Landowners' Riches:

The Distribution of Agricultural Subsidies*

Barry K. Goodwin, Ashok K. Mishra,
and
François Ortalo-Magné

November 2002

PRELIMINARY, DO NOT CITE WITHOUT AUTHORS' PERMISSION

Abstract

The U.S. has a long history of providing generous support for the agricultural sector. A new omnibus package of farm legislation (the 2002 Farm Bill) will provide in excess of $190 billion in financial support to U.S. agriculture, an increase of $72 billion over existing programs. Our paper is concerned with the distribution of these benefits. Farm subsidies make agricultural production more profitable by increasing and stabilizing farm prices and incomes. If these benefits are expected to persist, farm land values should capture the subsidy benefits. We use a large sample of individual farm land values to investigate the extent of this capitalization of benefits. Our results confirm that subsidies have a very significant impact on farm land values and thus suggest that landowners are the real benefactors of farm programs. As land is exchanged, new owners will pay prices that reflect these benefits, leaving the benefits of farm programs in the hands of former owners that may be exiting production. Approximately 45% of U.S. farmland is operated by someone other than the owner. We report evidence that owners benefit not only from capital gains but also from lease rates which incorporate a significant portion of agricultural payments even if the farm legislation mandates that benefits must be allocated to producers. Finally, we examine rental agreements for farmers that rent land on both a cash and share basis. We find evidence that farm programs that are meant to stabilize farm prices provide a valuable insurance benefit.

---

*We thank Erzo G.J. Luttmer, Jim Shilling and Fred Sterbenz for useful discussions and comments. Steve Malpezzi helped us with the data collection. Goodwin is Andersons Professor at the Ohio State University, Mishra is an economist with the Economic Research Service of the USDA, and Ortalo-Magné is Jean Monnet Lecturer at the London School of Economics and visiting associate professor at University of Wisconsin. Direct correspondence to Goodwin at 229 AA, 2120 Fyffe Road, The Ohio State University, Columbus, OH 43210, (614) 688-4138, E-mail: `goodwin.112@osu.edu`.

Landowners' Riches:
The Distribution of Agricultural Subsidies

1 Introduction

A recent news report posed the following question. What do basketball star Scottie Pippen, publisher Larry Flynt, and stockbroker Charles Schwab all have in common? The surprising answer is that all are recipients of farm program subsidies.[1] Other recipients include former Worldcom chief Bernard Ebbers, part-owner of a farm that received nearly $4 million in subsidies between 1996 and 2000, the Chevron corporation, the Caterpillar corporation, and several other Fortune 500 companies. Between 1996 and 2000, farm subsidy checks were mailed to 31 addresses in Beverly Hills, 31 addresses in Vail, 752 addresses in Miami, and 803 addresses in Washington, D.C. The fact that support for U.S. "farmers" is often directed to individuals and corporations that seem to be some distance from the farm has been the topic of considerable debate in recent years, in particular since congressional support for U.S. agriculture continues to expand. President Bush signed a new omnibus package of farm program support on May 13, 2002. Over the next ten years, the new Farm Bill will transfer about $190 billion from U.S. taxpayers to the farm sector. On a per year basis, this is more than what the federal government spends on community and regional development, and about the same magnitude as expenditures on general science, space and technology.[2]

To the extent that eligibility for government benefits is tied to the ownership or operation of certain assets, the market values of these assets will reflect expected future benefits. Such is the case with farmland. USDA statistics indicate that 45.3% of U.S. farmland is operated by someone other than the owner (USDA-NASS, 1999). Mishra et al. (2002) report that, contrary to conventional wisdom, most agricultural landlords (57%) are non-farm corporations or individuals that work in or are retired from nonfarm-related activities.

[1] "Farm Subsidies Help Those Who Help Themselves," a Fox News report by William LaJeunesse, July 15, 2002. This article is available from `http://www.foxnews.com/story/0,2933,57602,00.html`. These statistics are all drawn from the Environmental Working Group's farm subsidy database (`www.ewg.org`).

[2] Source: Budget of the United States Government, http://w3.access.gpo.gov/usbudget/.

In light of these facts, a fundamental question arises regarding the distribution of farm support programs and the extent to which those who operate the farms actually receive the benefits. This is a critical issue, not only for policymakers but also for farm operators who should understand the extent of their gain from the various programs they tend to support.

The relevant question is, of course, who are the policies intended to benefit? The capture of agricultural benefits by farmland is problematic if the policies aim to support farmers and these farmers do not own their land when the policies are announced. To the extent that (young) expanding farmers are paying for the expected policy benefits in the farm assets they acquire, the present value of future benefits is captured by the (old) sellers. New owners only benefit from surprise increases in public transfers. Given the large share of U.S. farm land that is operated by tenant farmers, the extent to which lease rates capture program benefits is also important to the distribution of these benefits.

The concern with the capture of agricultural policy benefits by the initial land owners is not new. A number of papers have attempted to estimate the capitalization of aggregate agricultural transfers into farmland values.[3] These papers suffer from a number of shortcomings which we are able to address here through an empirical analysis of a unique set of farm-level data. We contribute to the understanding of the distribution of farm subsidies in several ways. First, we are able to investigate the differential impact of the principal farm programs because we are able to observe the breakdown of government payments at both the farm and the county level. Second, because we know the location of each farm, we are able to control for non-agricultural pressures on the land and determine how they affect its value. Third, we observe not only land values but also the terms of lease arrangements and rates. This puts us in a unique position to be able to assess directly the extent to which owners and farmer operators share the benefits of various agricultural programs, a useful complement to the indirect assessment we obtain from investigating land values. Finally, variations in the

[3]See Barnard et al. (1997), Goodwin and Ortalo-Magné (1992), Ryan et al. (2001), Shertz and Johnston (1997), Shoemaker et al. (1990), and Weersink et al. (1999). These papers only examine aggregate policy effects on land values, thus ignoring the myriad effects of different programs. In addition, the extraction of policy benefits through lease arrangements has not been previously evaluated.

difference between cash lease rates and share lease rates enables us to investigate the extent to which the market values the insurance features built into some farm programs, a feature ignored by the literature.

Our analysis makes use of a data set drawn from an annual survey of approximately 10,000 farms per year over the 1998-2001 period. This period was characterized by a variety of different farm programs, including some which were not connected in any way to market conditions or production, at least in theory. At the other extreme are output price-support payments which are intimately tied to contemporaneous market conditions. We find that payments that are decoupled from output and supposed to be transitory yield the smallest effects on land values. Payments that may signal future benefits, even in cases where they are not a permanent part of farm legislation, have stronger effects. Price-support payments have the strongest effects.

U.S. farm legislation typically intends benefits to be “shared” between the owner and operator of a farm. Under cash lease arrangements, the entire subsidy is sent to the operator. However, the law does not regulate lease rates; they are set by the market. Our empirical analysis indicates that owners extract a large proportion of farm benefits from tenants through the lease rates. From the study of lease rates, we also find that programs with strong insurance objectives, such as output-price support payments, significantly affect the gap between cash and share lease rates. In particular, the share rate premium is significantly diminished by programs that serve to lower the risk associated with uncertain farm earnings. This finding provides direct evidence of the land market pricing the insurance component of agricultural policy.

Accounting for the benefits of decreased earnings volatility raises two issues with the traditional approach to the assessment of the contribution of agricultural policy to farm land values. First, the insurance feature of several governmental programs raises questions about the traditional implicit assumption that a dollar of transfer today conveys the same information about future transfers, regardless of market conditions and local agricultural output characteristics. Instead, a low price support payment this year may be due to high world market prices and thus in no way indicates a decrease in the expected stream of long run benefits from the price support program. Second, those government transfers whose level are negatively correlated with farm earnings from the market decrease the volatility of farm land returns. They must therefore decrease the discount rate required to hold farm land and thus the discount rate applied to earnings from the market. Hence, regression estimates of the contribution of market earnings to the value of land depend on the policy environment. In particular, it is wrong to assume that such estimates would not change to reflect a more volatile environment if price support programs were to be dismantled.

The remainder of our paper is organized as follows. The next section gives a brief overview of the history and nature of U.S. farm programs. Section three discusses issues pertaining to model specification, estimation, and measurement of the relevant variables. The fourth section presents the results of our empirical analysis and discusses their implication for the distribution of agricultural policy benefits. The final section offers some concluding remarks.

2 A Brief Overview of U.S. Farm Policy

Most U.S. farm programs have their origins in the New Deal legislation of the Great Depression. A variety of price and income support programs have been used over time to increase and stabilize farm earnings. These programs are revised approximately every 6 years by an omnibus "Farm Bill" package of legislation. In addition to this major package of farm programs, support is provided through a number of other legislative channels. This is the case with farm programs such as crop insurance and conservation measures. On a regular basis, agriculture also benefits from ad hoc support (though emergency bills) that is not a part of any budgeted legislation.

Over most of its history, U.S. agricultural policy has used price supports coupled with production controls, with the declared objective to provide income support to the farm sector. Some support was made on the basis of a need for "parity" with the high relative agricultural prices of 1910.[4] In more recent times, support was provided only to program crops (corn, wheat, cotton, rice, grain sorghum, rye, barley, and oats). Deficiency payments, determined by the difference between market and target prices, were paid to producers on the basis of their "base acreage." This base acreage reflected historical production (in most cases, acreage during the 1980s). The fact that price supports were tied to historical production patterns implied a lack of planting flexibility for producers. In addition, soybeans, a major U.S. crop, was largely omitted in provisions for support due to the fact that it was not an important crop when most farm programs began.

In 1996, Congress agreed to what was intended to be a major overhaul of U.S. farm policy– the Farm and Agricultural Improvement and Reform (FAIR) Act. This Act is also known as the "Agricultural Market Transition Act" or AMTA. The nomenclature "Reform" and "Market Transition" was meant to indicate a major shift in policy away from government involvement and toward market oriented policies. Eligibility for price support was no longer based upon historical production—producers were free to plant whatever crops they desired. Soybeans were made eligible for price supports, which were now provided through the "Loan Deficiency Payment" (LDP) program. LDP payments were made on the basis of the difference between market and support prices (the loan rates). The rhetoric accompanying the Act implied, in principle at least, that the legislation signaled a transition to an environment with limited government support. To compensate producers over this transition, a program of direct payments to those producers with base acreage (historical rights to program benefits) was instituted. These payments were known as AMTA or Production Flexibility Contract (PFC) payments. By design, AMTA payments were completely decoupled from the market—the only requirement for receiving AMTA payments was that the producer (or landowner) had to have base acreage. Eligibility for such payments in no way depended upon current production patterns. In some cases, payments were made on land no longer in production. The AMTA payments were set to decline each year until the FAIR Act expired in 2002.

[4] Though any link with market and production conditions in 1910 would seem difficult to make, arguments in favor of such "parity" pricing is still heard on occasion in farm policy debates.

Ad hoc disaster assistance has been a fixture in U.S. agricultural policy for many years. Periods of drought or poor market conditions frequently trigger ad hoc assistance labelled disaster payments. Under provisions of other farm legislation (the Crop Insurance Reform Act of 1994), Congress stated an intention to make subsidized insurance the only mechanism for providing disaster relief.[5] However, localized droughts and low market prices led Congress to rapidly retreat from this position and conclude that the support provided to farmers under the FAIR Act was not sufficient. Ad hoc assistance, in the form of yield compensations and payments for low market prices, termed "market loss assistance" (MLA), were then instituted. Market loss assistance payments totaled over $19 billion between 1998 and 2001. These payments were made on the basis of eligibility for AMTA payment benefits. Once again, since AMTA payments are based on historical production, there was no explicit eligibility requirement that recipients be actively involved in agricultural production.

A number of other programs have been important to agricultural policy. For example, a considerable amount of farm land (approximately 30 million acres) has been removed from production through the Conservation Reserve Program (CRP). The CRP pays producers annual rents to place their land in reserve under a ten-year lease agreement. Some of this land has base acreage associated with it, thus leading to AMTA and MLA payments on land in conservation reserve. In order to be eligible for the CRP program, land must be "erodible" and environmentally fragile. Such land is typically of a lower value in terms of crop production.

According to analysis by the Congressional Budget Office (CBO), the 2002 Farm Bill will increase federal spending on agriculture by $72 billion over the next ten years. Congressional debate over the 2002 farm legislation and the generous level of support that emerged from these deliberations has made clear Congress's intent to continue taxpayer support for agriculture. Lawmakers chose to provide generous increases in support with no real limits

[5]As an aside, an interesting policy situation exists for crop insurance, which recently has returned $1.90 in indemnity payments for every dollar of premiums paid by farmers. This program, also in existence since the 1930s, runs hand-in-hand with ad hoc disaster assistance—a form of free insurance. Note that disaster assistance is an obvious impediment to a well-functioning insurance program.

on the amount of payments an individual can receive; the issue of the distribution of farm subsidies remains.

3 Modelling Framework

3.1 The Income Approach to Farm Land Valuation

All government transfers help the farmers in at least one of two ways: by raising the returns to farming and by decreasing the volatility of these returns. The LDP and DP programs have major insurance components. The $AMTA$ payments are lump sum transfers determined by farmers' activities prior to their implementation. The same is true with CRP payments; they are lump sum additions to the return of farming that are uncorrelated with present or future earnings from the market. In addition to all these transfers, farm land also gives the farmer the opportunity to generate non-agricultural earnings. The jackpot is to own land in an area under strong urban pressure with friendly zoning authorities, hence providing the opportunity to realize substantial capital gains by converting the land to residential or commercial use.

The value of a parcel of land is the present discounted value of expected cash flows from agricultural activities plus the value of the option to convert the land to non-agricultural use.

$$V_0 = \mathrm{E}_0 \left[\sum_{t=1}^{\infty} \frac{MKT_t + \mathrm{LDP}_t + DP_t + \mathrm{AMTA}_t + \mathrm{CRP}_t}{(1+r)^t} \right] + CONV_0, \quad (1)$$

where MKT and DP denote earnings from the market and from disaster payments, $CONV$ is the value of the conversion option, and r is the discount factor. The discount factor reflects the risk of the overall portfolio of individual streams of cash flow. This risk is not simply the sum of the individual risks because of the non-zero covariance, by design, between MKT payments, LDP and DP.

As mentioned earlier, AMTA and CRP are, for the most part, lump sum transfers whose levels are independent of current and future earnings from MKT, LDP and DP, and from each other. We can therefore rewrite equation (1) as

$$V_0 = \mathrm{E}_0 \left[\sum_{t=1}^{\infty} \frac{MKT_t + \mathrm{LDP}_t + DP_t}{(1+r_1)^t} + \frac{\mathrm{AMTA}_t}{(1+r_2)^t} + \frac{\mathrm{CRP}_t}{(1+r_3)^t} \right] + CONV_0, \tag{2}$$

where r_1, r_2 and r_3 denotes the discount factors for output related earnings, AMTA payments, and CRP payments, respectively.

Implicit in equation (2) is the assumption of a constant discount rate. If we are willing to assume that farm land buyers and sellers expect the various earnings to grow at a constant rate, then the regression coefficients we will obtain will be the inverse of the capitalization rates, or cap rates. The valuation formula can indeed be re-arranged as

$$V_0 = \mathrm{E}_0 \left[\frac{MKT_1}{\kappa_1} + \frac{\mathrm{LDP}_1}{\kappa_2} + \frac{DP_1}{\kappa_3} + \frac{\mathrm{AMTA}_1}{\kappa_4} + \frac{\mathrm{CRP}_1}{\kappa_5} \right] + CONV_0 \tag{3}$$

where the cap rates are denoted by $\kappa^{\cdot}$. It is easy to check that if a stream of cash flows is expected to grow at the constant rate g and is discounted at the constant rate r, then the cap rate κ satisfies $\kappa = r - g$.

To estimate the contribution of each source of value in equation (3), we need estimates of expected next period cash flows for each source of agricultural earnings. This raises a serious measurement issue. As mentioned above, it has been largely ignored in the literature which tends to rely on current payments as an indicator of future payments. This is the issue to which we now turn.

3.2 Measuring Expected Cash Flows

Let us suppose that agents correctly assess the true determinants of land values but the econometrician, working with actual realizations of policy outcomes from year to year, is unable to observe these determinants. Instead the econometrician relates the observable annual realizations of market and policy outcomes to land prices. In this case, the econometrician is confronted with the classical errors in the explanatory variables problem. Errors-in-variables results in an attenuation bias that forces coefficients toward zero and thus yields inconsistent estimates.[6] This problem is compounded by the fact that the government operates more than one program of payments, hence suggesting that traditional empirical approaches suffer from multiple explanatory variables observed with error.

A complicating factor arises in that the errors applying to observed policy benefits may be correlated in a typical sample. This correlation may assume two different forms—correlation of the errors across different programs (for a given farm) and correlation of errors across different farms in a sample. Both circumstances are likely to exist when one considers a pooled cross-section of farms (as is the case in our empirical analysis). Consider a case of two programs—price supports and market loss assistance payments. The extent of support provided from the government is likely to vary considerably from year to year according to market conditions. Low price years realize larger payments for both programs. Thus, the errors associated with using realized benefits are likely to be highly correlated across the programs. The correlation could also be negative. Consider the case of yield disaster relief and price supports. In low yield years, market prices are likely to be high and thus price support payments will be low, though disaster benefits will be higher to compensate for the production shortfalls.

Another form of correlation is likely to be relevant when a pooled sample of individual farms is considered. Since realized program benefits are dependent upon aggregate market conditions, the errors are likely to be highly correlated across observational units (farms) in a given year. In a sample consisting of only a few years of data, the correlation across farms increases the estimation error and may further exaggerate the bias; year-to-year shocks may not average out when only a few years are observed.

Furthermore, if realizations are highly correlated across units within a year, parameter estimates may shift considerably from year to year. If only a few years are observed, the

[6]This problem is analogous to the standard omitted variable problem, where the omitted factor is the difference between what is observed and the true, latent value.

estimates from a pooled sample may be sensitive to events in the years observed and thus may vary substantially across years and be more variable in a pooled analysis.[7]

The standard approach to addressing this problem is to obtain instruments or proxy variables for those latent variables that are measured with error. An instrument should be correlated with the variable of interest but uncorrelated with the error pertaining to the observation. We represent the expected payment benefits by constructing average values of each relevant policy variable over the preceding four years. This approach raises one complicating factor. As we discuss in detail below, our data set is not a true panel in that a different set of farms is sampled each year—meaning that repeated observations for an individual farm are not available.

To represent expected payment benefits, we utilize the four-year average value of real payments per acre *in the county* where the individual farm is located.[8] We argue that this is a superior measure of long-run expected benefits as compared to realized payments because values in the county more closely represent the long-run potential benefits associated with agricultural policy. Payments on an individual farm, in contrast, may reflect individual policy choices and characteristics of the farm operation. Transfer of the land to a new operator may result in different subsidy realizations (for example, because of a different crop mix) which are better represented using county-level averages.[9]

We adopt a number of different historical averages to represent expected policy benefits. We use a four year average of county level total payments in our aggregate policy models. In contrast, because LDPs were not the main instrument for providing price support prior to the 1996 Farm Bill, we use a two-year average for LDP payments at the county level. We should note that this errors in variables problem does not apply to all sources of government

[7] See Goodwin et al. (forthcoming) for a quantitative assessment of this issue in the farm land valuation context.

[8] A standard instrumental variables estimation approach is also feasible, though the fact that payment realizations in any given year may be very weakly tied to long run expected benefits makes the utility of such an approach limited. This problem is exacerbated in a short sample when realizations are highly correlated in the cross section, as is true in our application.

[9] Observations for an individual farmer in a particular year might reflect crop rotation patterns. We expect county level acreage to be more reflective of the expected crop mix.

subsidies. Subsidies provided through AMTA payments and rents earned on land enrolled in the CRP program are known with certainty a priori. It is only those payments that are triggered by market and production events (price supports and disaster payments) that must be proxied.

3.3 Data

The primary source of our farm-level data is a data set collected from a large sample of farms through the USDA's National Agricultural Statistics Service (NASS) Agricultural Resource Management Survey (ARMS) project. The ARMS survey is a large probability-weighted, stratified sample of about 8,000-10,000 U.S. farms each year. The survey collected detailed government payments information for individual farm program benefits as well as extensive farm and operator characteristics for the years 1998-2001. Thus, our empirical analysis focuses on these years. All monetary values in our sample were adjusted to 1999 equivalent real values by deflating by the producer price index. Given the short nature of our sample, such deflation had only minor effects on the results.

Besides detailed farm earnings, the survey also reports farm land value. Farm operators are asked to estimate the market value on December 31 of each year of their land, dwellings, and other farm buildings and structures. We restrict our attention to the value of land only (excluding trees and orchards).[10]

To focus on policies directed at crop farms, we excluded any farms from our analysis for which livestock product sales accounted for over 50% of overall farm sales. We also excluded farms for which incomplete data were available. This left us with a small number of extreme outlier observations (land values exceeding of $10,000 per acre). Such extreme observations represent non-typical agricultural properties, such as vineyards and properties with characteristics (e.g., river-side properties) not recorded in the survey. Only a very small number of observations (0.8% of our total) were excluded on this basis. In the end, we were left with 13,603 individual farm-level observations.

[10] Farmers do not have any incentive to distort their response to the survey.

A variety of sources were used to collect county-level observations on crop acreage and state level prices (NASS statistics) and data relevant to county population and housing market trends (Census data). Summary statistics and definitions for the key variables of our analysis are presented in Table 1. Population weights were used in deriving summary statistics to account for the stratified nature of the ARMS survey. Table 1 also includes farm-average payment receipts, taken from unpublished Farm Service Agency (USDA-FSA) records, across years and program types.

To the variables aimed at capturing expected cash flows from farming, we added three factors intended to represent the additional value of land in areas facing non-agricultural pressures. First, to represent residential housing pressures, we include the total value of housing permits issued in the county in which the farm is located. The permit data were collected from the U.S. Census Bureau. The permits apply to all forms of residential housing, both single family and multi-family dwellings, and describe the total value of construction. We also include a series of discrete indicator variables (obtained from the USDA) that represent the extent of urbanization for each county. The ordinal ranking ranges from 1=rural to 5=urban. In that very few farms were located in the most urban counties (less than 0.3%), we combined urban categories 4 and 5 into a single category. Finally, we considered population growth rates and the population density (people per square mile) in each county in the preceding year.

We are interested in evaluating the differential effects of benefits provided by the government versus those returns generated by the market. Of course, a risk-neutral farmer will not care where a dollar comes from, though alternative sources of revenue may have different levels of risk, thus affecting the preferences of a risk-averse farmer. We acknowledge at the outset that any representation of market earnings should not be interpreted as a measure of the market returns that would be generated in the *absence* of farm policy. Returns in such a situation are difficult to assess, especially in light of the long history of government involvement in U.S. agriculture. Likewise, the relevance of such a consideration is limited—it is unlikely, to the authors at least, that the U.S. government will completely remove policies that currently support agriculture. Having acknowledged these limitations, we construct a measure of net returns from the market using a combination of data drawn from several sources.

Market revenues are straightforward to measure using county-level acreage allocations and state level prices. We use weights constructed from a consideration of the allocation of crop land within a county across different crops to evaluate the revenue potential of each acre. Costs are much more difficult to measure. For wheat, barley, and cotton, we use national, annual average per-acre costs drawn from our survey of farms. For corn and soybeans, we use cost budgets collected from the Iowa State University Extension Service. For grain sorghum (a minor crop), we use a cost figure collected from budgets generated by the University of Arkansas Extension Service. These crops overwhelmingly account for the the majority of U.S. crop-farm earnings and thus other minor crops were not included in the cost and returns measures. In every case, costs figures exclude the charge for land. These figures are weighted using county-level acreages in the same fashion as for earnings. The difference between county-level revenues and costs are used to construct a measure of returns from the market. It should be noted that, at least for our period of study, these returns are especially variable across counties (according to crop mix) but are less volatile over time. The county level costs and returns for each crop are adjusted to reflect higher or lower yields on the farm relative to the county by multiplying each crop component by the ratio of the farm's yield over the county average yield. This adjustment reflects the fact that some farms are better than others in terms of their yield performance.[11]

4 Empirical Results

A number of important econometric issues underlie our empirical analysis. An important characteristic of the ARMS data relates to the stratified nature of the sampling used to

[11]Some farmers may achieve higher earnings than others from the same parcel of land due to their own knowhow. We have no way of disentangling farm versus farmer specific effects so we implicitly assume that any difference between the performance of the farm and average performance of farms in the county is due to unobserved differences in the productivity of the land.

collect the data. Two estimation approaches have been suggested for problems such as this involving stratification. The simplest involves a jacknife procedure, where the estimation data are split into a fixed number of subsamples and the estimation is repeated with each subsample omitted. Under the jacknife approach, the sample is divided into m subsamples.[12] The model of interest is estimated m times using weighted regression procedures with each of the respective subsamples omitted from the estimation data. A simple expression for the variance is then taken by considering the variability of the estimates across each of the replicated estimates. The jacknife weights included in the ARMS data are constructed for use with the entire sample. In that we are using a subset of the overall sample, these weights are not appropriate for our use since the pre-defined jacknife groupings would leave one with unbalanced jacknifed groups.

An alternative approach involves repeated sampling from the estimation data in a bootstrapping scheme. Ideally, rather than random sampling from the entire estimation sample, an appropriate approach to obtaining unbiased and efficient estimation results involves random sampling from individual strata (see, for example, Deaton (1997)). In the ARMS data, however, this is not possible since the strata are concealed from the researcher. The database does, however, contain a population weighting factor, representing the number of farms in the population (i.e., all U.S. farms) represented by each individual observation. This can be used in a probability-weighted sampling scheme whereby the likelihood of being selected in any given replication is proportional to the number of observations in the population represented by each individual AMRS observation.

We utilize a probability-weighted bootstrapping procedure. The specific estimation approach involves selecting N observations (where N is the size of the survey sample) from the sample data. The data are sampled with replacement according to the probability rule described above.[13] The models are estimated using the pseudo sample of data. This process is repeated a large number of times and estimates of the parameters and their variances are

[12] Estimation programs created by the USDA use 15 subsamples.

[13] It should be acknowledged that our approach may result in less efficient estimates than would be the case were sampling from individual strata possible. This could occur in cases where inferences are being made about variables used in designing the stratification scheme in that such information is being ignored by

given by the mean and variance of the replicated estimates. We utilize 2,500 replications in the applications which follow.

4.1 Land Prices

Our analysis of the determinants of land values is conducted in three segments. In the first, we consider models that aggregate all program payments into a single category. Such a model is useful in that it provides a summary of the impacts of additional federal subsidy dollars on land values at the margin. Two version of this model are considered. The first uses actual, observed payments for each farm. The second uses county-level data to assess the total, expected per-acre receipts from farm program payments. Note again that expected payments are represented using the average in the county over the preceding four year period. The results are presented in Table 2. The model using observed farm-level payments (Model I) indicates that $1 of farm payments tends to add $5.40 per-acre to the value of farm land. The effect is highly statistically significant. This suggests a considerable degree of capitalization of farm benefits into land values. This occurs in spite of the fact that some of these benefits were unplanned (ad hoc, off-budget, and not a part of the Farm Bill legislation) and others were, in principle at least, intended to be transitory. It is interesting to compare the effects of government payments on land values to the effects of market returns. The results indicate that an additional $1 obtained from the market would raise land values by $5.56, a figure very similar to what is implied for subsidy payments.

We have argued that the use of observed payments may result in an attenuation bias that forces the implied capitalization rates toward zero. As an alternative, we have argued that a measure of expected payments may be preferred. Model II replaces the total realized payment measure with the four-year average measure noted above. As expected, the results suggest larger effects of agricultural policy on agricultural land values. An additional $1 of government payments raises land values by $6.09 per acre.

---

not drawing from individual strata. To the extent that this is relevant to our analysis, the t-ratios reported below represent conservative estimates.

Models I and II also confirm the importance of many other factors in the determination of agricultural land values. Both models suggest that a higher population growth rates and a greater density of population tends to raise land values. A 1% increase in the population growth rate adds $64 per acre to the value of farm land. Likewise, each additional person per square mile in the country adds $1.85 in value to an acre of land. Urban pressures also exert a strong influence on land values. The most rural counties were chosen as our default category. When compared to these rural counties, counties with urban indicators of 2, 3, and 4-5, had land values that were $312, $499, and $304 greater per acre, respectively. On top of this, we find that each additional $10 million in new construction adds about $4.77-$5.01 per acre to farm land values. Although these urban effects are interesting in their own right, it is important that they be accounted for (a step that has generally been neglected in previous analyses) in order to obtain accurate measures of the policy effects on land values.[14]

A second segment of the analysis breaks out the overall government payments into their individual components, generated from different programs. Recall that we have argued that it is likely that different policies will have different effects on land values. Models III and IV use actual and our measure of expected payments, respectively. It is also possible that policies may have different effects in different regions having varied agricultural characteristics. Thus, Model V is confined to the corn belt region—the major agricultural region of the U.S.[15] The model of observed payments suggests that an additional dollar of LDP payments (price supports) will increase land values by $7.02 per acre. When realized payments are replaced by the two year average at the county level, the LDP effect rises to $9.05, again perhaps reflecting the attenuation biases inherent in using observed payments in any given year on an individual farm. Using observed disaster and market loss assistance payments, an additional dollar of payments raises land values by $7.16 per acre.

[14] Hardie et al. estimate the effects of urban pressure on agricultural land. They are not concerned, however, with the contribution of agricultural policy to the returns from land.

[15] To be more precise, we use the "Heartland" region, a designation based upon the USDA's farm resource region groupings. These regions are defined to be relatively homogenous in terms of agricultural characteristics.

When county-average disaster payments are used, a very different effect on land values is indicated. Higher levels of disaster relief tend to be correlated with lower land values. This result is somewhat puzzling, though it most certainly is driven by the fact that disaster payments are a signal of localized problems—they are targeted to areas that have realized revenue shortfalls. This may make it difficult to disentangle the individual effects.[16]

The observed AMTA payments also exert a significant effect on agricultural land values. An additional dollar in payments raises land values by \$3.80-\$4.40 per acre. If land market agents truly believed that these payments were transitory, as the 1996 legislation implied, this would seem to be too large of an effect. It is likely that AMTA payments were a signal of future benefits to be paid on a decoupled basis. Indeed, in its generosity, Congress not only continued these payments under the 2002 legislation. More importantly, the new Farm Bill made soybean acreage eligible for AMTA payments. Thus, our results suggest that agents anticipated such legislative actions—threats to terminate this avenue of support with the expiration of the 1996 legislation were heavily discounted.

Another interesting finding is that land enrolled in the conservation reserve program has a significantly lower value, in spite of the fact that the government pays rents on such land. The models suggest that a dollar of CRP payments lowers land values by \$15.82-\$17.20 per acre. This negative effect is certainly in line with expectations since enrollment of such land necessarily removes it from production for a ten year period. In addition, land that is eligible for enrollment is likely to be of a lower quality and more erodible than land that is not eligible.

The effects of urban pressures, population growth, housing pressures, and market returns are very similar to what was implied by the aggregated policy models. In every case, urban pressures play an important role in determining agricultural land values. In every case, the effects are statistically significant.

[16] In an extension to this model, we are attempting to disentangle these effects by separating yield disaster payments from market loss assistance disaster payments, which may have different effects. Such a separation is not possible at the farm level since the survey lumps all disaster payments together.

Different farms face different opportunities in terms of crops but also in terms of government programs. In our language, farm land in different farms might be made up of different securities and hence represent a different portfolio. This suggest that coefficients may differ if we focus our estimation on a restricted, homogeneous group of farms. Model V uses only data for the corn belt region of the U.S. (comprised of Iowa, Illinois, Indiana, Ohio, and small portions of surrounding states). The results are quite similar, especially for policy effects. An additional dollar of LDP payments adds $7.20 per acre to land values. Disaster payments reduce land values and CRP rents are associated with substantially lower land values. AMTA payments have a significant positive effect on land values in the corn belt, raising values by $5.51 per acre. The effects of population and population growth are somewhat smaller in the corn belt. Likewise, with the exception of the most urban areas, the urban indicators reflect a smaller impact on agricultural land values.

In all, the results confirm that government payments exert a significant effect on land values. The (marginal) rates of capitalization suggest that in the current policy context, a dollar in benefits typically raises land values by $5-$7 per acre. This response certainly suggest that agents expect these benefits to be sustained for some time. In terms of the implications for the distribution of farm program benefits, our results confirm that a substantial share of the benefits is captured by landowners. Recall that, in many cases, landowners may be a very different entity than farmers. Farmers wishing to expand or enter production will realize much smaller benefits than the policy rhetoric tries to substantiate.

4.2 Lease Rates

The results on farm land values provide evidence that land captures policy benefits as land values are enhanced by the subsidies. When the farm operators owns his land, the transfers go to him. Likewise in the absence of effective limits on payments (there are none) the larger his farm the more he gets.[17] However, as we have noted above, about 45% of U.S. farmland is operated by someone other than the owner.

This raises the important question—how do the generous provisions for support of agriculture affect those farmers that rent the land used in production? Likewise, how much of the support goes to landowners? Again, the stated intent of the legislation is a "fair and equitable" sharing of program payments, with an owner that shares no risk (i.e., rents under cash lease arrangements) receiving none of the benefits. The real answer to this question lies in an evaluation of the terms of lease arrangements—do lease rates reflect policy benefits? If, as we have demonstrated, the value of land is increased by policy transfers, given that value of land is a present discounted value of expected cash flows plus an option to convert, one would expect that lease rates reflect payments from the government. Lease rates provide direct evidence on the proportion of farm payments passed on to landlords, something much more difficult to assess from land values.

For those farmers in our sample that were engaged in renting land, we were able to obtain the rental rates paid per acre for land rented under both share and cash arrangements. It is likely that some frictions exist in lease arrangements for farm land, since these arrangements may not be negotiated every year. In this light, it may take some time for lease markets to respond to increases or decreases in the level of support provided to producers, in particular for cash leases. On the other hand, we should expect share lease payments to reflect the ex-post contribution of every single source of agricultural earnings. Share rents are indeed paid at the end of the season, once all uncertainty has been realized. Share lease payments are supposed to reflect the agreed proportion of cash flows from all sources of earnings related to the farming of the land, including government payments, though again share arrangements may be subject to the terms set through individual negotiations.

[17] The extent to which farm program payments should be limited was an important point of considerable debate in the 2002 deliberations. Any support based on production (such as LDP payments) will naturally favor larger producers. In the end, the only effective limit is that a "farmer" with more than $2.5 million per year in average adjusted gross income is excluded from the benefits *unless* at least 75% of the $2.5 million came from farm earnings.

We considered regressions of cash and share rents, respectively, against the factors expected to be relevant to land values and rents, including the indicators of expected payments. The results are presented in Table 4. Perhaps the strongest evidence of the effectiveness of landlords in extracting program benefits from their tenants is offered in the case of cash rents. Our analysis indicates that an additional dollar per acre of LDP payments raises the cash rental rate by $0.53, suggesting that landlords are quite effective in extracting a considerable share of the benefits. This occurs in spite of the legislative mandate that such payments be directed to the operator. In the case of AMTA payments, an additional dollar per acre of AMTA payments raises cash rents by $0.69 per acre. Thus, an even greater proportion of AMTA payments, which are known with certainty prior to the execution of any lease arrangement, goes to the landowners. Disaster payments tend to lower cash rental rates. An additional dollar of expected disaster payments lowers cash lease rates by $0.29. Likewise, CRP payments are associated with substantially lower cash lease rates. Again, this result is not surprising given use restrictions and lower quality of such land.[18] Indicators of urban pressures are generally not significant. As expected, higher levels of market returns are correlated with higher rental rates.

Table 4 also presents regression results for a model of share rents. The results are quite similar, though important differences exist. LDP payments appear to exert a smaller effect on share rental rates, with an additional dollar raising share lease rates by $0.24. AMTA payments have a significant effect on share lease rates, raising the rate by $0.78 for each dollar of payments. In contrast to the case of cash rents, disaster payments tend to raise share lease rates, by $0.72 per acre for each dollar of payments.

It is difficult to directly compare the effects of various factors on cash and share lease rates since the lease rates are drawn from different farms. However, such a comparison is possible for the subsample of farms (3,946 farms) that rented land both by cash and share rental arrangements. The difference in lease rates for such farms is the focus of our analysis in the following section.

[18]Unless such land is to be used for nonagricultural purposes, such as hunting, it is unlikely that it would be rented to a producer. Thus, CRP payments are likely a signal of low quality land.

4.3 Insurance Benefits

The typical approach to the assessment of the total contribution of agricultural policy to land values relies on the coefficient from the land value regressions. This is problematic for two reasons. The first one, usually mentioned in the literature, is due to the fact that regressions yield the effects of the marginal dollar for each type of policy. The literature has however overlooked the second reason. If we think about land as a portfolio of securities each delivering its stream of cash flow, it is obvious that the risk of the portfolio depends on the covariance of the various underlying securities. Therefore, eliminating one or more of the underlying securities will affect the risk of the portfolio. In terms of the analysis of the policy contribution, this implies that eliminating a policy which provides an insurance benefit will not only decrease expected returns, it will also increase the volatility of the remaining (market) returns. In other words, we should expect the coefficient on market earnings to decrease in response to an increase in uncertainty. The capitalization rate of earnings will be lower reflecting the higher opportunity cost of capital for an asset with more volatile earnings.

This raises the following question: if there is a theoretical argument in favor of an insurance component to the contribution of agricultural policy to land, can we find evidence from the market that it matters quantitatively? Unfortunately, there are no counties targeted by the ARMS survey that exempt all farmers from the benefits of agricultural policy. However, as we have noted above, farm land is rented under both cash lease and share lease arrangement. Cash lease rate are set ex-ante while the share payment depends upon the actual earning of the parcel, thus implying a risk sharing arrangement.

The main program designed to reduce the variability of farm earnings and insure the cash flow to farmers is the LDP program. Disaster payments are much more ad hoc and suffer from the uncertainties of the political process from year to year. If the insurance component matters, we should find that higher LDP payments should be correlated with a lower risk premium on rental arrangements. This risk premium is represented by the difference in cash and share lease rates on the subset of 3,946 farms that report renting land under both types of arrangements. By committing to an ex-ante fixed payment, the farmer provides insurance to his landlord for which we should expect him to be rewarded (unless we observe cash rents only when the farmer is not risk averse).

To evaluate this risk premium, we regressed the share-cash rental premium on policy benefits and other factors suspected to be relevant to rental rates. These results are also presented in Table 4. We find that LDP payments do indeed tend to significantly decrease the share-cash premium. An additional dollar in expected LDP payments lowers this premium by \$0.61. In contrast, ad hoc disaster payments tend to raise the premium. This is certainly to be expected since disaster payments naturally are targeted toward areas more likely to suffer a disaster (a crop failure)—in other words, riskier areas. With the exception of market returns, which have a modest positive influence on the share-cash differential, the other factors do not appear to significantly affect the risk premium implied by the rental rate differences.

5 Concluding Remarks

Policy rhetoric often justifies Farm Bill expenditures with the argument that impoverished farmers are in need of governmental support to remain in business. This view is pervasive outside of Washington. For example, consider the annual "Farm Aid" events intended to draw attention to the plight of the American farmer. Our analysis challenges this view. We demonstrate that land owners capture substantial benefits from agricultural policy. This is particularly problematic given that in many cases land owners are different from the farmers whose plight we are told we should be concerned with.

Of course, many farmers are also landowners and thus have an important stake in maintaining agricultural policy benefits. A farmer that purchased land which reflected the value of anticipated benefits would certainly suffer a damaging capital loss if such support were to be withdrawn. Furthermore, all farmers have a strong interest in congressional surprises whereby more transfers are allocated than anticipated by the land market. As owners they benefit from the unexpected capital gains. The 2002 Farm Bill with its large increase in support expenditures may have been one such nice surprise.

Tenants also gain from positive surprises as long as lease rates do not adjust instantaneously. However, the 2002 Farm Bill seems to have shut down this avenue for a temporary increase in the share of transfers captured by farm operators. One valuable provision of the bill is that it offers to farmers the opportunity to update the factors which determine the level of some of the payments they receive. The power to decide whether or not to update has been given to the owners of the land, not the operator. Not surprisingly, tenant farmers are complaining that land owners are using this opportunity to impose a renegotiation of the existing leases that did not foresee the generosity of the 2002 Farm Bill.

References

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

Deaton, A. S. *The Analysis of Household Surveys: A Microeconometric Approach to Development Policy* Baltimore: Johns Hopkins University Press and the World Bank, 1997.

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

Goodwin, B. K., A. K. Mishra, and F. Ortalo-Magné. "Explaining Regional Differences in the Capitalization of Policy Benefits into Agricultural Land Values," Chapter in *Government Policy and Farmland Markets: Implications of the New Economy,* Eds. C. Moss and A. Schmitz, Ames, IA: Iowa State University Press, in press.

Hardie, I. W., T. A. Naryan and B. L. Gardner. "The Joint Influence of Agricultural and Nonfarm Factors or Real Estate Values: An Application to the Mid-Atlantic Region," *American Journal of Agricultural Economics* 83(2001):120-132.

Mishra, A. K., H. S. El-Osta, M. J. Morehart, J. D. Johnson, and J. W. Hopkins. "Income, Wealth, and the Economic Well-Being of Farm Households," USDA, Economic Research Service, Agricultural Economic Report No. 812, July 2002.

Ryan, J., C. Barnard, R. Collender, and K. Erickson. "Government Payments to Farmers Contribute for Rising Land Values," *Agricultural Outlook* June-July, 2001, pp. 22-26.

Shertz, L. and W. Johnston. "Managing Farm Resources in the Era of the 1996 Farm Bill," Staff Paper No. AGES 9711, December 1997, Washington, USDA-ERS.

Shoemaker, R., M. Anderson, and J. Hrubovcak. "U.S. Farm Programs and Agricultural Resources," Agriculture Information Bulletin No. 1990, September 1990, Washington, USDA-ERS.

Weersink, A., S. Clark, C. G. Turvey, and R. Sarker. "The Effect of Agricultural Policy on Farmland Values," *Land Economics* 75(August 1999):425-39.

Table 1. Variable Definitions and Summary Statistics

Variable	Definition	Mean	Std. Dev.
Value	$/acre reported value	1435.4700	11661.8500
Total Payments	Total program payments	35.5857	275.0026
Mean Total Payments	Average total payments over preceding 4 years	37.5198	305.4024
LDP	Loan deficiency payment receipts ($/acre)	13.4265	169.0008
Disaster	Disaster payment receipts ($/acre)	6.1012	102.7038
CRP	CRP payment receipts ($/acre)	1.2710	45.2617
AMTA	AMTA payment receipts ($/acre)	12.5311	133.1091
Mean LDP	Average LDP payments over preceding 2 years	15.0362	200.6129
Mean Disaster	Average Disaster payments over preceding 4 years	8.0834	90.0668
Population	County average persons per square mile	80.4959	1429.2200
Population Growth Rate	Population growth rate (proportion)	0.3458	21.7087
Urban $_2$	1 if Urban category 2, 0 otherwise	0.1113	3.0641
Urban $_3$	1 if Urban category 3, 0 otherwise	0.1136	3.0908
Urban $_{45}$	1 if Urban category 4 or 5, 0 otherwise	0.1351	3.3302
Housing Permits	Total value of housing permits ($ten-million)	3.3071	106.0131
Market Return	Average return using market prices	20.5896	455.2781
Cash Rental Rate	Cash rental rate ($/acre)	60.9207	520.3213
Share Rental Rate	Share rental rate ($/acre)	90.8526	1158.5400
Share-Cash Difference	Share-cash rental rate difference ($/acre)	34.8825	1293.3400
U.S. Average, 1998			
AMTA	$/farm for farms receiving	3626.09	5580.36
Disaster	$/farm for farms receiving	1596.03	5409.57
Market Loss Assistance	$/farm for farms receiving	1765.33	3072.35
LDP	$/farm for farms receiving	3065.49	6076.42
U.S. Average, 1999			
AMTA	$/farm for farms receiving	3373.57	5941.03
Disaster	$/farm for farms receiving	6631.56	6083.59
Market Loss Assistance	$/farm for farms receiving	3478.80	5773.27
LDP	$/farm for farms receiving	5805.11	9687.33
U.S. Average, 2000			
AMTA	$/farm for farms receiving	3259.65	5589.97
Disaster	$/farm for farms receiving	4453.46	4783.49
Market Loss Assistance	$/farm for farms receiving	4028.95	5965.30
LDP	$/farm for farms receiving	6109.53	9559.20
U.S. Average, 2001			
AMTA	$/farm for farms receiving	2690.20	4733.47
Disaster	$/farm for farms receiving	6505.10	7604.74
Market Loss Assistance	$/farm for farms receiving	3255.48	5299.70
LDP	$/farm for farms receiving	6170.27	6939.24

Table 2. Aggregated Policy Models of Land Value Determinants:

Parameter Estimates and Summary Statistics

Variable	Model I	Model II
Intercept	808.3932	775.6243
	(19.0329)*	(15.6736)*
Total Payments	5.4036	
	(0.4937)*	
Mean Total Payments		6.0902
		(0.3812)*
Population	1.8571	1.8426
	(0.2347)*	(0.2367)*
Population Growth	64.7655	57.7059
	(4.9556)*	(4.9894)*
$Urban_2$	312.1999	304.9882
	(30.4229)*	(29.8719)*
$Urban_3$	498.5101	516.4412
	(41.0672)*	(40.7747)*
$Urban_{45}$	303.8665	290.7366
	(51.2658)*	(50.6040)*
Housing Permits	4.7769	5.0110
	(2.6694)*	(3.0486)*
Market Return	5.5625	5.5633
	(0.2474)*	(0.2376)*
Number of Observations	13,603	13,603
R^2	0.2566	0.2653

[a] Numbers in parentheses indicate statistical significance at the $\alpha = .10$ or smaller level.

Table 3. Disaggregate Policy Models of Land Value Determinants

(U.S. and Heartland Regions): Parameter Estimates and Summary Statistics

Variable	Model III	Model IV	Model V
Intercept	839.2467	869.5799	1083.1000
	(16.6984)*	(15.7413)*	(29.9022)*
LDP	7.0195		
	(0.6910)*		
Mean LDP		9.0471	7.1979
		(0.7917)*	(1.3902)*
Disaster	7.1600		
	(1.1493)*		
Mean Disaster		−4.7531	−6.5483
		(1.8480)*	(3.8987)*
CRP	−15.8161	−17.1994	−19.4049
	(1.6775)*	(1.6522)*	(1.8886)*
AMTA	3.8046	4.4000	5.5105
	(0.6814)*	(0.6930)*	(1.0587)*
Population	1.8357	1.8283	0.8645
	(0.2325)*	(0.2328)*	(0.2968)*
Population Growth	64.3271	62.0081	21.1991
	(4.9389)*	(4.9316)*	(9.8773)*
$Urban_2$	307.2501	309.5498	198.4065
	(30.1725)*	(29.7035)*	(39.8393)*
$Urban_3$	494.0376	516.2642	494.4873
	(40.7594)*	(40.6575)*	(59.8004)*
$Urban_{45}$	287.5609	284.5653	453.6068
	(51.2871)*	(50.7983)*	(64.5146)*
Housing Permits	4.8583	5.6337	2.8406
	(2.5963)*	(2.8092)*	(5.2764)
Market Return	5.6028	5.6395	6.3319
	(0.2417)*	(0.2421)*	(0.4325)*
Number of Observations	13,603	13,603	4,599
R^2	0.2645	0.2656	0.1990

[a] Numbers in parentheses indicate statistical significance at the $\alpha = .10$ or smaller level. Models III and IV correspond to entire U.S. while Model IV corresponds to the Heartland region.

Table 4. Models of Land Rental Rate Determinants:

Parameter Estimates and Summary Statistics

Variable	Cash Rents	Share Rents	Difference
Intercept	38.3035	54.7436	18.5531
	(1.0445)*	(2.7395)*	(4.0126)*
Mean LDP	0.5261	0.2448	−0.6117
	(0.0380)*	(0.1484)*	(0.2343)*
Mean Disaster	−0.2917	0.7167	1.5311
	(0.0953)*	(0.3844)*	(0.6120)*
CRP	−0.5197	−0.8692	−0.7892
	(0.1208)*	(0.3121)*	(0.3647)*
AMTA	0.6877	0.7806	0.2520
	(0.0597)*	(0.1510)*	(0.2610)
Population	−0.0017	0.1075	0.1162
	(0.0037)	(0.0574)*	(0.0724)
Population Growth	1.0475	0.6651	0.5714
	(0.3403)*	(0.6369)	(0.9533)
$Urban_2$	1.4754	14.3245	13.1163
	(1.7789)	(8.1201)*	(11.7120)
$Urban_3$	−0.4555	8.6911	11.9550
	(1.5206)	(6.8805)	(10.1192)
$Urban_{45}$	1.6098	−10.9351	−13.7343
	(1.8301)	(6.8567)	(9.8205)
Housing Permits	−0.1523	−0.3019	−0.4106
	(0.0697)*	(0.5433)	(0.7564)
Market Return	0.4076	0.5745	0.1467
	(0.0151)*	(0.0279)*	(0.0405)*
Number of Observations	9,605	6,007	3,946
R^2	0.2135	0.0928	0.0165

[a] Numbers in parentheses indicate statistical significance at the $\alpha = .10$ or smaller level.