancs 2014). Kirwan (2009) uses long (five-year) differences. Hendricks, Janzen, and Dhuyvetter (2012) rely on the partial adjustment framework to estimate long-run impacts. The year-to-year variation in subsidies exploited by these studies is often small so rental rates may be slow to adjust or not adjust at all to maintain rent at a round number. The more relevant counterfactual is how rents adjust in the long run to large changes in subsidy rates given adjustments in contracts and customary rents. We exploit large cross-sectional variation in subsidy rates which inherently captures a long-run effect without having to explicitly specify the dynamic process (Pesaran and Smith 1995).[10]

Aggregation

The fifth challenge is to have data at the appropriate level of aggregation. Kirwan and Roberts (Forthcoming) assume that rents are established at the field-level. Estimates with aggregate data (i.e, at the farm or county level) are biased if fields with above-average rental rates also have above-average subsidies or if rent is averaged across subsidized and unsubsidized farmland and subsidies are averaged across all rented and owner-operated cropland.

[10] Lence and Mishra (2003) also exploit cross-sectional variation in rents but only in Iowa so they do not exploit large differences in subsidy rates due to commodity favoritism.

Kirwan and Roberts (Forthcoming) find that farm-level estimates of the incidence are roughly twice as large as field-level estimates.

An important assumption made by Kirwan and Roberts (Forthcoming) is that rental rates are field specific. We argue, however, that a single rental rate is likely to be established for all acreage within a tenant-landlord relationship so that the relevant unit of analysis is all acreage within the tenant-landlord agreement.[11] Under this alternative assumption, field-level subsidy rates vary more than the average tenant-landlord subsidy rate creating classical measurement error and attenuated coefficients with field-level data. Furthermore, Kirwan and Roberts (Forthcoming) find that the effect of subsidies on rental rates is smaller for larger farms which is consistent with tenant market power or consistent with more measurement error for larger farms that rent larger areas of land in each landlord relationship.

It may also be the case that rental rates depend on customary arrangements within a particular region. For example, Young and Burke (2001) note that cropshare agreements have different splits across different regions as would be predicted by conventional theory, but the agreements rarely vary within a geographic region even though soil quality clearly varies within a region. Young and Burke (2001) suggest that this occurs because contracts tend to cluster around a few discrete values and because contracts tend to conform to the customary local arrangements. In this case, the cash rental rate depends on the average direct payments within the region and the field-level direct payment is a noisy approximation of direct payments in the region resulting in attenuation bias.

[11] Aggregate statistics indicate that it is likely that a large portion of tenant-landlord relationships include multiple fields. According to the 2014 TOTAL (Tenure, Ownership, and Transition of Agricultural Land) survey, landowners that rent more than 200 acres represent 70% of all land rented in the South, Plains, and Midwest states. Landowners that rent more than 500 acres represent 45% of all land rented in this region. Furthermore, landowners renting to a single tenant represent 62% of all land rented in this region.

Econometric Model

Our identification strategy uses two-stage least squares (2SLS) to estimate the effect of direct payments on rental rates. Our second stage equation of interest is

$$(1) \quad Rent_i = \beta_1 + \beta_D DirectPmts_i + f\left(\boldsymbol{\beta}_R, MktReturns_i\right) + \beta_A ACRE_i + \varepsilon_i,$$

where $Rent_i$ is the average cash rental rate per acre for cropland in county i, $DirectPmts_i$ is the average direct payment subsidy per acre, $MktReturns_i$ is the expected market returns for cropland, $f(\cdot)$ is a function of expected market returns that is potentially nonlinear, $\boldsymbol{\beta}_R$ is a vector of parameters in the nonlinear function of expected market returns, $ACRE_i$ is the proportion of cropland enrolled in the ACRE program, and ε_i is the variation in rental rates from other unobserved factors. The objective of our paper is to estimate β_D, which represents the proportion of direct payments captured in rental rates. The first stage equation is

$$(2) \quad DirectPmts_i = \alpha_1 + \alpha_{CR} CottonRice_i + f\left(\boldsymbol{\alpha}_R, MktReturns_i\right) + \alpha_A ACRE_i + u_i,$$

where $CottonRice_i$ is the proportion of cropland with cotton or rice base acres.

First, consider why ordinary least squares (OLS) estimates of equation (1) are likely biased. For OLS to estimate the causal parameter β_D, direct payments per acre must be uncorrelated with the variation in rental rates not explained by our measure of market returns and ACRE enrollment (i.e, $Cov(DirectPmts_i, \varepsilon_i) = 0$). Given that we are unlikely to perfectly measure expected market returns and expected ACRE payments, this assumption is unlikely to hold. Any variability in returns not captured by our controls is included in the error term (i.e., an omitted variable) and is likely correlated with direct payments. The bias of OLS could be upwards or downwards depending whether the bias from omitted market returns or omitted ACRE payments dominates.

The bias of OLS may not be large when the sample includes counties that have differing amounts of base acreage in cotton or rice. Angrist (1998) shows that regression estimates an average coefficient where more weight is given to observations with a greater variance of direct payments conditional on the controls. The variance of direct payments is greatest between counties that have different amounts of cotton or rice base acreage. Therefore, OLS identifies the incidence of direct payments on rents primarily using the variation in direct payments due to commodity favoritism.

To further alleviate concerns about omitted variable bias, we consider 2SLS. Consistency of 2SLS requires two assumptions: (i) the first stage relationship between the instrument and the endogenous regressor exists and (ii) the exclusion restriction holds. The first assumption requires that $\alpha_{CR} \neq 0$. Furthermore, finite sample bias can exist if the relationship between the instrument and endogenous regressor is not sufficiently strong (Bound, Jaeger, and Baker 1995). In our case, the relationship between the share of cropland with cotton or rice base acreage and direct payments is strong as we show in our results.

The exclusion restriction in our model requires that the variation in rental rates left over after parsing out expected market returns and enrollment in ACRE cannot be correlated with the proportion of cropland with cotton or rice base acreage (i.e., $Cov(CottonRice_i, \varepsilon_i) = 0$). This assumption requires that our estimates of expected market returns are not systematically over or underestimated for counties with cotton or rice. The consistency of OLS requires that expected market returns are measured perfectly, which is a much more stringent assumption. For example, the exclusion restriction does not require that we perfectly measure the difference in market returns between two neighboring counties, but simply that on average we correctly measure the difference in market returns between counties with and without cotton and rice base acres.

The exclusion restriction also requires that there is nothing systematically different about counties producing cotton or rice apart from direct payments, expected market returns, and ACRE enrollment that would affect the rental rate. The exclusion restriction would be violated if, for example, there were cultural differences such that counties with cotton or rice base acres had more or less competitive rental markets.

The exclusion restriction is unlikely to hold perfectly in most applications and there are reasons to think that it might be violated in our model. Following Conley, Hansen, and Rossi (2012), equation (1) can be rewritten as

$$(3) \quad Rent_i = \beta_1 + \beta_D DirectPmts_i + f\left(\boldsymbol{\beta}_R, MktReturns_i\right) + \beta_A ACRE_i + \gamma CottonRice_i + \varepsilon_i,$$

where the exclusion restriction imposes $\gamma = 0$. Intuitively, γ represents the expected value of the difference in cash rent in a county where all cropland had cotton or rice base acres and the cash rent in a county that had no cotton or rice base acres—controlling for differences in direct payments, our measure of expected market returns, and ACRE enrollment.[12] The difference in cash rental rates represented by γ could occur because we have not completely controlled for differences in expected market returns or due to differences in the rental markets between counties with cotton or rice base acres and those without cotton or rice base acres.

When $\gamma \neq 0$, then the probability limit of 2SLS is written as $\hat{\beta}_D \xrightarrow{p} \beta_D + \gamma/\alpha_{CR}$ in our case where β_D, γ, and α_{CR} are scalars (Conley, Hansen, and Rossi 2012). The probability limit of 2SLS makes clear that the bias from violations of the exclusion restriction depends on the strength of the first stage relationship (see also Bound, Jaeger, and Baker 1995). Small deviations from the exclusion restriction can induce large bias when the first stage

[12]Let ρ_i represent the variation in rent not explained by $DirectPmts_i$, $MktReturns_i$, and $ACRE_i$,

$$Rent_i = \beta_0^\rho + \beta_D^\rho DirectPmts_i + f\left(\boldsymbol{\beta}_R^\rho, MktReturns_i\right) + \beta_A^\rho ACRE_i + \rho_i.$$

Then we can write $\gamma = E\left[\rho_i | CottonRice_i = 1\right] - E\left[\rho_i | CottonRice_i = 0\right]$.

relationship is weak and conversely relatively large deviations from the exclusion restriction may have a smaller effect on bias when the first stage relationship is strong. In practice, there is often a tradeoff between the plausible exogeneity of an instrument and the strength of the first stage relationship. We choose an instrument that has a strong first stage relationship but where the exclusion restriction is unlikely to hold perfectly.

To account for potential deviations from the exclusion restriction, we construct revised standard errors using the framework of Conley, Hansen, and Rossi (2012). We do not know the true value of γ but we make an assumption about likely values, essentially imposing a prior distribution for γ. We assume that $\gamma \sim N(0, \delta^2)$, where δ is the standard deviation of likely values of γ. We do not have any prior beliefs about whether γ is more likely to be positive or negative so we assume γ has mean zero. Imposing prior beliefs about the distribution of γ is more general than the standard 2SLS approach that imposes the prior belief that $\gamma = 0$. When γ is assumed to be normally distributed, Conley, Hansen, and Rossi (2012) show how to easily calculate a revised variance matrix by using a large sample approximation that assumes uncertainty about γ is of the same order of magnitude as sampling uncertainty. Conley, Hansen, and Rossi (2012) refer to this approach as a local-to-zero approximation.[13] In the results section, we discuss our specific prior beliefs about γ.

Data

First, we describe our data sources and the construction of variables and then show summary statistics and data visualizations.

[13] Another approach proposed by Conley, Hansen, and Rossi (2012) is to use Bayesian analysis that incorporates prior information about γ. A full Bayesian analysis also requires priors about other model parameters though. Conley, Hansen, and Rossi (2012) suggest that the Bayesian and local-to-zero approaches are likely to give similar results in large samples so we simply use the local-to-zero approach. Another alternative approach proposed by Conley, Hansen, and Rossi (2012) is to use only a support assumption for γ and construct the union of confidence intervals. The disadvantage of this approach is that the confidence intervals are likely to be large since it gives equal weight to all potential values of γ, even those at the extremes that seem unlikely. The local-to-zero approach gives tighter confidence intervals by assuming a normal distribution for the potential values of γ.

Data Description

We restrict our analysis to counties in four farm resource regions as defined by U.S. Department of Agriculture (2015): the Northern Great Plains, Prairie Gateway, Heartland, and Mississippi Portal. Altogether, the four regions in our analysis account for roughly 66% of U.S. cropland area.

Our dependent variable is the average cash rental rate for cropland in 2012. County-level data on the cash rental rate ($/acre) for irrigated and nonirrigated cropland are obtained from National Agricultural Statistics Service (NASS) survey data. We construct the average cash rental rate as irrigated rent times the share of cropland irrigated plus nonirrigated rent times the share of cropland nonirrigated. The share of cropland irrigated for each county is the ratio of harvested irrigated cropland to total cropland in 2012 obtained from the Census of Agriculture.[14] In some cases, we only have data on irrigated or nonirrigated rental rates. Often this occurs because a large majority of the cropland is either irrigated or nonirrigated. We use the nonirrigated rental rate as the county average when less than 10% of the county is irrigated and use the irrigated rental rate when more than 75% of the county is irrigated.

Data on direct payments and base acres enrolled in farm programs are obtained from the Farm Program Atlas from U.S. Department of Agriculture (2012). For our key explanatory variable, we construct direct payments per cropland acre as total direct payments in 2009 divided by total cropland acres in 2012. The proportion of county cropland that has cotton or rice base is calculated as the direct payment cotton and rice base acres divided by total cropland acres in 2012. Base acres enrolled in the ACRE program are also obtained from the Farm Program Atlas in order to calculate the proportion of cropland enrolled in ACRE.

[14]In many cases the Census does not report irrigated acreage in a county because it could risk disclosing an individual respondent's data. If irrigated acreage was not reported for 2012, then we use the average irrigated acreage from 2002 and 2007. If irrigated acreage was not reported for 2002, 2007, or 2012 then we assume zero irrigated acres.

We use the following equation to calculate the average expected market returns at the county level:

$$(4) \qquad MktReturns_i = (1-\phi_i)\sum_c \frac{acres_{ci}}{\sum_c acres_{ci}}\left[\frac{1}{5}\sum_{t=2008}^{2012}(Revenue_{cit} - Cost_{crt})\right],$$

where $MktReturns_i$ is the average expected market returns for county i, ϕ_i is the proportion of cropland in summer fallow in county i, $Revenue_{cit}$ is the expected revenue for crop c in county i in year t, $Cost_{crt}$ is the cost of production for crop c in ERS farm resource region r in year t, and $acres_{ci}$ are the average acres planted to crop c in county i. The crops considered for calculating expected market returns are corn, cotton, rice, soybeans, sorghum, and wheat. We use average expected returns over the past 5 years—but including 2012—to approximate the market returns relevant for setting cash rental rates in 2012. An alternative would be to calculate a measure of expected market returns for 2012 only; however, we expect that cash rents are fairly sticky and do not fully adjust each year in response to different prices so market returns in previous years affect the current cash rental rate.

For all crops, except cotton, expected revenue is calculated as $Revenue_{cit} = Price_{cst} \times Yield_{cit}$, where $Price_{cst}$ is the price for crop c in state s in year t and $Yield_{cit}$ is the trend yield for crop c in county i in year t. State-level marketing-year prices are obtained from NASS. If the state-level price for a crop is missing in a particular year, it is replaced by the average price in all states with data in that year. The trend yield is estimated from county-specific linear trend regressions using data from 1980 to 2012. We only estimate trend yield if there are 20 or more observations for a county and if there was at least one yield observation from 2007 to 2012. We use trend yields rather than observed yields because cash rents depend on expected market returns and average realized returns in the five-year period could deviate substantially from expected market returns if weather was especially good or poor.

For cotton, expected revenue includes revenue from cotton lint and cottonseed production. The revenue from cotton lint production is calculated the same as for other crops. Cottonseed prices are also state-level prices. NASS does not, however, report county-level cottonseed production. We assume cottonseed yield is 1.62 times the cotton lint trend yield based on data in U.S. Department of Agriculture (2014).[15]

For all crops, production expenses are obtained by farm resource region from U.S. Department of Agriculture (2014). We include all operating costs and allocated overhead but exclude the opportunity cost of land (i.e., land rent). U.S. Department of Agriculture (2014) provides cost estimates for the following regions for each commodity: soybeans in all regions, corn, wheat, and sorghum in the Heartland, Prairie Gateway, and Northern Great Plains; cotton in the Heartland, Prairie Gateway, and Mississippi Portal; and rice in the Mississippi Portal.[16] For corn, wheat, and sorghum expenses in the Mississippi Portal, we use expenses from the Heartland. For rice expenses in the Heartland, we use expenses from the Mississippi Portal.[17] Using expenses from neighboring regions ensures that we have cost estimates in every county where we have trend yield and acreage data for a commodity.

In equation (4), we average market returns across crops where we weight by the share of acreage planted to each crop ($\frac{acres_{ci}}{\sum_c acres_{ci}}$). The acres planted to the crop is the 2008 to 2012 average planted acreage. If acreage data are missing for a particular crop in all years, then we assume the crop is not produced in the county. If acreage data are available but trend yield is not available for the crop, then we set acreage for that crop equal to zero.

Equation (4) assumes that the returns from summer fallowed land are zero. We obtain 2012 acres in summer fallow from the Census of Agriculture and divide it by cropland acres

[15]The ratio of cottonseed yield to cotton lint yield is equal to 1.62 for every year between 2007 and 2012 in the Prairie Gateway and Mississippi Portal according to ERS costs and returns.

[16]ERS only provides cost estimates up to 2010 for sorghum in the Heartland. We calculate the average ratio of sorghum costs from 2003 to 2010 between the Prairie Gateway and Heartland to impute costs in the Heartland for 2011 and 2012. From 2003 to 2010, costs ranged 8–15% larger in the Heartland. On average, costs are 10% larger in the Heartland for sorghum.

[17]There are only a few counties in the southern portion of the Heartland region where rice is produced.

to calculate ϕ_i. Annual data do not exist at the county level for summer fallow acreage so ϕ_i is constant over time.

We drop observations from our sample if we have estimates of market returns from less than 25% of total cropland.[18] Counties that are dropped are likely those counties where other crops comprise a major portion of cropland area and our measure of market returns may not be representative for these counties. In the sample used for econometric analysis, expected market returns accounts for more than 50% of cropland area for 81% of counties.

Our econometric analysis also excludes observations if market returns are greater than \$325/acre. There is only one county with more than 1% cotton or rice base acreage that has market returns greater than \$325/acre while there are 255 counties with less than 1% cotton or rice base acreage. Including observations with market returns greater than \$325/acre creates a problem where—for this portion of the data—we have little overlap between counties with and without cotton or rice base acres. In a later section, we explore the robustness of our estimates to different specifications for dropping counties.

Alternatively, we could estimate expenses using county level data from the Census of Agriculture similar to the approach taken by Goodwin, Mishra, and Ortalo-Magné (2011). One problem with using Census data is that the Census does not differentiate expenses for crop production. For example, expenses for machinery rent and utilities also account for expenses for livestock production. Therefore, expenses from the Census will be systematically biased estimates of crop production expenses depending on the amount of livestock production in the county.

Data Summary and Visualization

Table 1 shows summary statistics for the variables used in our econometric analysis. Panel A shows summary statistics for counties with less than 1% of cropland with cotton or rice base acres (461 counties) and panel B for counties with more than 1% of cropland with

[18] That is, we add $\sum_c acres_{ci}$ and summer fallow acreage and divide by total cropland acres and drop the observation if the proportion is less than 0.25.

cotton or rice base acres (178 counties). The mean value for direct payments for the counties with negligible cotton or rice base ($10.92) is lower than for those counties with cotton or rice base ($19.35). The mean values for cash rent and market returns are higher in counties with negligible cotton or rice base acreage. Enrollment in the ACRE program was greater in counties with negligible cotton or rice base acreage. Among those counties with cotton or rice base, the proportion of cropland with cotton or rice base acres differs substantially among counties with a mean of 0.32 and a standard deviation of 0.23.[19]

Figure 1 shows maps for cash rent, market returns, direct payments, and the proportion of cropland with cotton or rice base acres. The light grey area shows those counties that are not included in one of the four farm resource regions included in our sample. The dark grey area shows those counties that had missing data for one of the variables used in the econometric analysis. Missing data usually occurred because county-level cash rent was not reported or market returns could not be calculated because trend yield or acreage data were missing. The light blue area shows those counties that were dropped from our analysis because either market returns were calculated for less than 25% of cropland area or market returns exceeded $325/acre.

High cash rental rates are concentrated in the area surrounding the Corn Belt and Mississippi Portal and rental rates are smaller moving west to the plains states (figure 1a). Market returns generally follow a similar pattern as the cash rental rate (figure 1b). Direct payments, however, are much larger in the Mississippi Portal region and portions of Texas compared to the Northern regions (figure 1c). The larger direct payments are directly related with the proportion of cropland with cotton or rice base acres (figure 1d).

Figure 2 shows a scatterplot of the data used in our econometric analysis for the relationship between market returns and the average cash rental rate. Purple circles indicate counties with less than 1% cotton or rice base and orange diamonds indicate counties with more than 1% cotton or rice base. The clear positive relationship between returns and the rental rate

[19] Cotton or rice base acres exceeded cropland acreage in one county. This may have occurred if cropland area decreased from the time base was established.

provides some support for the accuracy of our measurement of market returns—though not necessarily eliminating omitted variable bias.

The most important observation from figure 2 is that conditional on the same market returns, counties with cotton or rice base acres tend to have higher rental rates. Furthermore, from figure 3, we see that conditional on the same market returns, counties with cotton or rice base acres tend to have much larger direct payments. These simple observations from the data provide suggestive evidence that direct payments are at least partially captured in the rental rate.

Econometric Results

Next, we show the econometric results that conduct more rigorous tests than the graphical evidence above and estimate the proportion of direct payments reflected in rental rates. We first show OLS results which we argue are likely biased, then we show our preferred 2SLS results and robustness checks.

OLS Results

Table 2 reports OLS results for the effect of direct payments on rental rates. The different columns report estimates where we control for market returns with different polynomial specifications. The R^2 indicates that our regression is able to explain roughly 73% of the variation in cash rents.

Each of the specifications in table 2 give similar estimates of the incidence. For the linear functional form (column 1), for example, the coefficient on direct payments indicates that cash rents increase by \$0.51 for every dollar of direct payments. For all three specifications in table 2, we reject the null hypotheses of $\beta_D = 0$ and $\beta_D = 1$ at the 5% level. Our standard error of the coefficient on direct payments (≈ 0.15) is similar in magnitude to the standard error on direct payments in Kirwan and Roberts (Forthcoming) for soybeans (≈ 0.11).

The coefficient on market returns in column (1) of table 2 indicates that cash rents increase by \$0.36 for an additional dollar of market returns. In theory, this coefficient should also represent the amount that landowners would capture from a purely coupled subsidy. Our result is consistent with Alston (2010) who finds that standard economic theory suggest landowners receive about \$0.39 from a pure output subsidy under plausible parameters with a range from \$0.19 to \$0.62 under alternative parameter assumptions. An important caveat, is that our coefficient on market returns could be biased downward to the extent that we have measurement error in expected market returns. However, our coefficient is much larger than estimated by Kirwan (2009) and Hendricks, Janzen, and Dhuyvetter (2012)—0.03 and 0.11, respectively.[20] Goodwin, Mishra, and Ortalo-Magné (2011) estimate a coefficient on markets returns of about 0.12–0.16 depending on their specification. The estimate of Goodwin, Mishra, and Ortalo-Magné (2011) is likely biased downwards given that they use an historical average of actual returns from crop and livestock production. Our coefficient on market returns is similar to Lence and Mishra (2003).

As expected, the coefficients on the proportion of cropland enrolled in ACRE indicate that cash rents are larger in counties with more land enrolled in ACRE, *ceteris paribus* (table 2). Direct payments per cropland acre within a county decrease as more area is enrolled in ACRE because farmers had to reduce their direct payments in order to enroll in the ACRE program. However, farmers may have still expected to receive some subsidy payments from ACRE and so the coefficient on ACRE reflects this value.

Table 3 reports OLS results with alternative specifications. Column (1) in each of the panels shows a simple bivariate relationship between cash rent and direct payments with different samples. Columns (2)-(4) show results with linear, quadratic, and cubic controls for market returns.

[20] Kirwan (2009) and Hendricks, Janzen, and Dhuyvetter (2012) both include revenues and costs as separate variables. Here we cite the coefficient on revenues from these articles which is larger in absolute magnitude than the coefficient on costs in both cases.

Panel A in table 3 shows results that omit the proportion of cropland enrolled in ACRE as a control. The coefficient in column (1) shows that OLS is biased upwards substantially when controls for market returns and ACRE enrollment are omitted. It is not surprising that the coefficient on direct payments exceeds 1 in the simple bivariate regression. Cash rental rates are larger than direct payments per acre and direct payments are positively correlated with market returns. So the coefficient on direct payments in the bivariate regression reflects the impact of subsidies and market returns on rental rates. Consistent with our discussion in the model section, results in columns (2)-(4) show smaller estimates of the incidence of direct payments on rents when we omit the control for ACRE enrollment.

Panel B in table 3 shows regression results using data from only those counties with negligible cotton or rice base acreage. These results do not exploit the variability in direct payments due to commodity favoritism. The coefficients on direct payments in columns (2)-(4) are much larger than those in table 2, consistent with a large omitted variable bias when we do not exploit the variability from commodity favoritism.

Panel C in table 3 shows results using only counties with more than 1% cotton or rice base acreage. Since the proportion of cropland with cotton or rice base acreage varies across these counties, OLS exploits—at least in part—the variability in direct payments due to variation in cotton and rice base acreage. Therefore, estimates in panel C should have less bias than those in panel B. Indeed, OLS estimates in columns (2)-(4) of panel C are much smaller than in panel B and are slightly larger than OLS estimates for the entire sample in table 2. The main disadvantage of the OLS estimates in panel C is that the standard errors increase to about 0.25 compared to 0.15 in table 2 since we only have 178 observations in panel C.

The main concern with OLS estimates in table 2 and panel C of table 3 is that there could still be some remaining unobserved heterogeneity affecting rental rates that is also correlated with direct payments. For example, if we have omitted some variability in market returns, then OLS estimates are biased upwards. Alternatively, if we have not sufficiently controlled for expected ACRE program payments, then OLS estimates are likely biased downwards. Next, we consider an instrumental variables approach to exploit the variability in direct payments due to favoritism for cotton and rice.

2SLS Results

Table 4 reports our first-stage regression results. Not surprisingly, the share of cropland with cotton or rice base acreage has a large impact on direct payments even after controlling for market returns and ACRE enrollment. The results indicate that direct payments are roughly \$37/acre larger if all of the cropland in a county has cotton or rice base acreage relative to a county with no cotton or rice base acreage. This is a large difference in payments, given that the average direct payments in counties with less than 1% cotton or rice base is only \$11/acre in our sample (see table 1).

Our first-stage results also indicate no evidence of a weak instrument problem. The F-statistics for the coefficient on our instrument exceed 300 for all specifications. This suggests minimal finite sample bias for instrumental variables (Staiger and Stock 1997). The strong relationship between the instrument and direct payments also means that violations of the exclusion restriction have a relatively smaller impact on our estimate of the incidence than if we had a weak instrument.

Table 5 reports estimates of the incidence using 2SLS. Heteroskedasticity-robust standard errors are reported in parentheses under each coefficient. Standard errors that allow for a potential violation of the exclusion restriction are reported in brackets under each coefficient. We place asterisks next to the standard errors to indicate the statistical significance for each type of standard errors.

We relax the exclusion restriction using the local-to-zero approximation proposed by Conley, Hansen, and Rossi (2012) and impose the prior distribution $\gamma \sim N(0, \delta^2)$. We assume γ has mean zero because we do not have a prior on whether cash rents are likely to be systematically higher or lower in counties with cotton or rice base acreage after accounting for subsidies and our measure of market returns. We assume $\delta = 5$. This assumption implies that we have 95% confidence that the value of γ is between -9.8 and +9.8. This allows for the possibility that cash rents in counties with cotton or rice base acres on all cropland could be \$9.80/acre greater (or less) than in counties with no cotton or rice base acres due to factors not accounted for in our regression. The mean cash rental rate for counties with more than 1% cotton or rice base acres is \$62/acre, so our prior on γ allows for a substantial violation of the exclusion restriction. Of course, our assumption of a normal distribution assumes that γ is most likely close to zero.

Column (1) in table 5 indicates that cash rents increase by \$0.81 for every dollar of direct payments. The coefficient on direct payments is larger with 2SLS than with OLS indicating that the bias from omitted variables in OLS was downward. The most likely explanation is that our control for ACRE enrollment does not sufficiently control for the expected payments from the ACRE program and these unobserved payments bias OLS downwards. The p-value for a test for endogeneity that is robust to heteroskasticity is reported near the bottom of table 5 (see Wooldridge 2010). The test rejects the null hypothesis of exogeneity for each specification.

The heteroskedasticity-robust standard error for the coefficient on direct payments is 0.17, only slightly larger than 0.15 from the OLS model. Accounting for a potential violation of the exclusion restriction, the standard error increases to 0.22 (standard error in brackets in column 1). With either type of standard error, we reject the null hypothesis that $\beta_D = 0$ but fail to reject the null that $\beta_D = 1$ at the 5% level.

Including quadratic and cubic controls for market returns does not dramatically alter the coefficient on direct payments (columns 2-3). The coefficient on market returns is similar with 2SLS compared to OLS (compare tables 5 and 2).

Robustness

In the supplementary appendix, we report results from several different robustness checks and describe the specifications for the robustness checks in more detail. Table A1 shows results if we use different thresholds for the proportion of cropland area that is accounted for in our estimate of market returns or different thresholds for market returns to maintain overlap in our sample. We also consider estimating the model for all observations with nonmissing data. The coefficient on direct payments in these specifications varies between 0.76 and 0.83, so these assumptions make little difference to our results.

Table A2 in the supplementary appendix shows results if we calculate the variables in our analysis differently. The first column shows results if we use cropland used for crops (i.e., the sum of harvested, failed, and summer fallowed cropland) rather than total cropland area to derive per acre estimates. The coefficient on direct payments is 0.75. Our results are also not highly sensitive to using a 4 or 3-year historical average of market returns instead of a 5-year average. If we calculate market returns over the period 2009–2012 instead of 2008–2012, then the coefficient on direct payments is 0.75 and the coefficient on market returns is 0.33. If we use the period 2010–2012 for market returns, then the coefficient on direct payments is 0.89 and the coefficient on market returns is 0.28.

Table A3 in the supplementary appendix shows results using the rental rate from different years. Using rental rates from 2011 and market returns from 2007–2011, the coefficient on direct payments is 1.02. Using rental rates from 2010 and 2009 and the respective five year periods for market returns, the coefficients on direct payments are 1.37 and 1.31, but the difference from 1 is not statistically significant. The coefficient on market returns in each of these specifications ranges from 0.34 to 0.40. Using rental rates from earlier years gives a larger estimate of the incidence. Estimates from earlier years could be problematic because there was a sharp increase in agricultural returns in 2008 due to an increase in commodity prices so a five-year historical average of expected market returns is less likely to represent market returns in these prior years than for 2012 used in our main specification.

Discussion and Conclusion

Our preferred estimate of the incidence of direct payments on rental rates is 2SLS with a linear control for market returns (column 1 of table 5) and assuming the instrument is only plausibly exogenous (standard error in brackets). This specification isolates the variability in direct payments due to commodity favoritism, but without strictly imposing the exclusion restriction.

Our preferred specification indicates that \$0.81 of every dollar of direct payments is captured by landowners through adjustments in the rental rate in the long run. Standard economic theory suggests that subsidies not tied to production should be completely reflected in rental rates ($\beta_D = 1$) and our econometric estimates are not able to reject this null hypothesis, though the evidence suggests slightly less than full incidence on rental rates. We also estimate that about \$0.36 of every dollar of expected market returns accrues to landowners through higher rental rates in the long run, which is also consistent with economic theory.

According to the 2012 TOTAL Survey, about 46% of cropland in the United States is rented by non-operator landlords.[21] Assuming that the incidence of direct payments is similar across different types of rental rate agreements, our estimate indicates that of the annual \$4.7 billion of direct payments in the 2008 Farm Bill, about \$1.75 billion ($1.75 = 4.7 \times 0.46 \times 0.81$) was captured by non-operator landlords.

Kirwan (2009) and Kirwan and Roberts (Forthcoming) estimate that only about \$0.25 is captured by landowners and Hendricks, Janzen, and Dhuyvetter (2012) estimate about \$0.37 is captured by landowners in the long run. Many other studies also estimate a small incidence (Breustedt and Habermann 2011; Ciaian and Kancs 2012; Kilian et al. 2012; Michalek, Ciaian, and Kancs 2014; Herck, Swinnen, and Vranken 2013).

[21] About 57% of cropland in the United States is rented and roughly 81% of rented cropland is rented by non-operator landlords. Note that the percent of cropland rented (57%) is larger than the percent of all agricultural land rented (39%).

Overall, we argue that exploiting large differences in subsidy rates across regions with different commodities provides a more plausible estimate of the incidence of direct payments on rental rates in the long run. It could be that rental rates capture little of the difference in direct payments that occur over time or across areas with similar commodities. But rental rates may capture most of the large, persistent difference in direct payments that occurs between regions. A rationale for this distinction in incidence for different types of changes in subsidies is that rental rates may be set by customary local arrangements that are slow to adjust and tend toward rounds numbers. The impact of persistent differences in direct payments is arguably most relevant for policy analysis that seeks to understand the ultimate beneficiaries of these programs.

Agriculture Risk Coverage (ARC) and Price Loss Coverage (PLC) payments are similar to direct payments in that they are tied to base acres and base yields rather than current production. Our estimates indicate that non-operator landlords are likely to capture a large portion of ARC and PLC payments. One caveat is that ARC and PLC payments are uncertain because they depend on market prices and—for ARC—yields. Future research could explore the impact of payment uncertainty on the incidence of subsidies.

We began this paper by noting that the politics of government interventions depend as much on the distribution of benefits and costs as the overall change in social welfare. Our empirical results indicate that there is a tradeoff between reducing trade distortions (i.e., transferring with less deadweight loss) and transferring benefits to farm operators. Subsidies tied directly to production are trade distorting, but non-operator landlords only capture roughly 36% of the benefits on rented land. Subsidies not tied to production are less trade distorting, but non-operator landlords capture roughly 81% of the benefits on rented land.

References

Alston, J.M. 1986. "An Analysis of Growth of U.S. Farmland Prices, 1963–82." *American Journal of Agricultural Economics* 68:1–9.

—. 2010. "The Incidence of U.S. Farm Programs." In V. E. Ball, R. Fanfani, and L. Gutierrez, eds. *The Economic Impact of Public Support to Agriculture*. New York, NY: Springer, vol. 7, chap. 5, pp. 81–105.

Alston, J.M., and J.S. James. 2002. "The Incidence of Agricultural Policy." In B. L. Gardner and G. C. Rausser, eds. *Handbook of Agricultural Economics*. Elsevier, vol. 2, Part B, pp. 1689–1749.

Angrist, J.D. 1998. "Estimating the Labor Market Impact of Voluntary Military Service Using Social Security Data on Military Applicants." *Econometrica* 66:249–288.

Barnard, C.H., G. Whittaker, D. Westenbarger, and M. Ahearn. 1997. "Evidence of Capitalization of Direct Government Payments into U.S. Cropland Values." *American Journal of Agricultural Economics* 79:1642–1650.

Bound, J., D.A. Jaeger, and R.M. Baker. 1995. "Problems with Instrumental Variables Estimation When the Correlation Between the Instruments and the Endogeneous Explanatory Variable is Weak." *Journal of the American Statistical Association* 90:443–450.

Breustedt, G., and H. Habermann. 2011. "The Incidence of EU Per-Hectare Payments on Farmland Rental Rates: A Spatial Econometric Analysis of German Farm-Level Data." *Journal of Agricultural Economics* 62:225–243.

Bryan, J., B. James Deaton, and A. Weersink. 2015. "Do Landlord-Tenant Relationships Influence Rental Contracts for Farmland or the Cash Rental Rate?" *Land Economics* 91:650–663.

Burt, O.R. 1986. "Econometric Modeling of the Capitalization Formula for Farmland Prices." *American Journal of Agricultural Economics* 68:10–26.

Ciaian, P., and d. Kancs. 2012. "The Capitalization of Area Payments into Farmland Rents: Micro Evidence from the New EU Member States." *Canadian Journal of Agricultural Economics* 60:517–540.

Conley, T.G., C.B. Hansen, and P.E. Rossi. 2012. "Plausibly Exogenous." *Review of Economics and Statistics* 94:260–272.

Floyd, J.E. 1965. "The Effects of Farm Price Supports on the Returns to Land and Labor in Agriculture." *Journal of Political Economy* 73:148–158.

Gardner, B.L. 1987. "Causes of U.S. Farm Commodity Programs." *Journal of Political Economy* 95:290–310.

Goodwin, B.K., A.K. Mishra, and F. Ortalo-Magné. 2011. "The Buck Stops Where? The Distribution of Agricultural Subsidies." Working Paper No. 16693, National Bureau of Economic Research, Cambridge, MA.

Goodwin, B.K., and F. Ortalo-Magné. 1992. "The Capitalization of Wheat Subsidies into Agricultural Land Values." *Canadian Journal of Agricultural Economics* 40:37–54.

Hendricks, N.P., J.P. Janzen, and K.C. Dhuyvetter. 2012. "Subsidy Incidence and Inertia in Farmland Rental Markets: Estimates from a Dynamic Panel." *Journal of Agricultural and Resource Economics* 37:361–378.

Hendricks, N.P., and D.A. Sumner. 2014. "The Effects of Policy Expectations on Crop Supply, with an Application to Base Updating." *American Journal of Agricultural Economics* 96:903–923.

Hennessy, D.A. 1998. "The Production Effects of Agricultural Income Support Policies under Uncertainty." *American Journal of Agricultural Economics* 80:46–57.

Herck, K.V., J. Swinnen, and L. Vranken. 2013. "Capitalization of Direct Payments in Land Rents: Evidence from New EU Member States." *Eurasian Geography and Economics* 54:423–443.

Ifft, J., T. Kuethe, and M. Morehart. 2015. "The Impact of Decoupled Payments on U.S. Cropland Values." *Agricultural Economics* 46:643–652.

Just, D.R., and J.D. Kropp. 2013. "Production Incentives from Static Decoupling: Land Use Exclusion Restrictions." *American Journal of Agricultural Economics* 95:1049–1067.

Just, R.E., and J.A. Miranowski. 1993. "Understanding Farmland Price Changes." *American Journal of Agricultural Economics* 75:156–168.

Kilian, S., J. Antón, K. Salhofer, and N. Röder. 2012. "Impacts of 2003 CAP Reform on Land Rental Prices and Capitalization." *Land Use Policy* 29:789–797.

Kirwan, B., and M.J. Roberts. Forthcoming. "Who Really Benefits from Agricultural Subsidies: Evidence from Field-Level Data." *American Journal of Agricultural Economics*, in press.

Kirwan, B.E. 2009. "The Incidence of US Agricultural Subsidies on Farmland Rental Rates." *Journal of Political Economy* 117:138–164.

Latruffe, L., and C. Le Mouël. 2009. "Capitalization of Government Support in Agricultural Land Prices: What do We Know?" *Journal of Economic Surveys* 23:659–691.

Lence, S.H., and A.K. Mishra. 2003. "The Impacts of Different Farm Programs on Cash Rents." *American Journal of Agricultural Economics* 85:753–761.

Michalek, J., P. Ciaian, and d. Kancs. 2014. "Capitalization of the Single Payment Scheme into Land Value: Generalized Propensity Score Evidence from the European Union." *Land Economics* 90:260–289.

Patton, M., P. Kostov, S. McErlean, and J. Moss. 2008. "Assessing the Influence of Direct Payments on the Rental Value of Agricultural Land." *Food Policy* 33:397–405.

Perry, G.M., and L.J. Robison. 2001. "Evaluating the Influence of Personal Relationships on Land Sale Prices: A Case Study in Oregon." *Land Economics* 77:385–398.

Pesaran, M.H., and R. Smith. 1995. "Estimating Long-run Relationships from Dynamic Heterogeneous Panels." *Journal of Econometrics* 68:79–113.

Schlegel, J., and L.J. Tsoodle. 2008. "Non-Irrigated Crop-Share Leasing Arrangements in Kansas." Staff Paper No. 08-03, Kansas Farm Management Association.

Staiger, D., and J.H. Stock. 1997. "Instrumental Variables Regression with Weak Instruments." *Econometrica* 65:557–586.

Tsoodle, L.J., B.B. Golden, and A.M. Featherstone. 2006. "Factors Influencing Kansas Agricultural Farm Land Values." *Land Economics* 82:124–139.

U.S. Department of Agriculture. 2014. "Commodity Costs and Returns." Economic Research Service, Available at: `http://www.ers.usda.gov/data-products/commodity-costs-and-returns.aspx`, Accessed March 24, 2014.

—. 2012. "Farm Program Atlas." Economic Research Service, Available at: `http://www.ers.usda.gov/data-products/farm-program-atlas.aspx`, Accessed March 24, 2014.

—. 2015. "Farm Resource Regions." Economic Research Service, Available at: `http://www.ers.usda.gov/media/926929/aib-760_002.pdf`, Accessed May 9, 2015.

—. 2016. "U.S. and State Farm Income and Wealth Statistics." Economic Research Service, Available at: `http://ers.usda.gov/data-products/farm-income-and-wealth-statistics.aspx`, Accessed February 9, 2016.

Weersink, A., S. Clark, C.G. Turvey, and R. Sarker. 1999. "The Effect of Agricultural Policy on Farmland Values." *Land Economics* 75:425–439.

Wooldridge, J.M. 2010. *“Econometric Analysis of Cross Section and Panel Data”*, 2nd ed. Cambridge, MA: MIT press.

Young, H.P., and M.A. Burke. 2001. “Competition and Custom in Economic Contracts: A Case Study of Illinois Agriculture.” *American Economic Review* 91:559–573.

Tables

Table 1: Summary Statistics

	Observations	Mean	Std. Dev.	Min	Max
Panel A. Counties with Less than 1% Cotton or Rice Base					
Cash Rent ($/acre)	461	98.28	50.30	10.50	237.13
Direct Payments ($/acre)	461	10.92	3.91	1.21	22.20
Market Returns ($/acre)	461	168.59	105.76	-76.94	324.44
Proportion ACRE	461	0.10	0.12	0.00	0.62
Panel B. Counties with More than 1% Cotton or Rice Base					
Cash Rent ($/acre)	178	62.09	36.35	10.50	145.00
Direct Payments ($/acre)	178	19.35	12.41	2.44	60.88
Market Returns ($/acre)	178	60.73	94.86	-143.22	315.29
Proportion ACRE	178	0.03	0.09	0.00	0.59
Proportion Cotton or Rice	178	0.32	0.23	0.01	1.10

Table 2: OLS Results for the Incidence of Direct Payments on Cash Rental Rates

	(1)	(2)	(3)
Direct Payments	0.509	0.546	0.545
	(0.151)**	(0.148)**	(0.149)**
Market Returns	0.360	0.270	0.247
	(0.010)**	(0.017)**	(0.011)**
Market Returns^2		0.000	0.001
		(0.000)**	(0.000)**
Market Returns^3			-0.000
			(0.000)**
Proportion ACRE	27.857	26.442	26.667
	(8.628)**	(8.632)**	(8.505)**
Intercept	29.245	29.989	28.170
	(2.038)**	(1.907)**	(2.074)**
Observations	639	639	639
R^2	0.729	0.737	0.739

Standard errors in parentheses represent heteroskedasticity-robust standard errors. Asterisks * and ** denote significance at the 10% and 5% levels, respectively.

Table 3: OLS Results with Alternative Specifications

	(1)	(2)	(3)	(4)
Panel A. Omit ACRE Control				
Direct Payments	1.443	0.426	0.468	0.467
	(0.249)**	(0.148)**	(0.144)**	(0.146)**
Market Returns	No	Linear	Quadratic	Cubic
Proportion ACRE	No	No	No	No
Observations	639	639	639	639
R^2	0.058	0.725	0.734	0.735
Panel B. Counties with Less than 1% Cotton or Rice Base				
Direct Payments	7.690	1.534	1.377	1.507
	(0.424)**	(0.396)**	(0.429)**	(0.441)**
Market Returns	No	Linear	Quadratic	Cubic
Proportion ACRE	No	Yes	Yes	Yes
Observations	461	461	461	461
R^2	0.358	0.748	0.749	0.752
Panel C. Counties with More than 1% Cotton or Rice Base				
Direct Payments	1.749	0.712	0.724	0.622
	(0.230)**	(0.245)**	(0.251)**	(0.252)**
Market Returns	No	Linear	Quadratic	Cubic
Proportion ACRE	No	Yes	Yes	Yes
Observations	178	178	178	178
R^2	0.357	0.594	0.599	0.613

Standard errors in parentheses represent heteroskedasticity-robust standard errors.
Asterisks * and ** denote significance at the 10% and 5% levels, respectively.

Table 4: Proportion of Base Acres Cotton or Rice and Direct Payments (First-Stage)

	(1)	(2)	(3)
Proportion Cotton or Rice	36.653	36.635	36.875
	(2.036)**	(2.023)**	(2.055)**
Market Returns	0.029	0.030	0.036
	(0.002)**	(0.004)**	(0.004)**
Market Returns^2		-0.000	-0.000
		(0.000)	(0.000)**
Market Returns^3			0.000
			(0.000)**
Proportion ACRE	2.061	2.072	2.117
	(1.447)	(1.451)	(1.404)
Intercept	5.742	5.734	6.159
	(0.307)**	(0.316)**	(0.335)**
F-Statistic ($H_0 : \alpha_{CR} = 0$)	324	328	322
Observations	639	639	639
R^2	0.687	0.687	0.691

The dependent variable is direct payments per acre. Standard errors in parentheses represent heteroskedasticity-robust standard errors. Asterisks * and ** denote significance at the 10% and 5% levels, respectively.

Table 5: Two-Stage Least Squares Results for the Incidence of Direct Payments on Cash Rental Rates

	(1)	(2)	(3)
Direct Payments	0.807	0.861	0.835
	(0.174)**	(0.171)**	(0.176)**
	[0.221]**	[0.219]**	[0.222]**
Market Returns	0.355	0.262	0.240
	(0.010)**	(0.017)**	(0.011)**
	[0.011]**	[0.018]**	[0.012]**
Market Returns^2		0.000	0.001
		(0.000)**	(0.000)**
		[0.000]**	[0.000]**
Market Returns^3			-0.000
			(0.000)**
			[0.000]**
Proportion ACRE	32.173	30.927	30.802
	(8.624)**	(8.611)**	(8.448)**
	[8.629]**	[8.616]**	[8.453]**
Intercept	25.595	26.178	24.665
	(2.037)**	(1.942)**	(1.981)**
	[2.183]**	[2.094]**	[2.149]**
P-value for test of endogeneity (H_0=exogeneity)	0.003	0.001	0.005
Observations	639	639	639

Standard errors in parentheses represent heteroskedasticity-robust standard errors imposing the exlcusion restriction. Standard errors in brackets are calculated allowing for a potential violation of the exclusion restriction (Conley, Hansen, and Rossi 2012). The test for endogeneity is conducted using the results that impose the exclusion restriction and the test is robust to heteroskedasticty.
Asterisks * and ** denote significance at the 10% and 5% levels, respectively.

Figures

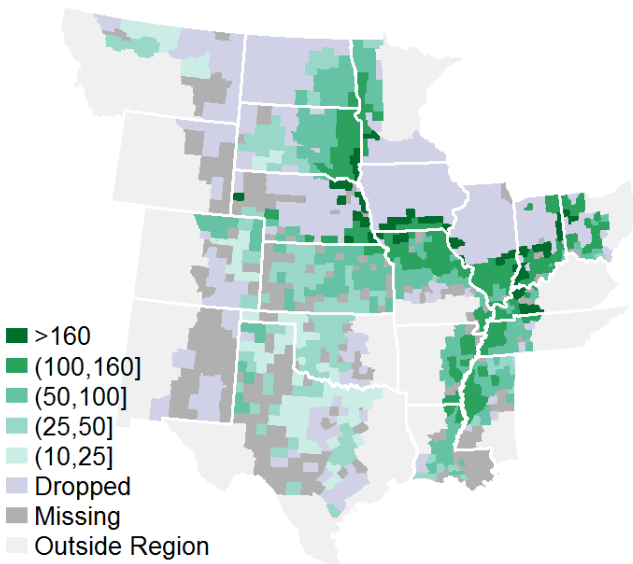

(a) Cash Rent

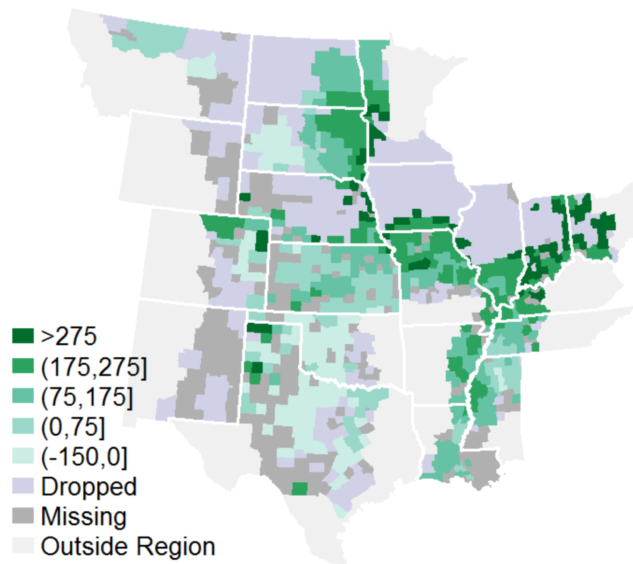

(b) Market Returns

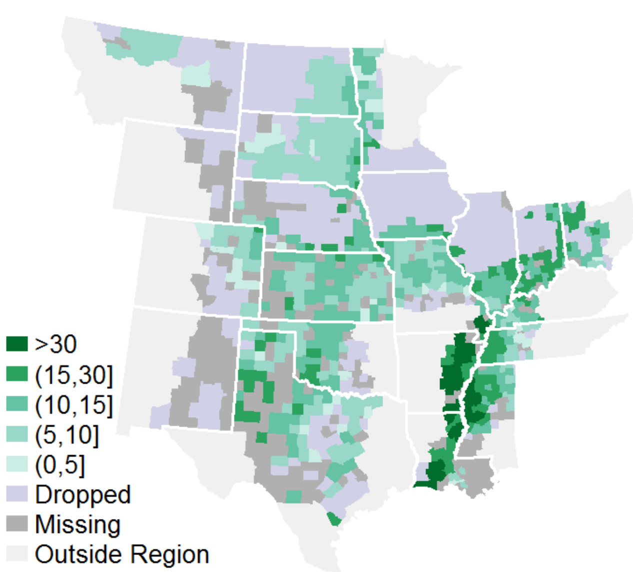

(c) Direct Payments

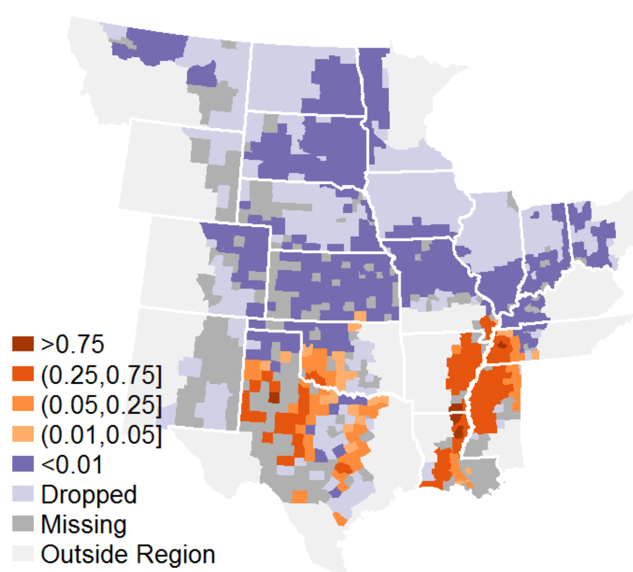

(d) Proportion Cotton or Rice Base Acres

Figure 1: Maps of Key Variables

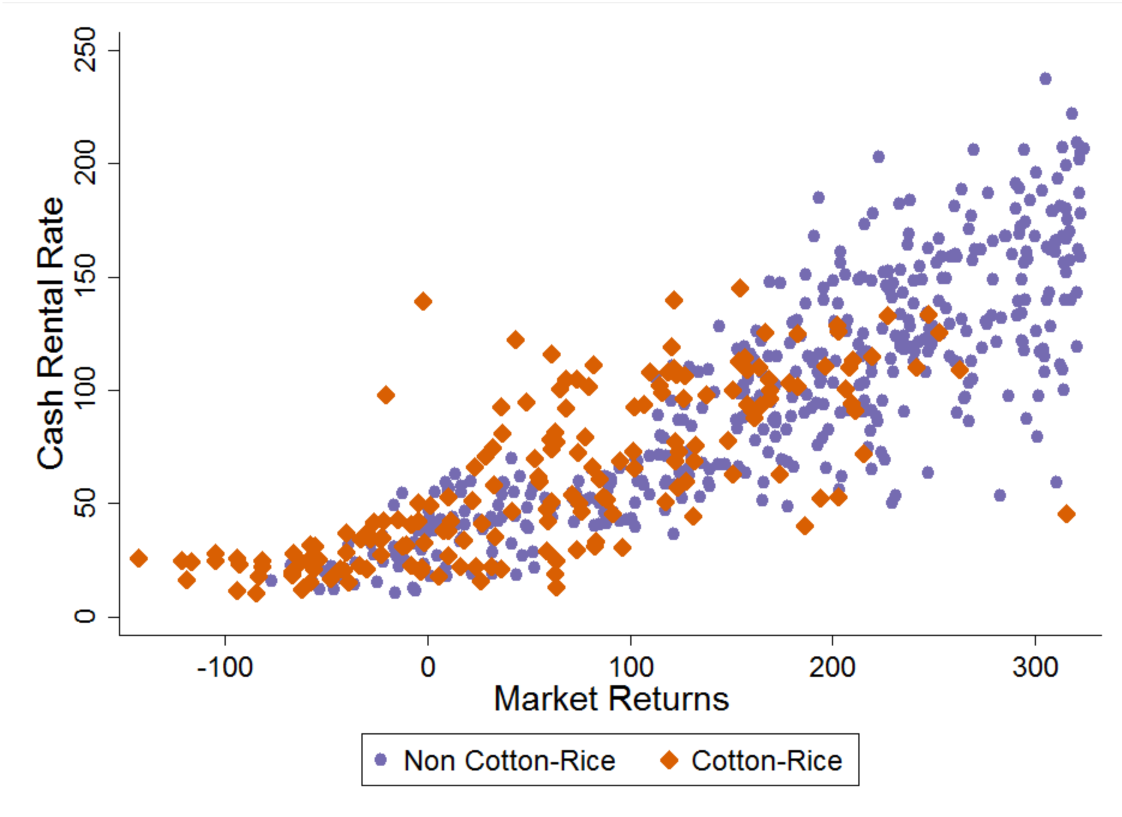

Figure 2: Rents and Market Returns

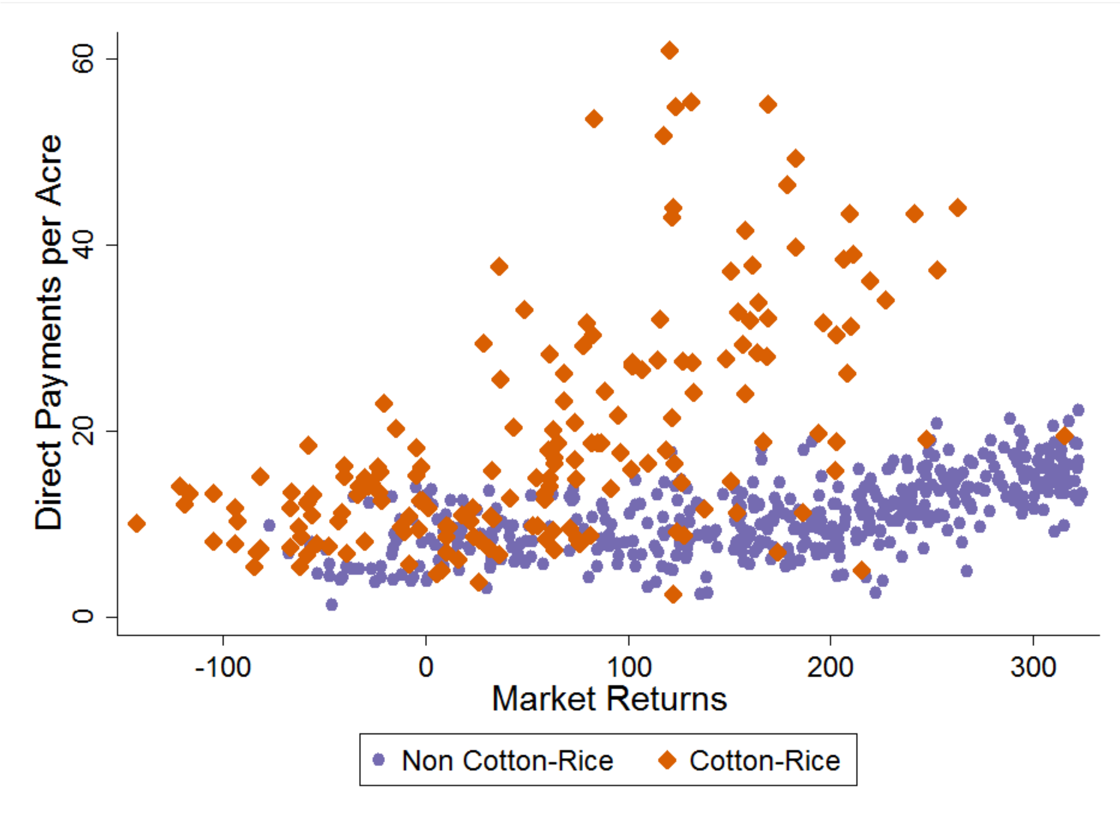

Figure 3: Direct Payments and Market Returns

Supplementary Appendix

Table A1 shows 2SLS results when we restrict our sample in different ways. Results in column (1) restrict the sample to those counties where the crop acreage used to calculate the market returns is at least half as large as total cropland area for the county. In the paper, we include all counties where the crop acreage used to calculate the market returns is at least 25% as large as total cropland area. Column (1) uses 519 observations instead of the 639 observations in the main paper. Column (2) only considers counties with market returns less than $275/acre instead of $325 used in the paper. Column (2) only eliminates one county with substantial cotton or rice base acreage but 84 counties with little cotton or rice base (see figure 2 in the paper). Column (3) only considers counties with market returns between -$80/acre and $275/acre. Figure 2 in the paper shows that only counties with greater than 1% cotton or rice base acres have market returns less than -$80/acre and mostly only counties with negligible cotton or rice base have market returns greater than $275/acre. So column (3) considers the portion of the sample with the most overlap. Column (4) includes all counties with nonmissing data. That is, column (4) does not drop counties if crop acreage used to calculate market returns represents a small portion of cropland nor does it drop counties with market returns above a certain threshold. For brevity, we only report results with a linear control for market returns.

Table A2 shows 2SLS results when we use different methods of constructing variables used in the econometric analysis. Column (1) shows results when we use cropland used for crops rather than total cropland area to calculate variables. Cropland used for crops is calculated from the Census of Agriculture and is the sum of harvested, failed, and summer fallowed cropland. Cropland used for crops is then used (i) to calculate the share of cropland irrigated that affects the average rental rate, (ii) to calculate direct payments per acre, (iii) to calculate the share of cropland fallowed, (iv) to calculate the share of cropland with cotton or rice base acres, and (v) to calculate the share of cropland enrolled in the ACRE program. The problem with using cropland used for crops is that one of the three components is often not reported at the county-level to avoid disclosing individual information. Therefore, we have to assume zero failed or zero summer fallowed acres when it is not reported, but this may underestimate cropland area used for crops. In several cases, total base acres greatly exceeds acreage of cropland used for crops which does not seem likely. Therefore, we drop counties where total base acres are more than 1.5 times as large as acreage of cropland used for crops.

Columns (2) and (3) in table A2 use different historical periods to calculate the market returns. In the main paper, we average expected market returns over the 5-year period of 2008–2012. In column (2) we average expected market returns over the 4-year period of 2009–2012 and in column (3) we average over the period of 2010–2012.

Table A3 reports results if we use the rental rate from different years. In the main paper, we use rental rates from 2012. In column (1) we use rental rates in 2011 and market returns from the period 2007-2011. Similarly, columns (2) and (3) use rental rates from 2010 and 2009 and market returns from the respective 5-year periods.

Table A4 reports 2SLS results if we do not control for ACRE enrollment. These are the exact same results as in table 5 of the main paper except that the results in table A4 omit ACRE enrollment.

Table A1: 2SLS Estimates with Different Samples

	(1)	(2)	(3)	(4)
			-$80/acre<	
	Returns	Returns	Returns	
	>50% Cropland	<$275/acre	<$275/acre	All Data
Direct Payments	0.789	0.827	0.758	0.826
	(0.188)**	(0.170)**	(0.171)**	(0.181)**
	[0.237]**	[0.218]**	[0.217]**	[0.224]**
Market Returns	0.350	0.344	0.357	0.395
	(0.011)**	(0.011)**	(0.011)**	(0.009)**
	[0.012]**	[0.011]**	[0.012]**	[0.010]**
Proportion ACRE	33.351	26.527	27.166	56.565
	(10.358)**	(8.392)**	(8.280)**	(9.380)**
	[10.358]**	[8.397]**	[8.284]**	[9.415]**
Intercept	28.410	26.501	25.238	19.971
	(2.640)**	(1.969)**	(1.966)**	(1.880)**
	[2.808]**	[2.121]**	[2.130]**	[2.001]**
Observations	519	554	541	1002

Standard errors in parentheses represent heteroskedasticity-robust standard errors imposing the exlcusion restriction. Standard errors in brackets are calculated allowing for a potential violation of the exclusion restriction (Conley, Hansen, and Rossi 2012). Asterisks * and ** denote significance at the 10% and 5% levels, respectively.

Table A2: 2SLS Estimates with Different Methods of Constructing Variables

	(1) Cropland Used for Crops	(2) Returns 2009-2012	(3) Returns 2010-2012
Direct Payments	0.753	0.754	0.893
	(0.167)**	(0.158)**	(0.151)**
	[0.203]**	[0.210]**	[0.204]**
Market Returns	0.363	0.325	0.282
	(0.010)**	(0.009)**	(0.008)**
	[0.010]**	[0.010]**	[0.008]**
Proportion ACRE	25.901	31.875	30.314
	(7.411)**	(8.341)**	(8.289)**
	[7.413]**	[8.346]**	[8.293]**
Intercept	24.687	23.912	17.727
	(2.196)**	(1.897)**	(1.890)**
	[2.354]**	[2.060]**	[2.050]**
Observations	635	626	594

Standard errors in parentheses represent heteroskedasticity-robust standard errors imposing the exlcusion restriction. Standard errors in brackets are calculated allowing for a potential violation of the exclusion restriction (Conley, Hansen, and Rossi 2012).
Asterisks * and ** denote significance at the 10% and 5% levels, respectively.

Table A3: 2SLS Estimates with Rental Rates from Different Years

	(1) 2011	(2) 2010	(3) 2009
Direct Payments	1.018	1.366	1.312
	(0.169)**	(0.188)**	(0.189)**
	[0.221]**	[0.238]**	[0.233]**
Market Returns	0.341	0.344	0.398
	(0.010)**	(0.011)**	(0.013)**
	[0.011]**	[0.013]**	[0.015]**
Proportion ACRE	22.409	7.058	11.647
	(7.114)**	(6.729)	(7.000)*
	[7.120]**	[6.737]	[7.008]*
Intercept	26.742	30.442	37.744
	(1.969)**	(2.082)**	(2.295)**
	[2.138]**	[2.297]**	[2.498]**
Observations	637	598	692

Standard errors in parentheses represent heteroskedasticity-robust standard errors imposing the exlcusion restriction. Standard errors in brackets are calculated allowing for a potential violation of the exclusion restriction (Conley, Hansen, and Rossi 2012).
Asterisks * and ** denote significance at the 10% and 5% levels, respectively.

Table A4: 2SLS Estimates with no Control for ACRE Program Enrollment

	(1)	(2)	(3)
Direct Payments	0.646	0.708	0.683
	(0.168)**	(0.166)**	(0.172)**
	[0.217]**	[0.216]**	[0.220]**
Market Returns	0.361	0.265	0.244
	(0.010)**	(0.017)**	(0.011)**
	[0.011]**	[0.018]**	[0.012]**
Market Returns^2		0.000	0.001
		(0.000)**	(0.000)**
		[0.000]**	[0.000]**
Market Returns^3			-0.000
			(0.000)**
			[0.000]**
Intercept	29.665	30.099	28.602
	(1.805)**	(1.731)**	(1.714)**
	[1.982]**	[1.914]**	[1.923]**
Observations	639	639	639

Standard errors in parentheses represent heteroskedasticity-robust standard errors imposing the exlcusion restriction. Standard errors in brackets are calculated allowing for a potential violation of the exclusion restriction (Conley, Hansen, and Rossi 2012).
Asterisks * and ** denote significance at the 10% and 5% levels, respectively.