Staff Paper 04-9 November 2004

Staff Paper Series

Evidence of a market effect from conservation easements

Steven J. Taff

DEPARTMENT OF APPLIED ECONOMICS
COLLEGE OF AGRICULTURAL, FOOD, AND ENVIRONMENTAL SCIENCES
UNIVERSITY OF MINNESOTA

Staff Paper 04-9 November 2004

Evidence of a market effect from conservation easements

Steven J. Taff

Steven J. Taff is an associate professor and extension economist in the Department of Applied Economics, University of Minnesota. The analyses and views reported in this paper are those of the author. They are not necessarily endorsed by the Department of Applied Economics or by the University of Minnesota.

This research was funded by the U.S. Department of Agriculture and the Minnesota Agricultural Experiment Station. Thanks to members of the Minnesota Association of Assessing Officers and to Department of Revenue Property Tax Division staff for preparing the dataset analyzed in this report.

Copies of this publication are available at http://agecon.lib.umn.edu/. Information on other titles in this series may be obtained from: Waite Library, University of Minnesota, Department of Applied Economics, 232 Classroom Office Building, 1994 Buford Avenue, St. Paul, MN 55108, U.S.A.

Abstract

Both federal income tax law and local property tax assessment practice are premised on the "theory" that permanent conservation restrictions reduce property values, while shorter term restrictions do not. An analysis of 190 recent sales of Minnesota agricultural land with varying cropping rights restrictions supports only a portion of this theory. Both permanent and short-term restrictions were negatively and significantly associated with per-acre sales prices. The former effect is consistent with the theory and is statistically meaningful, but the dataset contained too few sales with permanent restrictions to warrant a strong conclusion. The latter effect is inconsistent with the theory, but is strongly supported by the data..

Introduction

Easements have been an increasingly popular conservation tool over the past few decades. They are presumed to be both less expensive to acquire and less intrusive into the workings of a private property system, while at the same time accomplishing many if not all of the objectives achieved by full public ownership. There are two major types of conservation easement: those for cropping rights and those for development rights. The former might be for either a short duration or permanent. Strictly speaking, shorter-term restrictions are not "easements" in the conventional sense, so, to avoid confusion in this paper, I follow Viaggi and Taff (2004) and refer to all use-restrictions, whatever their duration, as "conservation contracts" or "conservation restrictions."

I examine the ways in which cropping rights restrictions influence farmland markets. Why do we care about the market effects? Because these effects (if any) reverberate through both our income and property tax systems, as well as indirectly into the administration of easement programs. Property tax assessors want to know whether or not to change the estimated market value of properties on which a restriction has been placed. Appraisers are asked to value the restriction itself, to be used as the donation value for any income tax deduction. (None of the contracts examined in this report were donated, but appraisers are known to use purchased conservation contract properties as comparable sales.) Administrators want to know how much contracts will cost. Presumably, they will have to pay at least as much as the property value goes down (if at all), in order to make the contract seller financially whole.

Shultz and Taff (2004) found significant value reductions associated with USFWS wetland easements in North Dakota. Here, I look at the market effects of two programs that are prominent in Minnesota: the federal Conservation Reserve Program (CRP), which pays participating landowners a fixed annual amount for ten years for not growing crops; and the state's Reinvest in Minnesota (RIM) easement program which pays landowners up-front for permanent conservation use of cropland. (A blended CRP/RIM program, called the Conservation Reserve Enhancement Program (CREP), is too recent and too narrowly targeted for there to be enough sales in the study area to warrant treatment here.) The conservation restrictions under each are essentially the same: enrolled land cannot be cropped and must be put into some sort of persistent cover crop such as native grasses or trees.

The "theory" of conservation restrictions and property markets

The story we tell about land values is that the buyer calculates how much to pay in order to secure rights to a stream of future income. For agricultural land, the convention is to calculate the expected annual net return (profits) and discount future payments by some amount to take account of uncertainty, risk, capital costs, etc. This "income approach" to land valuation is commonly used by appraisers, but it cannot be used by tax assessors (in Minnesota, at least), who are restricted to look only to comparable sales, using a "market approach." We usually assume, however, that buyers and sellers employ at least an implicit income approach to their personal valuations of a property.

A property restricted by a conservation contract ought to sell for less than an otherwise identical but unrestricted property, because the restriction precludes certain future uses that might have economic value greater than those that remain with the property. Such an assumption underlies the whole notion of paying for the contract (or of allowing an income tax deduction in the case of a donated easement) in the first place. If the restriction does not bind, if it does not have any future (negative) economic consequences, how can we justify paying for it in the first place?

The expected income stream of unrestricted properties is a perpetual series of annual net profits from agricultural production. (All the contracts examined here are required to be on cropland, and agriculture is presumed to be the highest and best use (HBU) because that's what the unrestricted properties are in fact being used for.)

A conservation contract restricts these future uses to some use that is presumably of lesser economic magnitude, because it is not the HBU for the property. The seller of RIM-restricted land has already received the one-time payment for the contract, so the buyer secures only the rights to the future lower income stream. If our conventional asset pricing theory holds, the buyer will want to pay less for a RIM-restricted property than for an unrestricted property of similar location and agronomic characteristics.

CRP contracts are a little more complicated. The buyer of a CRP-restricted property secures rights to the remaining annual CRP payments, plus, when the contract expires, a thenceforth perpetual stream of unrestricted-use income, minus any costs of converting from the conservation use to the new (presumably agricultural) use. Because the seller of the property was also the seller of the CRP contract to the government, we assume that the net returns of the annual CRP payments, minus entry and exit conversion costs, exceeded the unrestricted-use income over the contract period; otherwise, the seller would not have entered the CRP. So, conventionally, we assume that the income stream of the buyer of a CRP-restricted property will result in a purchase price that is not lower than that for a non-restricted property.

This "theory" is built into law in several ways. The theory is rarely tested, however, because we usually lack a sufficient number of property sales to statistically examine market effects. But sometimes we get lucky.

Important for our purposes is that CRP contracts are not considered property transfers under US tax law. Annual payments are taxed as ordinary income, just like cash rent on these formerly agricultural properties. And Minnesota property tax assessors are instructed not to reduce the assessed value of CRP-restricted properties, under the premise that the conservation restrictions are short term only and that the underlying market value of the property is not affected by short-term contracts. Conversely, the perpetual RIM contracts are treated as property transfers, the proceeds from the contract sale are treated as would be proceeds from a full-title transaction, and Minnesota assessors are instructed to adjust the assessed market value accordingly. A new dataset may permit us to determine how big the conservation adjustment should be. Let's look at the markets.

Data

In many hedonic price studies, we use market data to infer nonmarket values such as water quality. In those studies, actions of real estate buyers and sellers form the basis of inferences about the population as a whole, not just about market participants. But here we're interested only in how "the market" works, so we use observed sales (and, implicitly, the preferences of property buyers and sellers) as a base from which to infer characteristics of all properties in the market, including—especially—those that did not sell.

Sales data in general are only "opportunistic samples" in that we can't be sure they fully represent the unobserved population of all sales in a geographic area. Only 1-2% of Minnesota's farmland sells in any given year, so the "sample" we have here must be used with caution. In Taff (2004) I discuss the implications of this sampling problem.

Until 2003, official farmland sales records in Minnesota did not reveal whether or not the property was encumbered by a conservation restriction. Indeed, in many instances, such sales, if known at all, were discarded from the "arms-length" sales data set that is used for property tax equalization judicial hearings and that are reported through Minnesota Land Economics web site (*http://www.apec.umn.edu/landeconomics*). For this study, however, county assessors in three assessment regions went back to 2000-2003 files and pulled sales for which they had record of either a CRP or a RIM conservation contract. (There were only 10 sales with CREP contracts; these were excluded from further analysis.) These cannot be considered all the sales that were encumbered in those regions, but they are thought to constitute a substantial portion of that unknown population.

These sales files were subsequently processed by Department of Revenue experts to adjust the prices for terms and time of sales, so that they could be properly contrasted to prices used in the regular sales data series. Sales with per-acre prices of over $10,000 were discarded, as were sales that had more than one type of contract. This resulted in a conservation contract dataset of 190 sales, located in the Central, South West, and West Central regions shown in Figure 1. Not all counties in the regions reported contracts. These contract sales were then paired with a stratified (by location and year) sample of the same size from the non-eased dataset reported in Minnesota Land Economics and analyzed in Taff (2004).

Figures 2-4 show the distribution of the sales by year, location, and conservation contract type. Figure 5 shows the proportion of each property that is restricted by the conservation contract. It's clear that it is insufficient to speak simply of restricted and unrestricted properties: properties restricted by contracts may still retain a high proportion of unrestricted land.

Figure 1: Location of study area

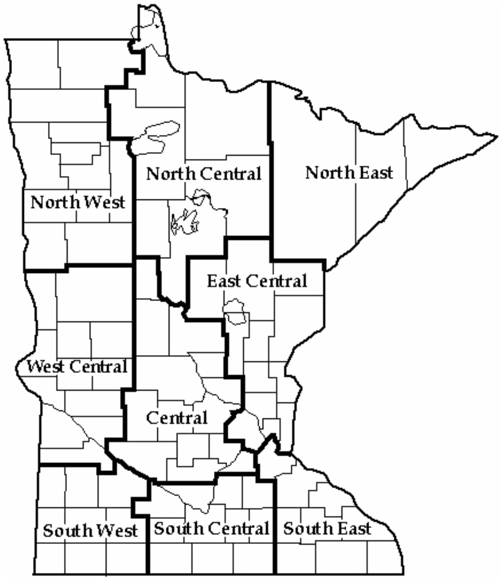

Figure 2: Type of conservation restriction by year of sale

	2000	2001	2002	2003	Total
None	19	60	64	47	190
CRP	13	50	52	41	156
RIM	6	10	12	6	34
Total	38	120	128	94	380

Figure 3: Location of sales by year of sale

	2000	2001	2002	2003	Total
Central	12	44	45	33	134
South West	13	34	48	29	124
West Central	13	42	35	32	122
Total	38	120	128	94	380

Figure 4: Type of restriction by location of sale

	Central	South West	West Central	Total
None	103	38	49	190
CRP	22	74	60	156
RIM	9	12	13	34
Total	134	124	122	380

Figure 5: Proportion of property subject to conservation restriction

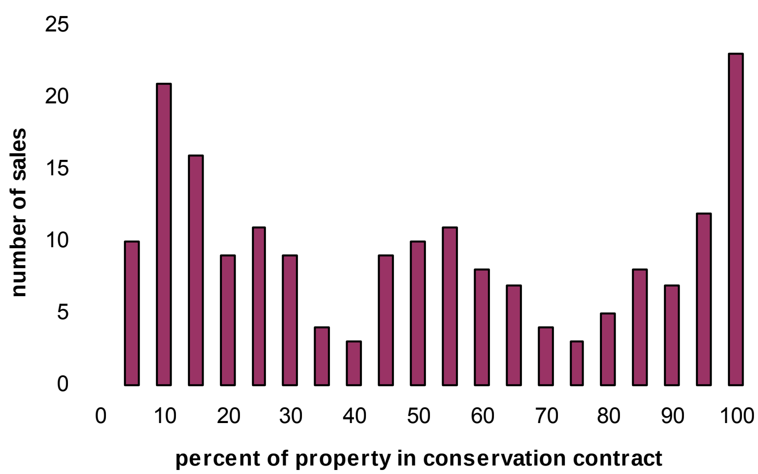

Analysis

The combined 380 sale dataset is the basis for the analysis that follows. What we want to know is simple: to what extent do conservation restrictions affect property sales prices? These sales are examined with three analytic tools—scatter plots, box plots, and regression models—that permit us to speak of certain aggregate and average relationships among them. Any market influence from conservation contracts that we might find here are averages, adjusted for various institutional, location, or physical characteristics of the examined properties. We then assert that: (1) the sales we examined are "representative" of the market as a whole, so that (2) we can use any measured influences as a best estimate of influences that will be seen in transactions yet to be made.

Scatter plots

First, let's get a sense of how restricted properties compare to unrestricted properties with respect to the size of property and the productivity of the land. We know that each of these variables influences farmland sales prices, because we've been tracking these relationships for many years (Taff, 2004). In these charts (Figures 6-7), I've marked each sale by whether or not it is encumbered by a conservation restriction of any type. Neither figure seems to show a strong difference between restricted and unrestricted properties: there are black dots pretty much everyplace that there are white dots, although the mass of the contract sales in Figure 6 might lie lower in the price range than does the mass of the non-contract sales. There are fewer sales plotted in Figure 7 because productivity is known for only some of the properties (Figure 15). Productivity is measured by the University of Minnesota's crop equivalent rating (CER), which ranges from 0 to 100.

Figure 6: Per-acre sales prices and size of parcel, marked by presence or absence of conservation restriction

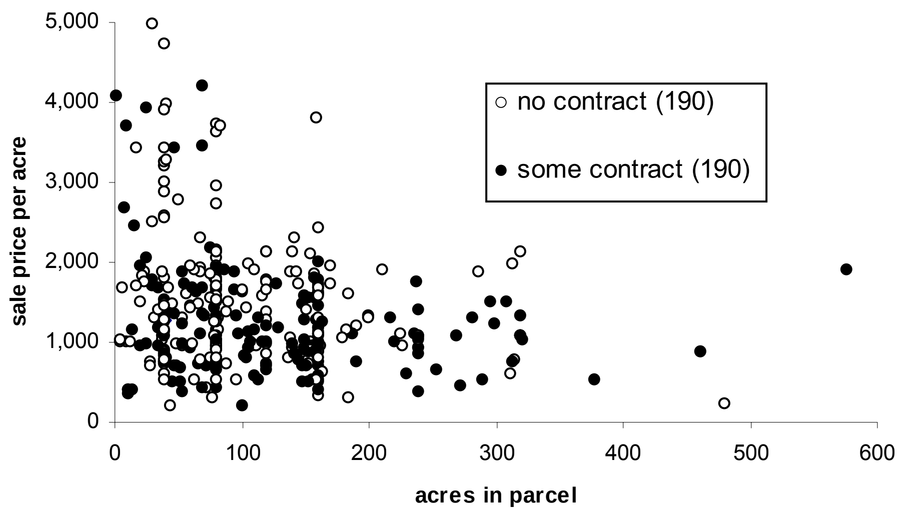

Figure 7: Per-acre sales prices and land productivity, marked by presence or absence of conservation restriction

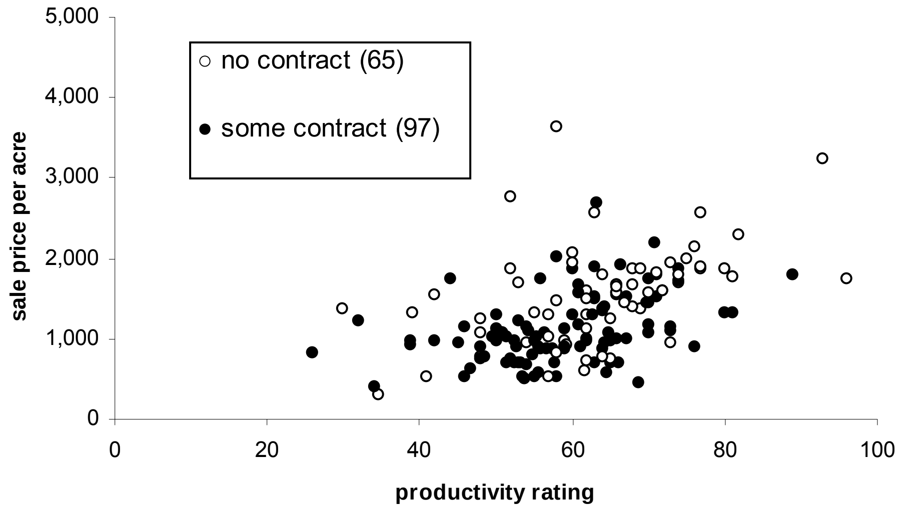

Box plots

If there is an influence of conservation restrictions on sales prices, it didn't jump out in the scatter plots. Let's look at the relationships in a different way. Figures 8-13 are box-and-whisker plots that show the separate distributions of per-acre prices, separated according to certain features. (For clarity, a few high price sales are excluded from these plots—but not from the regression models that follow.)

Each box enfolds the middle half of the sales price distribution, with the median price shown by a horizontal bar. The dotted lines and the markers span most of the rest of the price range. So, for example, Figure 8 shows that half the Central region sales were between $1,000 and $2,500 per acre, with a median sales price of about $1,500. Prices in the other two regions were, on the whole, lower, although there still were a few above $2,00 per acre in each region.

Figure 8: Distribution of per-acre prices by location of sale

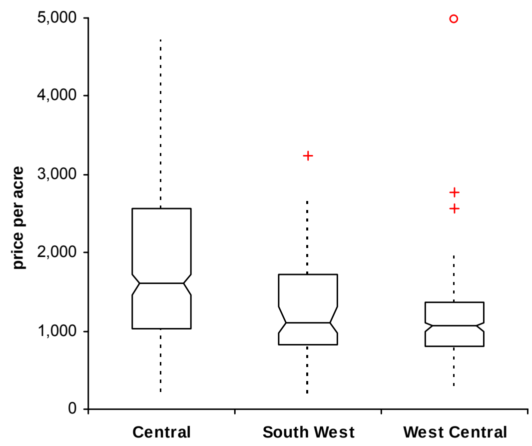

Figure 9: Distribution of per-acre prices by presence or absence of a contract

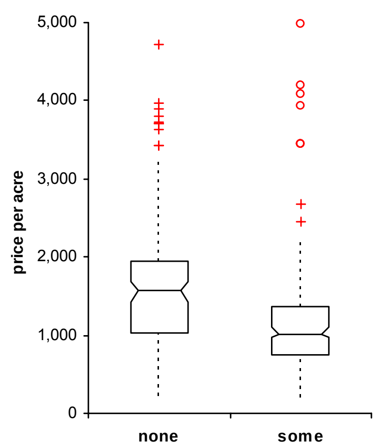

Figure 10: Distribution of per-acre prices by type of contract

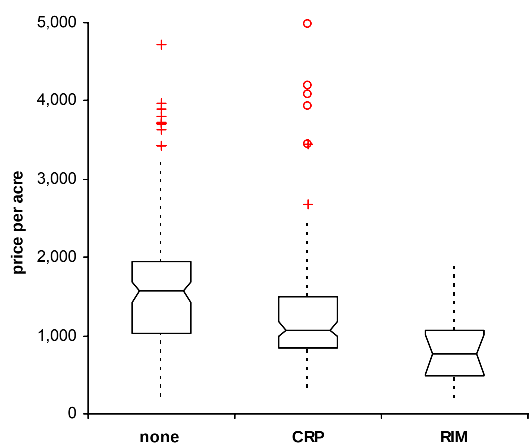

Both Figures 9 and 10 suggest that conservation contracts might make a difference: the boxes, which span half of the sales in each category, look like they're lower for conservation contracts (compared to no contracts) and for CRP and RIM contracts (compared to no contracts).

In Figures 11-14, we show how contract types differ (or don't differ) with respect to property size, cropland productivity, and tillable land proportion. (Recall that the number of properties for which we know the productivity is considerably lower than for other property characteristics. See below for a discussion of the implications.) There is an especially striking difference between contract types with respect to the proportion of the property that is tillable (Figure 13) and also the percent of the property that is restricted by contract (Figure 14).

Figure 11: Distribution of property size by type of contract

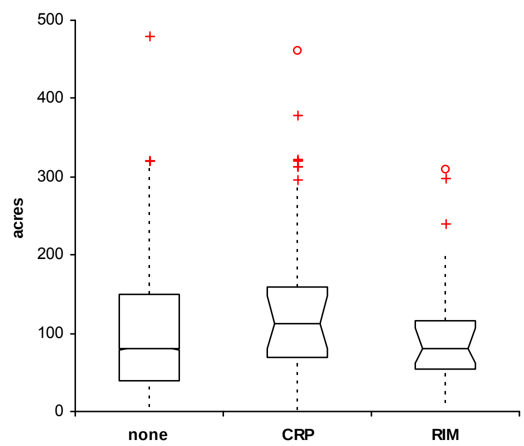

Figure 12: Distribution of cropland productivity by type of contract

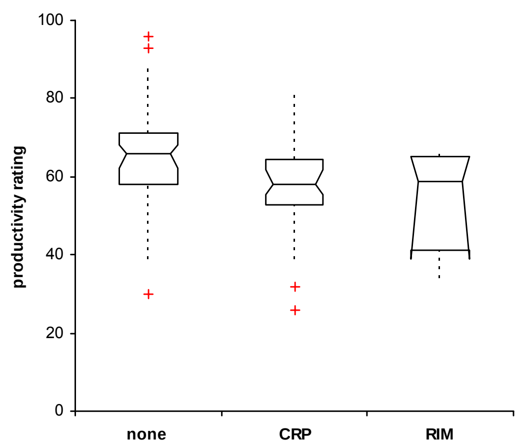

Figure 13: Distribution of tillable proportion by type of contract

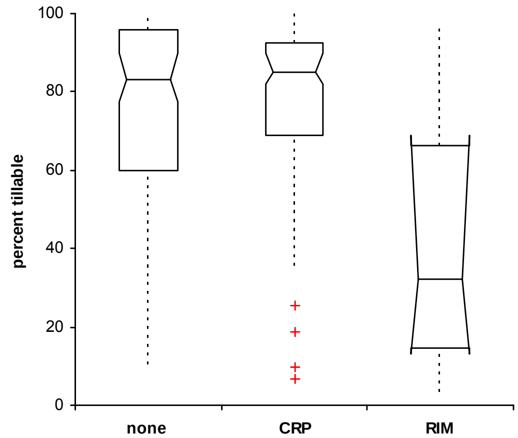

Figure 14: Distribution of contract proportion by type of contract

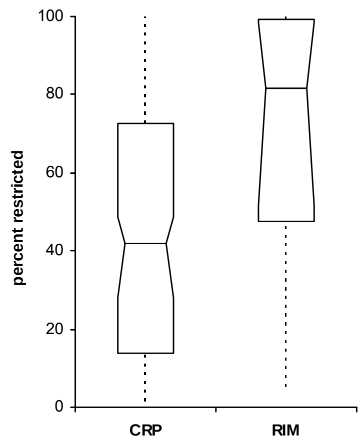

Figures 13 and 14 suggest that we need to be careful about how we add information about tillable land and restricted land to the model. Because each appears to be strongly correlated with contract type (long-term contracts show lower tillable and higher restriction proportions), adding them to a model that already has contract type information can mislead us: we won't be sure what the estimated coefficients are telling us. Tillable land is not so much a characteristic of the land itself as it is a measure of the legal uses for the land. Fields with RIM contracts are removed from the tillable land category by the local assessor when the contract is signed, because these field can no longer be used for crops, whatever their inherent agronomic capacity.

Adding information about productivity, on the other hand, appears to be less likely to confuse us. The addition of productivity information does come at a cost, however. We do not have these data for all the sales in the dataset (Figure 15). When we add it, we reduce the number of sales by which we can calibrate the model, from 380 to 162, and the number of sales with conservation contracts from 190 to 97. Note especially that the following regression models are calibrated with the smaller set of data. This means that we include information on only 12 properties with RIM contracts. While this can still lead to statistically reliable results, there is strong reason to suspect that our findings with respect to RIM contracts may be "fragile." I'll discuss later the effect of using instead the full sat set, without the productivity information.

Figure 15: Effect of limiting sales to those with productivity data (number of sales)

	Full	CER
None	190	65
CRP	156	85
RIM	34	12
Total	**190**	**162**

Regression models

I now turn to a more formal modeling of the price relationships hinted at in the charts. My approach is to start with a base model that makes use of available property characteristics to try to explain as much of the observed variation in sales prices as possible. Then I'll add information about conservation contracts to this model. If the additional information "improves" the model in a particular way, we will say that contracts "matter," that they influence sales prices. (More precisely, my approach is to statistically test whether or not we can reject the hypothesis that the conservation contract information *doesn't* matter.)

Here are the variables—in addition to conservation contract information—that we can use. Almost certainly there exist non-observed variables that, if known, should be included as well. The statistical procedures used here help us test for the possible influence of these unknown variables and to adjust the model to take into effect our not including them—to an extent. The key variables are:

Price: per-acre time- and terms-adjusted sale price per acre

Bareland factor (dummy), compared against property with improvements

Size of property, in number of acres

Year in which property sold factors, compared to 2000

Central and *West Central* region factors, compared to Southwest region

CER: average agronomic productivity (scaled 1-100)

Proportion, percent of property under conservation contract

CRP or RIM contract factor (dummy), compared to no contract

I start with a base model that regresses per-acre sales price against non-contract variables (Model 1). Both the price and size variables are logged to improve the linearity of the model. This initial base model has an R2 of .44, and most of the other variables are "significant" in the sense that each has a low p-value. Loosely, a low p-value means that the estimated coefficient probably does not have a value of 0. In this sense, a low p-value gives us some confidence—but not complete assurance—that the variables that matter in the model probably also matter in the world. (Examination of the statistical properties of this model show that both the OLS linearity and constant variance assumptions hold. Details available from the author.)

Model 1 is our base model, upon which I will now add information about contracts. If additional information about conservation contracts have "low" individual p-values, we'll say they matter, and we'll assert that the association is causal: that, all else equal, conservation contracts influence sales prices in the stated manner. In these explorations of the data, I'll not hold to a too-rigid threshold for statistical "significance." A p-value of 0.05 is the level below which we say a variable matters.

Whenever many models are calibrated against the same dataset, we can expect to see variation in their conclusions. Sometimes, chance alone can account for an observation that, say, a variable is significant in one model but not in several other models. In this study, I'm interested in the preponderance of the evidence, not in any particular model. If we see the same variable behaving in the same way in each of several models, we can start taking it seriously. Later, I'll combine the conclusions drawn from each of the several models, along with conclusion already drawn from examination of the scatter plots and box plots.

Model 1: Base model, not adjusted for conservation contracts

```
Response      = log[price]
Coefficient Estimates
Label                        Estimate       Std. Error    t-value    p-value
Constant                      6.44706       0.280264       23.004     0.0000
bareland                     -0.324954      0.0780720      -4.162     0.0001
log[size]                    -0.140277      0.0429920      -3.263     0.0014
{F}year[2001]                 0.110508      0.111480        0.991     0.3231
{F}year[2002]                 0.280442      0.109050        2.572     0.0111
{F}year[2003]                 0.387022      0.111027        3.486     0.0006
{F}[Central]                  0.216939      0.104229        2.081     0.0391
{F}[West Central]             0.128758      0.0620568       2.075     0.0397
cer                           0.0205576     0.00252559      8.140     0.0000

R Squared:                 0.443062
Sigma hat:                 0.350326
Number of cases:               380
Number of cases used:          162
Degrees of freedom:            153

Summary Analysis of Variance Table
Source         df       SS            MS           F    p-value
Regression      8   14.9381       1.86726      15.21    0.0000
Residual      153   18.7775      0.122729
 Lack of fit  151   18.7766      0.124348     276.29    0.0036
 Pure Error     2  0.000900113   0.000450056
```

Figure 9 suggests that the mere presence of a conservation restriction is associated with a lower-lying distribution of per-acre sales prices. Let's look at this more formally. Model 2 takes our base model and adds a variable called *contract*, which takes the value 1 if there is any sort of conservation restriction of any size on the property, and 0 otherwise.

The estimated coefficient of the *contract* variable can be interpreted as the average proportional change in price associated with the imposition of conservation restrictions. Addition of this variable increases the model R2 slightly, and the variable itself has a p-value of .0015. Simply knowing whether or not a sale has a conservation contract helps us better explain observed variation in (logged) per-acre sales prices: a contract lowers average price by 20%. (The estimated coefficient of a 0/1 variable such as *contract* is not precisely the proportional change in price, but the two are close enough for their difference to be erased by the rounding employed here.)

Model 2: Adjusting for contract presence

```
Response      = log[price]
Coefficient Estimates
Label                        Estimate        Std. Error    t-value    p-value
Constant                      6.64597        0.278872       23.832     0.0000
bareland                     -0.282287       0.0769101      -3.670     0.0003
log[size]                    -0.115439       0.0424254      -2.721     0.0073
{F}year[2001]                 0.102869       0.108217        0.951     0.3433
{F}year[2002]                 0.255762       0.106108        2.410     0.0171
{F}year[2003]                 0.364948       0.107968        3.380     0.0009
{F}[Central]                  0.146687       0.103463        1.418     0.1583
{F}[West Central]             0.0728041      0.0626650       1.162     0.2471
cer                           0.0174721      0.00263045      6.642     0.0000
contract                     -0.200696       0.0621006      -3.232     0.0015

R Squared:                  0.47887
Sigma hat:                  0.33999
Number of cases:               380
Number of cases used:          162
Degrees of freedom:            152

Summary Analysis of Variance Table
Source         df       SS            MS           F    p-value
Regression      9   16.1454       1.79393      15.52    0.0000
Residual      152   17.5702      0.115593
```

Model 3 looks at the influence of contract type, evaluated against the no-contract condition. For example, the estimated coefficient for the variable *{F}CRP* is the proportional effect on price of just the CRP contracts of any size, compared to sales with no contracts. Here, too, contracts matter: the coefficients for both contract types are negative with p-values below the threshold. CRP contracts on average reduce per-acre sales prices by 19%, while RIM contracts reduce prices by 32%.

Model 3: Adjusting for contract type

Response = log[price]

Coefficient Estimates

Label	Estimate	Std. Error	t-value	p-value
Constant	6.67903	0.279859	23.866	0.0000
bareland	-0.270574	0.0774281	-3.495	0.0006
log[size]	-0.122008	0.0427227	-2.856	0.0049
{F}year[2001]	0.116153	0.108638	1.069	0.2867
{F}year[2002]	0.269476	0.106582	2.528	0.0125
{F}year[2003]	0.371468	0.107956	3.441	0.0007
{F}[Central]	0.159800	0.103903	1.538	0.1261
{F}[West Central]	0.0724284	0.0625787	1.157	0.2489
cer	0.0170894	0.00264631	6.458	0.0000
{F}CRP	-0.185884	0.0632449	-2.939	0.0038
{F}RIM	-0.315452	0.114446	-2.756	0.0066

R Squared:	0.483737
Sigma hat:	0.339517
Number of cases:	380
Number of cases used:	162
Degrees of freedom:	151

Summary Analysis of Variance Table

Source	df	SS	MS	F	p-value
Regression	10	16.3094	1.63094	14.15	0.0000
Residual	151	17.4061	0.115272		

From Figure 5, we know that the sales show a wide range of contract coverage. In many cases, only a small portion of the property is restricted. In Model 4, I examine this by adding the ratio of the number of restricted acres to the total number of acres, multiplied by 100, a variable I call *proportion*. It, too, shows a negative coefficient and a low p-value. Each percent increase in restricted acres, all else equal, reduces average price by 0.35%.

Model 4: Adjusting for contract proportion

Response = log[price]

Coefficient Estimates

Label	Estimate	Std. Error	t-value	p-value
Constant	6.81918	0.291342	23.406	0.0000
bareland	-0.270490	0.0770455	-3.511	0.0006
log[size]	-0.149309	0.0416127	-3.588	0.0004
{F}year[2001]	0.106090	0.107698	0.985	0.3262
{F}year[2002]	0.269029	0.105395	2.553	0.0117
{F}year[2003]	0.367283	0.107405	3.420	0.0008
{F}[Central]	0.199232	0.100817	1.976	0.0499
{F}[West Central]	0.0788314	0.0616618	1.278	0.2030
cer	0.0161505	0.00275259	5.867	0.0000
proportion	-0.00350698	0.00101422	-3.458	0.0007

R Squared:	0.483677
Sigma hat:	0.338418
Number of cases:	380
Number of cases used:	162
Degrees of freedom:	152

Summary Analysis of Variance Table

Source	df	SS	MS	F	p-value
Regression	9	16.3074	1.81193	15.82	0.0000
Residual	152	17.4081	0.114527		

Model 5 combines the type of contract and the proportion of the property that is restricted. This permits us to examine the separate influences of contract type. The CRP contracts, on average, reduce sales prices by 0.31% for every percentage point increase in the property under CRP restrictions. RIM contracts reduce sales prices by 0.52% for each percentage point increase in RIM coverage.

Model 5: Adjusting for contract type and proportion

```
Response      = log[price]
Coefficient Estimates
Label                        Estimate       Std. Error    t-value    p-value
Constant                      6.87095       0.294327       23.345     0.0000
bareland                     -0.262312      0.0772679      -3.395     0.0009
log[size]                    -0.155696      0.0419184      -3.714     0.0003
{F}year[2001]                 0.117880      0.108037        1.091     0.2770
{F}year[2002]                 0.277110      0.105493        2.627     0.0095
{F}year[2003]                 0.372710      0.107374        3.471     0.0007
{F}[Central]                  0.201035      0.100706        1.996     0.0477
{F}[West Central]             0.0707733     0.0619699       1.142     0.2552
cer                           0.0155876     0.00279095      5.585     0.0000
{F}CRP.proportion            -0.00313653    0.00106124     -2.956     0.0036
{F}RIM.proportion            -0.00521708    0.00177738     -2.935     0.0039

R Squared:                 0.488323
Sigma hat:                 0.338006
Number of cases:               380
Number of cases used:          162
Degrees of freedom:            151

Summary Analysis of Variance Table
Source         df       SS           MS           F    p-value
Regression     10   16.4641       1.64641      14.41    0.0000
Residual      151   17.2515      0.114248
```

Figures 16 and 17 show the CRP and RIM proportion variables plotted against the price variable, adjusted for all other variable contributions. Such added-variable plots (Cook and Weisberg, 1999) can be used for, among other purposes, checking whether certain observations distort the central finding or if there is any striking nonlinearity in the relationship. The sloping lines are the estimated coefficients for the contract variables. The linear fit seems reasonable for the CRP variable, but I'm a lot less comfortable with the RIM variable, for the same reason that I discussed above: there are simply too few RIM contracts in the dataset as long as we include the productivity variable.

For the CRP variable, however, the linearity seen in Figure 16 means we can use the estimated coefficient for the conservation restriction variables as a "multiplier" for quick valuation purposes. For example, a property with a 25% CRP restriction would sell for 25*.31=7.75% less than would an otherwise similar non-restricted property. A property with a 50% CRP restriction would sell for 15.5% less than an unrestricted property.

Figure 16: Added variable plot for Model 6, CRP contract proportion

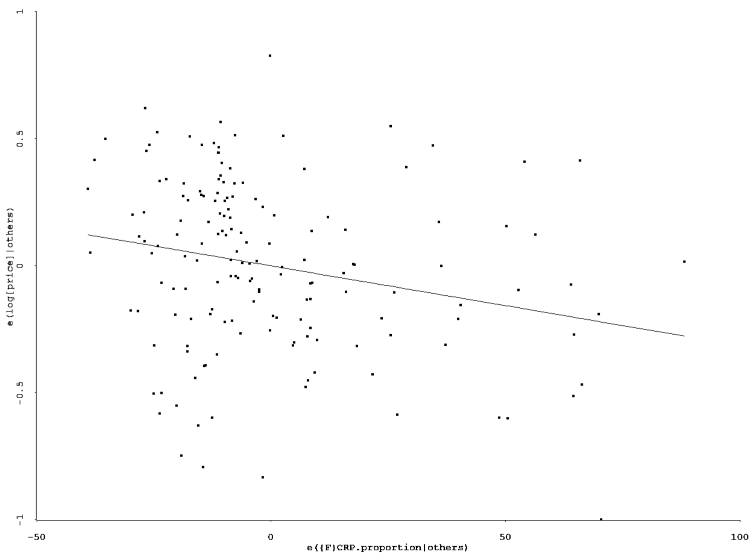

Figure 17: Added variable plot for Model 6, RIM contract proportion

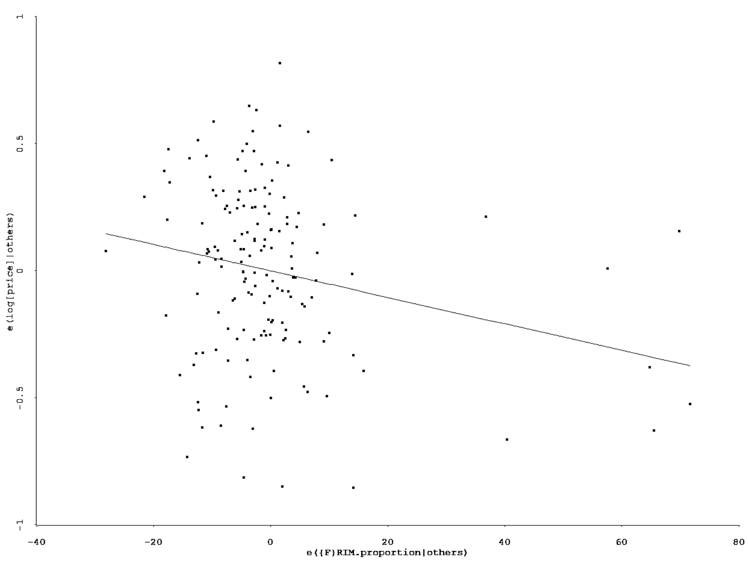

As noted earlier, including the productivity variable greatly reduces the number of sales against which to calibrate the models. Model 6 drops productivity from Model 5, thereby increasing the sample size from 162 to 380 and the number of RIM sales from 12 to 34. This gives us a hint of how the inclusion or exclusion of the productivity variable affects our conclusion about the influence of conservation contract information. The short answer: It does, but not by much. A formal F-test of this exclusion results in a very high p-value: we can reject the claim that the addition/subtraction of the productivity variable has no effect on our conclusions. The coefficients on the CRP variable are about the same, whichever model is used, but the RIM coefficient is unstable across the models, even when the expanded dataset is used.

Model 6: Not adjusting for productivity

```
Response      = log[price]
Coefficient Estimates
Label                       Estimate        Std. Error    t-value    p-value
Constant                     7.80332        0.185084       42.161     0.0000
bareland                    -0.319410       0.0660038      -4.839     0.0000
log[size]                   -0.154051       0.0317700      -4.849     0.0000
{F}year[2001]                0.204091       0.0904013       2.258     0.0246
{F}year[2002]                0.389885       0.0895703       4.353     0.0000
{F}year[2003]                0.512165       0.0933175       5.488     0.0000
{F}[Central]                 0.257269       0.0639277       4.024     0.0001
{F}[West Central]           -0.0515759      0.0615247      -0.838     0.4024
{F}CRP.proportion           -0.00356185     0.000864689    -4.119     0.0001
{F}RIM.proportion           -0.00871267     0.00118517     -7.351     0.0000

R Squared:                0.401845
Sigma hat:                 0.47978
Number of cases:              380
Degrees of freedom:           370

Summary Analysis of Variance Table
Source         df       SS            MS           F    p-value
Regression      9   57.2176       6.35751      27.62    0.0000
Residual      370   85.1698      0.230189
 Lack of fit  330   72.3997      0.219393       0.69    0.9574
 Pure Error    40   12.7701      0.319252
```

Conclusions

Figure 18: Summary of conservation contract price effects

Figure or model	Feature	Effect
Scatter plot 5	Price and size	Restricted properties have lower prices
Scatter plot 6	Price and productivity	None observed
Box plot 8	Location	Sales in West Central and Southwest have lower price than Central
Box plot 9	Price and presence or absence of contract	Restricted properties have lower prices
Box plot 10	Price and contract type	Both RIM and CRP properties have lower prices than do non-restricted properties
Box plot 11	Price and size	CRP contracts are on larger properties
Box plot 12	Price and productivity	RIM contracts are on lower productivity properties
Box plot 14	Contract proportion and contract type	RIM contracts have a higher proportion of property under contract
Model 2	Presence or absence of contract	Conservation contracts of any size lower price by 20%
Model 3	Type of contract	CRP contract of any size reduces price by 19%; RIM contract of any size reduces price by 32%
Model 4	Proportion of parcel under contract	Each percentage point increase in restricted portion reduces price by 0.4%
Model 5	Type of contract and contract proportion	For each percentage point increase in restricted portion, CRP reduces price by 0.3% and RIM reduces price by 0.5%.
Model 6	Type of contract and contract proportion; no productivity	For each percent increase in restricted portion, CRP reduces price by 0.4% and RIM reduces price by 0.9%.

Figure 18 pulls together all the findings so far. What do we conclude? The preponderance of the evidence suggests that, for this sample of sales, conservation contracts significantly and negatively influence sales prices. This holds for both long-term (RIM) and short-term (CRP) contracts. Let's call these the RIM Effect and the CRP Effect, respectively. Can these estimates be used for valuation purposes?

Figures 19 and 20 show the 95% confidence intervals for the conservation contract estimates in Models 2-3 and 4-6, respectively. I include the ranges for these same models absent the productivity variable, as in Model 6. The relative narrow range and similar magnitudes for the CRP Effects in each chart give us some confidence about using those numbers. For example, I'm pretty comfortable saying that the price effect of a CRP contract of any size lies between 6% and 30%, with a "best" estimate of 18%. Even that may be too wide a range for practical use, however: a larger sales dataset might reduce this range somewhat. I suggest some ways to get more data in the final section.

The range of the RIM Effect, in either chart, is so large as to be untrustworthy, in my judgment, even though these intervals are statistically correct. Even calibrating the model without the productivity variable, as in Model 6, so as to increase the number of RIM sales from 12 to 34, really does nothing to reduce the RIM Effect to a useful range.

Figure 19: Price effects of conservation contracts of any size (95% CI)

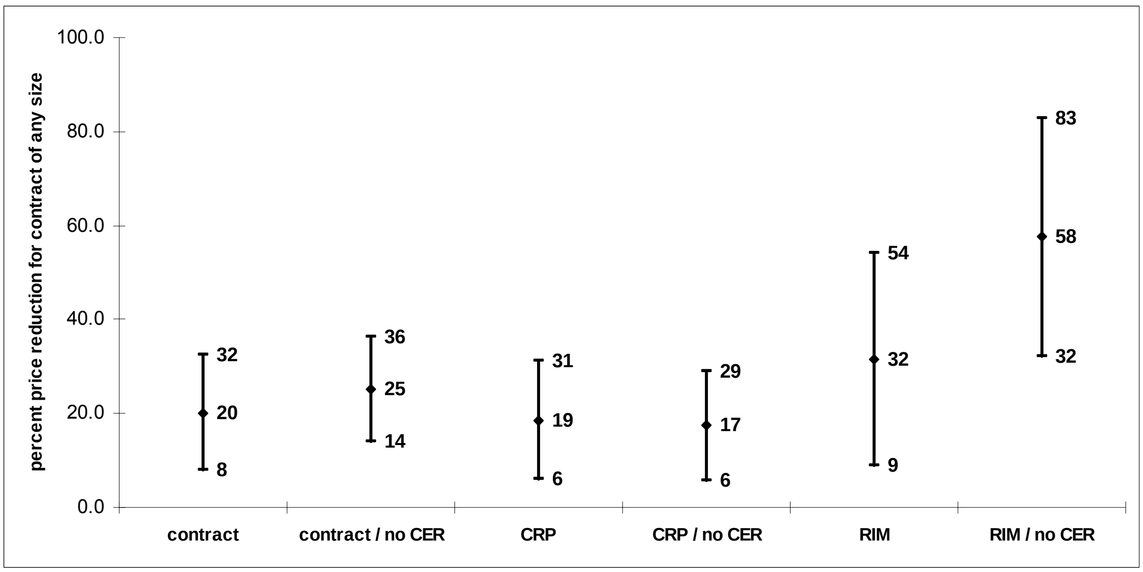

Figure 20: Price effects of one percentage point change in restriction (95% CI)

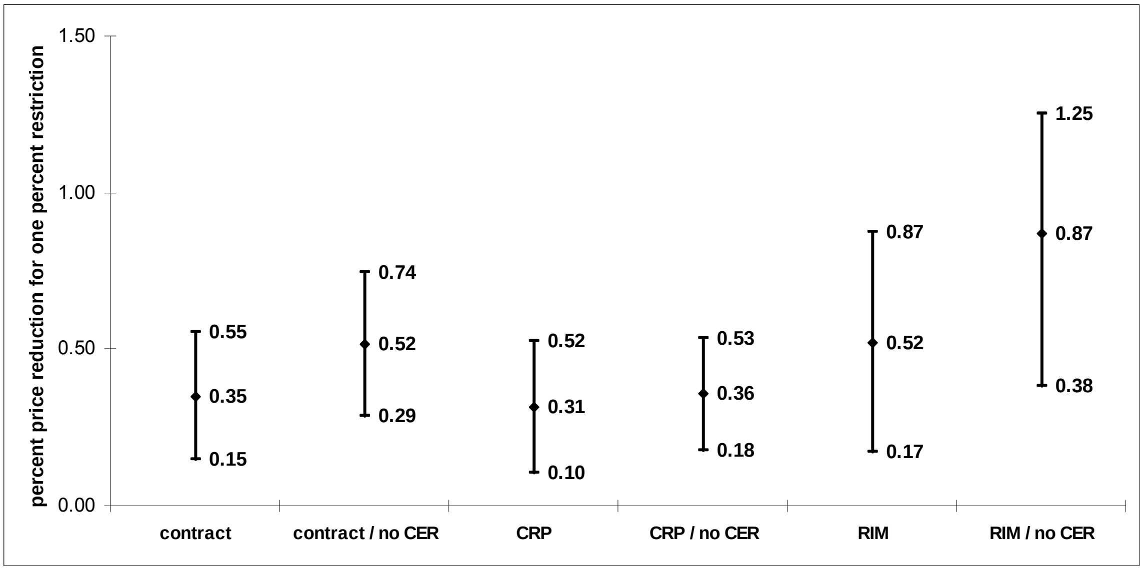

A negative and persistent CRP Effect is not what I expected. Short-term restrictions aren't supposed to influence land prices. In the next section, I'll see if I can explain the CRP Effect away. I'll conclude with a discussion of some of the implications of these findings.

What's going on here?

The CRP Effect contradicts the land valuation theory set out at the beginning of this report. Long-term contracts were expected to have negative effects on price, and that's what we found. Short-term contracts were expected to have no effects on price, but that's not what the data seem to be telling us. What's going on here?

Three are three possible categories of reasons why the world might not match our theory, why the CRP Effect was found:

1. The analysis could be wrong.
2. The data could be wrong.
3. The theory could be wrong

Let's work through each of these possibilities.

The analysis could be wrong

Maybe our theory is right, and the data would support it, but we misread the data. After all, there's a lot of statistical noise in the world; maybe our procedures for removing some of that noise were misapplied, resulting in our failing to find the "true" relationships between price and conservation contract information. Our model might have the wrong variables or it might handle variables in the wrong way.

Ideally, real estate models like those used here would examine the price effect of the variable of interest—contracts in our case—adjusted for characteristics of the buyers and sellers, of the market, and of the property itself. As discussed earlier, we lack useable information about buyers and sellers (other than the price at which they eventually settled, so we know that buyers would have paid at least that much or less, and sellers would have taken at least that much or more), so we can do no adjustment of the first type. I tried to deal with market differences by including time (year of sale) and location (region) adjusters. For the property itself, I used size and productivity. Perhaps my inclusion of these variables is flawed, or perhaps I should have included others.

Nor is there anything magic about the linear structure of the models I use here. While linear models permit the use of ordinary least squares regression, which has all sorts of convenient mathematical properties, and while linear models are the overwhelming favorite of statisticians and economists working in these areas, they still may be wrong. The data may really require a different approach. But there's nothing in the data to suggest that this might be the case. None of the models violate the standard OLS regression assumptions about linearity or non-constant variance, nor do they seem to require more complex structure than that employed here. (Details available from the author.)

The data could be wrong

Even if the model is correct, and even if the theory is correct, I might have been mislead by incorrect data. This could take several forms.

Coding errors. We always run the risk that errors in recording data may bias our results. What we traditionally do is hope/assume that these errors cancel each other out, so that our models remain un-afflicted by so-called errors in variables and that the error caused by unobserved but still important variables have certain defined properties.

Interpretation errors. We users of secondary data necessarily assume that our interpretation of the data is consistent with that of the person who entered the data. Disconnects can be obvious, or they can be hidden. For example, the recorded sales price may not be the "true" transaction price, not for reasons of tax dodging, but because of delayed payments in a contract for deed. (In this dataset, all reported sales prices have been adjusted for terms and also for time-of-year.) But there can be subtler problems. The CER variable, for example, is the weighted average for the entire property, not just for the cropped portion. If market participants are interested principally in the cropland portion, then the CER variable understates the relevant productivity measure for the sale. If the "true" effect of this variable is thereby incorrectly determined, then our model may incorrectly measure the effect of other variables such as conservation contracts.

Wrong assignment. The type and extent of the conservation contract is both the most important information for this study and the one that might be the most suspect. Until very recently, conservation contracts were not typically recorded on the CRV. (Although there has long been a place on the form for this information, its completion has until recently not been a priority for the Revenue Department officials who oversee the compilation of the data.) Too, the information is often not readily available to the person filling out the form, even if was relevant to the transaction itself. CRP contracts are administered by the local USDA FSA office, and their existence is not required to be recorded in the local property registration offices because only permanent contracts are considered to be property transfers. RIM (and CREP) contracts are considered to be property transfers, so the easement is officially recorded at the registration office, but that detail is not always picked up by the person filling out the form.

Two possible kinds of error could enter during this process. First, of course, local officials might have missed recording some sales that did have conservation contracts. Such sales would have entered into the regular agricultural land sales database and might by chance have been selected as a member of the control set when I sampled from it. So some of the sales I analyzed might be incorrectly listed as having no contracts, when in fact they did. (The reverse possibility, that some of the contract sales were incorrectly labeled, is even less likely, in my judgment.)

The other possible assignment error stems from confusion about what sorts of contracts are actually on the properties examined here. It's possible that CREP sales, for example, were mistakenly identified as CRP sales, or the reverse. Or a CREP sale might have been listed with the RIM sales, because RIM is a component of the contract.

All of these data errors could have occurred, but I judge their cumulative effect on model calibration to likely be negligible. Experienced professionals built this dataset.

The theory could be wrong

If our data are good, and our analysis is good, then our theory must be wrong. Could it be modified to account for the CRP Effect? I can think of at least four alterations that might make the Effect a logical conclusion, rather than something to be explained away. (These might deal with the theory, but they won't automatically help with the policies and practices in which the conventional theory is embedded.)

Buyers have different income expectations than sellers. Sellers presumably entered into a CRP contract thinking they were better off doing so than staying with crop production. Maybe the set of buyers in the market think that the remaining years of annual CRP payments, minus re-conversion costs, are actually lower than would be the case if the land had remained in crop production. (The buyers might think crop prices will be higher during the remaining CRP years than did the sellers; or they buyers might think their production costs would have been lower than were the sellers'.) Either way, the buyers would have different opportunity costs than did the sellers when they entered the CRP contract. Under reasonable assumptions, however, it can be shown that the buyer's expected income would have to be at least 50% higher than the remaining CRP payment to bring about even a modest negative CRP Effect. As it turns out, the single factor which leads to the most negative CRP Effect is when the buyer thinks that the post-CRP conversion cost is a couple hundred dollars per acre more than the seller implicitly calculated. This is highly unlikely except where trees were planted.

Buyers have different preferences than the sellers. Although all buyers of properties in this dataset stated on the Certificate of Real Estate Value (CRV) that they intended to keep the land in agriculture, maybe they're looking more at the non-agriculture characteristics of the RIM properties than do the buyers of non-contract land. For example, RIM buyers may have a high preference for included wetlands, whereas non-contract buyers would discount these non-tillable portions of the property.

Sellers of contracted lands are different from the sellers of non-restricted lands. Perhaps the sellers are strongly motivated to sell, they're no longer interested in farming (witness their readiness to enter into the CRP in the first place), so they'll settle for a lower price than still-"serious" farmers.

Properties with conservation contracts constitute a distinct land market. Perhaps a contract signals to potential buyers that there is something about the property that distinguishes it from regular cropland, even though the agronomic productivity is no different. If restricted lands are different in ways that we don't account for here, we perhaps err in comparing them to sales with no restrictions. Our findings would then be interpreted as simply reflecting differences between two different types of land, not between different sets of property rights. In that case, however, we'd not be likely to see the strong association not just with presence or absence of a contract but also with the amount of the property subject to the contract.

Policy implications

I think the RIM Effect shown in Models 5-6 is based on too few sales to support any strong recommendations. But the CRP Effect is based on plenty of sales, and it seems consistently strong. This could certainly provide justification for a landowner demanding a down-valuation during the course of a CRP contract, even though it flies in the face of conventional valuation theory. Minnesota (and many other states) requires that property be assessed at its market value, the assessor's estimate of what the land would sell for were it to sell today. These estimates are made each year, and they rise and fall as market conditions change. And we've just demonstrated—if you accept the results of this study—that the market says CRP land is worth less than non-restricted land. Should not these properties' assessed values then be reduced accordingly, at least during the course of the CRP contract?

I'm not prepared to recommend such a policy change yet, however. I'd like to see even more data, for several more locations, first. While we had access to 190 conservation contract sales, only 34 were for permanent RIM status, and this set was further reduced by the scarcity of observations for which we had the necessary productivity data. Perhaps our disconcerting findings are attributable to simple "micronumerosity" problems: more observations would correct our results. (Of course, simply adding observations will not in and of itself remove the CRP Effect, already "significant." If the new observations aren't different in complexion from those we already have, then a higher number of sales just results in a smaller standard error and, as a result, an even smaller p-value.) Additional sales might also provide sufficient sample size to analyze the influence of CREP easements separate from standard RIM easements, as well as assess the true size of the RIM Effect.

One way to get more sales to analyze would be to retroactively assign CERs to those existing sales for which we now lack this data. The problem is that many counties simply lack the data, either because they haven't conducted the necessary studies or because they lack the detailed soils surveys required to calculate CERs in the first place. However, as we saw in Model 6, the CER productivity data is not altogether critical to the analysis. Perhaps more useful would be inclusion of additional conservation contract information such as the number of years remaining in CRP contracts (to see how remaining payments are treated by the market) and the type of conservation practices installed on RIM properties (wetlands, trees, riparian buffers, etc.).

The best way to get more sales for systematic analysis is to (1) reach into more parts of the state (especially the South Central and the North West) and (2) wait until the 2004 farmland sales records, with their more carefully tracked conservation contract data, become available in Spring 2005.

To make conservation contract tracking more efficient, I have two recommendations:

1. For RIM and CREP contracts, the state's Board of Soil and Water Resources (BWSR), which administers these contracts, should require that every contract record have attached to it the county property tax office's property identification number. Then it would be easy to link contract information with sales information, either when the CRV is prepared or when analysts examine the compiled statewide sales files.

2. County USDA FSA offices should work with county assessors to ensure that CRP contracts can be cross-linked to property tax records, probably through conforming GIS platforms. FSA records are increasingly GIS-based, as are property tax records in many counties, but the two systems don't always match up. At the very least, county assessors should start to carry even "temporary" CRP contract information in the property tax files, so that contracted-land sales can be more readily identified when a CRV is prepared.

References

Cook, R. D. and S. Weisberg. 1999. **Applied Regression Including Computing and Graphs**. John Wiley and Sons.

Shultz, S. and S. J. Taff. 2004. *Implicit Prices of Wetland Easements in Areas of Production Agriculture*. **Land Economics**. 80(4), November.

Taff, S. J. 2004. *Minnesota Farm Real Estate Sales: 1990-2003*. University of Minnesota Department of Applied Economics. Staff Paper P04-7. April.

Viaggi, D. and S. J. Taff. 2004 *Public Acquisition of Property Rights To Serve Agriculture/Conservation Policy: Lessons from Italy and the US*. Paper presented to Ninth Joint Conference on Agriculture, Food, and the Environment. Conegliano, Veneto, Italy. September 3.