Optimal Sequential Plantings of Corn and Soybeans Under Price Uncertainty

Michael Livingston,♣ Michael J. Roberts,♠ and Yue Zhang♠
November 2012

Abstract

We examine crop choice and fertilizer applications as the solution to stochastic dynamic optimization problem for an infinite stream of discounted profits. The efficient decision rule depends on the stochastic evolution of commodity prices, fertilizer prices, and the agronomic effects of rotation versus monoculture. This decision rule includes an account of real option values associated with maintaining soil quality and disposition in an environment with highly uncertain future prices and irreversible past planting decisions. We parameterize a baseline model for a representative acre in Iowa and compare the model's predictions and profits to relatively naive, shorter-horizon decision rules, and a field managed with optimal fertilizer applications conditional on corn and soybeans always being rotated. We also examine the effects of a permanently larger premium on corn prices relative to soybean prices, which has been observed in locations near recently established ethanol plants. We then compare the various decision rules to actual crop choices in a panel of over 6500 Iowa plots during 1979-2007. As compared to less forward-looking objectives, we generally find the agronomic benefits of rotations coupled with real option values leads to a much more inelastic response of planting decisions to both transitory and permanent price changes. Always rotating, regardless of prices, is extremely close to optimal.

An early version of this article benefited greatly from discussions with John Rust and also from comments provided from attendees at the 2008 Agricultural and Applied Economics Association meeting in Orlando, Florida. Views expressed are the authors and do not necessarily represent those of the US Department of Agriculture.

♣ US Department of Agriculture, Economic Research Service, 1800 M Street NW, Washington DC 20036. Email: mlivingston@ers.usda.gov.

♠ Department of Agricultural and Resource Economics, North Carolina State University, Box 8109, Raleigh, NC. Email: michael_roberts@ncsu.edu and yzhang19@ncsu.edu.

1 Introduction

Corn and soybeans, the two largest crops in the United States and among the most important food staples in the world, are typically grown in rotation. That is, on any given parcel, a farmer will typically grow soybeans in one year and corn the next. For example, based on the US Department of Agriculture's ARMS Phase II survey data, an estimated 63 percent of acres planted to corn in Iowa in 2005 were planted with soybeans the year before. The main incentive to rotate crops comes from higher yields and lower input costs enjoyed by farmers who rotate in comparison to farmers with similar land that grow the same crop every year. Yields of rotated crops are higher because rotations reduce pest problems and enrich soils. Soybeans, for example, fix nitrogen that is used by the subsequent corn crop, thereby reducing fertilizer costs for corn (Mallarino, Ortiz-Torres & Pecinovsky 2005, Hennessy 2006).

While the basic agronomic benefits of rotations are well known and understood, the phenomenon presents a challenging stochastic dynamic economic problem surrounding optimal planting decisions in an environment when commodity and fertilizer prices are uncertain and highly autocorrelated over time. In this article we show how the solution to this problem has important implications for the supply of these important commodities. Specifically, we show that in comparison to the solution to the static problem, the forward looking, dynamic optimum is much less likely to deviate from rotating corn after soybeans and vice-versa, in response to price movements. This is largely true even when static optimizers take into account the agronomic effects of past plantings. Somewhat counterintuitively, a farmer applying a simple rule of always rotating, regardless of price, often performs better over the longer run than a farmer optimizing over a shorter time horizon. Stochastic-dynamic considerations may therefore help to explain an inelastic supply response of agricultural commodities or, more generally, farm management practices that nearly mimic agronomic maximization, at least with respect to crop choice.

A key feature of our solution is that it accounts for real option values connected with price uncertainty and irreversible past planting decisions. These option values are associated with choosing crop rotations that sustain soil disposition and reduce susceptibility to pest outbreaks over the long run, investments that pay off particularly well in the event commodity or input prices turn out unexpectedly high. Thus, deviating from an agronomically optimal rotation in response to a change in relative prices may be profitable in the short run, but may not be economically efficient over the longer run. These short-run versus long-run tensions imply a supply response in a dynamic model under uncertainty that is less responsive to temporary or permanent price shocks than is implied by static models.

Burt & Allison (1963) were first to consider crop rotations in a stochastic and dynamic programming (SDP) framework. Their article, among the first applications of SDP, considered the decision to leave a field fallow or to plant wheat in an environment where soil moisture evolved according to a Markov process. In subsequent work Burt and a few other scholars have considered specialized crop management decisions in contexts with random state variables pertaining to pest problems, irrigation water, and other agronomic factors. Many have used of SDP for modeling commodity storage problems and for agricultural, biological and ecological modeling (Kennedy 1986, Williams & Wright 1991, Deaton & Laroque 1992). But the literature has generally been slow and reluctant to use SDP to model planting decisions more broadly, especially as a positive model to describe farmer behavior. To our knowledge, this article is first to explicitly consider the consequences of price uncertainty on sequential crop planting decisions using an SDP framework.

Farmers maximizing the expected present value of profits will choose crops by weighing the agronomic benefits of rotations together with the current and expected distribution of prices. For example, high corn prices in the current year could make expected profits from planting corn after corn greater than planting soybeans after corn, despite the agronomic benefits of planting soybeans. Planting corn after corn this year, however, would reduce potential profitability of planting corn in the following year. A decision to plant corn after corn may or may not be worthwhile, but calculation of the optimal choice is not obvious or simple. The optimal decision depends on the relative prices of corn and soybeans in the current year, past plantings and the whole distribution of anticipated prices for corn and soybeans in futures years. Over a multiple-year horizon, price uncertainty matters too, even if farmers are risk neutral, due to real option values.

Crop rotations have been incorporated into some linear programming models (El-Nazer & McCarl 1986, Musser et al. 1985, Johansson, Peters & House 2007), but these models do not account for the sequential nature of planting decisions, or for option values related to price uncertainty. That is, these models take multi-year rotation rules as a single decision, with future prices assumed known in advance, as if the problem were static, nonsequential, and nonstochastic.[1] In some earlier models of aggregate supply, current supply is sometimes conditioned on past aggregate plantings, which may proxy for rotational effects, but there is no explicit account of parcel-specific choices (Eckstein 1984, Burt & Worthington 1988).

Given the importance of crop rotations has long been widely acknowledged, it is natural to ask why there has been so little modeling of rotational decisions using an SDP framework.

[1]There are, of course, great computational advantages to these alternative approaches.

The most obvious answer is that SDP modeling with more than one or two continuous state variables can be computationally expensive and difficult. But time and rapidly advancing computer technology have lowered the computational expense of larger-scale SDP problems. Also, methods in computational economics have matured, with several contributors making their computer code publicly available, further lowering the costs of implementing SDP models. While computational limits still constrain the complexity of these models, with a few simplifying assumptions, we are able to develop a model that captures the most salient aspects of the corn versus soybeans planting decisions that govern a significant share of highly productive cropland in the midwestern United States.

In the next section we describe the parameterization of the model. We then report econometric estimates of the key parameters, followed by a description of the algorithm used to solve the SDP. We then compare profits and crop choices of the optimal infinite-horizon policy to policies derived from assuming maximization over one- and two-year horizons rather than the infinite horizon SDP, as well as a policy that optimizes input use conditional on an "always rotate" rule-of-thumb. To explore the potential effects of recent ethanol expansion on planting decisions, we consider the effect of premium corn prices that have been documented for fields located near ethanol plants. Finally, we compare predictions from the infinite horizon model and alternative objectives to actual crop choices observed on parcels sampled by the National Resources Inventory.

2 A Stochastic Dynamic Model of Crop Choice

Consider planting decisions for a standardized unit of land. At a sufficiently small scale, planting decisions on an individual unit of land, such as a field, crop choice is a discrete decision, even though when aggregated across all units for a given farm, or in a county or state, the decision will approximate a continuous decision (e.g. what fraction and which parcels of land to allocate to corn, soybeans, wheat, and so on). To focus squarely on the issue of rotational dynamics, we assume no time or capital allocation constraints, spatial interaction effects, or farm-level capital or liquidity constraints that would force us to approach planting decisions at the farm level. Instead, each unit of land is regarded as an independent "profit center" and, by maximizing profit from harvesting crops on each unit of land, the farm maximizes its value as a whole.

There is some arbitrariness about how big or small an individual unit of land may be. It should usually be treated as a contiguous parcel of land that is typically growing only one crop per season. We might think of this as a "field" where, for agronomic reasons, it would not make sense for a farmer to plant different crops on the same field. If this is not the case, then the "field" should be conceptually subdivided into smaller parcels for which the farmer almost always plants only one crop or the other. For purposes of this first analysis we will abstract away from the size of the unit even though in a more general model unit size may be chosen simultaneously with crop choice. As such, we also abstract away from portfolio considerations that have been a traditional focus of planting decisions under uncertainty. The data to which we compare our model's predictions refer to specific points where the discrete crop choice is indicated. We report value and output on a per-acre basis.

Our yield data on crop rotations and input use, which we use to calibrate the model, are from an experiment station in Northeast Iowa. As such, the model we develop is most relevant to that region of the country and nearby regions with similar soils and climate. Iowa, being the largest corn and soybean producer in the United States, and likely the highest concentration of corn and soybean production in the world, is a salient focus.

The main features we want the model to be capable of predicting are (1) the pattern of planting, (2) the relative frequency of corn versus soybean plantings, (3) the sensitivity of optimal planting decisions to exogenous changes in commodity or input prices. There are two ways in which we examine the question of price sensitivity. First, prices change over time in a stochastic fashion. Second, evidence in previous research shows that corn prices are higher near recently established ethanol plants. A key focus will be on how decision rules differ depending on the planning horizon of the farmer, and which best predicts actual planting decisions.

A key simplifying assumption is to model crop revenues per acre, price multiplied by yield, rather than keeping prices and yields separate. This simplification accounts for price-yield correlations stemming from spatially-correlated weather and pest outcomes. Historical revenue-per-acre data also appear strongly stationary, despite a significantly increasing trend in yields and decreasing trend in real prices. Stationarity is necessary for describing the stochastic autoregressive processes for our state variables that we discuss below. This one simplification therefore solves both the problem of price-yield covariance and the problem of non-stationarity.

In general, a long history of crop choices and soil management choices may influence current expectations for yield and revenue. Incorporating a long crop history for a large number of possible crops would greatly increase the state space and render a solution computationally infeasible. Thus, to simplify the problem we consider economically efficient choice among just two crops, corn or soybeans, together with the optimal continuous application of nitrogen fertilizer, where we denote crop choice, i_t, and nitrogen fertilizer use, n_t, for a field of arable land, on which either soybean, $i_t = 0$, or corn, $i_t = 1$, may be produced. Although just two crop choices limits the scope of the model, these two particular crops are exclusively planted, or nearly so, on a large share of the nation's most productive cropland.

We consider models in which up to three years of crop history may affect current revenues. Crop choices more than three years past are not economically or statistically significant in regression analysis of yields from experimental field plots (we show this below). To account for this history, we denote a crop-specific adjustment factor $a^i(\mathbf{I}_t, n_t)$, where $\mathbf{I}_t$ denotes crop history at planting time in year t. This function gives a proportional adjustment to expected revenues for each crop i that depends on past plantings and nitrogen fertilizer application in the current year (n_t).

Field-level expected revenues are also tied to the stochastic evolution of prices and to the broader covariance between prices and yield. To account for both autocorrelation of prices and current-period covariances between prices and yields, we model current expected revenues for any given field as being tied to expectations about *state-level* revenues per acre. State-level revenues, like prices, are observed, publicly available and exogenous to field-level decisions. State-level revenues per acre equal the average price received in Iowa multiplied by the realized state-level yield. We denote these state-level revenues by r_t^c and r_t^s for corn and soybeans, respectively. State-level revenues per acre, like prices, display strong autocorrelation, so that past revenues strongly influence expectations about current revenues. *Field-level* revenues are given by the crop-specific adjustment factor multiplied by state-level revenues. The idea is that an individual farmer's planting and fertilizer application decisions will affect his or her own revenues by same proportion that it affects the farmer's yield. In other words, we assume field-level marginal yield effects are too small to affect price or yield at the state level.

$$\text{Field Revenues} = r_t^i a^i(\mathbf{I}_t, n_t) \tag{1}$$

The other critical factor affecting profits of corn relative to soybeans is nitrogen fertilizer prices, which are autocorrelated and show strong association with state-level corn and soybean revenues per acre (this is shown in an appendix). Unlike revenues per acre, however, fertilizer prices, denoted f_t, are observed at the time of the planting decision.

We assume expectations about about current and future profits are a function of past state-level revenues per acre for corn, past state-level revenues per acres for soybeans, current fertilizer prices, and planting history. The farmer makes optimal planting and fertilizer decisions while anticipating how current decisions affect both current and future profit opportunities.

To evaluate these expectations, we assume these three state variables follow a 1^{st}-order vector autoregressive process:

$$\mathbf{p}_{t+1} = \mathbf{a} + \mathbf{p}_t^T \mathbf{B} + \mathbf{u}_{t+1}, \tag{2}$$

where $\mathbf{p_t}$ is a vector comprised of lagged corn and soybean state-level revenues per-acre and current-period fertilizer prices:

$$\mathbf{p}_t = \begin{bmatrix} log(r^c_{t-1}) \\ log(r^s_{t-1}) \\ log(f_t) \end{bmatrix}, \tag{3}$$

$\mathbf{a}$ and $\mathbf{B}$ are a vector and matrix of parameters, and $\mathbf{u}_{t+1}$ is a vector of random innovations with a variance-covariance matrix equal to $\mathbf{\Omega}$. We assume multivariate normal random innovations, which approximately fits the data. We use a first-order model because standard model selection criteria (AIC and BIC) strongly preferred a first-order model. No higher-order terms showed statistical significance. This is fortunate because each additional lag in the vector autoregressive process would have added three continuous state variables to the SDP, which would be prohibitively expensive from a computational standpoint. (We describe the data below.)

To calculate current profits we simply subtract costs from revenues. The only costs we explicitly consider are nitrogen fertilizer expenditures. While other costs and inputs surely matter for field-level profits, nitrogen fertilizer inputs are the largest single variable expenditure that interacts strongly with the corn-or-soybeans crop choice. For one, nitrogen fertilizer is not typically applied to soybeans but is almost universally applied to corn. Also, there are strong substitution possibilities between fertilizer applications and rotational decisions. For example, corn monoculture requires greater levels of fertilizer to achieve the same yield as compared to corn following soybeans. Other inputs are generally of lesser expense or are typically similar regardless of whether corn or soybeans are planted. These other inputs include phosphorus, potassium, labor, land and capital utilization. The current profit function is thus given by:

$$\pi(i_t, n_t | \mathbf{I}_t, p_t) = r_t^i a^i(\mathbf{I}_t, n_t) - i_t f_t n_t. \tag{4}$$

The producer's objective is to maximize the expected present value of profit (4) over an infinite time horizon (we also consider shorter horizons), subject to the stochastic evolution of state-level revenues and fertilizer prices (2). We can write this infinite horizon problem using the recursive Bellman equation that relates the current value function to the future value function:

$$V(\mathbf{I}_t, \mathbf{p}_t) = \max_{i_t n_t} \mathbb{E}\left[\pi(i_t, n_t | \mathbf{I}_t, \mathbf{p}_t) + \theta V(\mathbf{I}_{t+1}, \mathbf{p}_{t+1})\right], \tag{5}$$

where V is the maximum expected present value of the field over an infinite horizon of optimal corn and soybean plantings. It is a function of three discrete and three continuous state variables: the previous three crop choices, past unadjusted (state-level) soybean and corn revenues, and current nitrogen fertilizer prices. We assume these variables provide all the information necessary to form expectations about end-of-season profit and the next period's discounted value function, where θ is the discount factor.

Given parameters for the autoregressive process of revenues, a functional form for the agronomically-based, revenue-adjustment functions, $a^i(\mathbf{I}_t, n_t)$, and a discount factor, θ, the dynamic optimization problem is fully defined. The solution to the model is a policy function that gives crop choice and nitrogen fertilizer applications as a function of the state variables: $i(\mathbf{I}_t, n_t)$ and $n(\mathbf{I}_t, n_t)$. These policy functions are the choices that satisfy the Bellman equation (5). Because the value function does not have an analytical representation, the model must be solved numerically. We use value function iteration to find these policy functions.

3 Solving the Model

Here we describe the numerical methods used to examine solutions to the SDP models of crop choice. For readers uninterested in the technical details of how these problems are solved, this section may be of little interest and safely skipped without significant loss of applied content. Similarly, readers familiar with these techniques are unlikely to learn new methods. We provide these details for completeness and to aid replication. Miranda & Fackler (2004) provide an in-depth discussion of the numerical methods we use. In addition, the computer code we use to approximate functions and perform quadrature to evaluate integrals was written by Paul Fackler, which can be downloaded from his website.

Each model we examine includes three or four state variables. The three state variables included in all models are: crop history, which embodies from one to three previous years of plantings; the previous year's average revenue received per corn acre, r^c_{t-1}; and the previous year's average revenue received per soybean acre, r^s_{t-1}. In models with just three state variables the fertilizer price is assumed fixed. In other versions of the model, the current year's nitrogen fertilizer price, f_t is the fourth state variable.

To solve the model numerically, we first divide the state space into a discrete number of points. The value function is approximated at each discrete point and we interpolate values in-between points using a linear spline. The vector $\mathbf{s}$ denotes the discrete evaluation points and $V(\mathbf{s})$ the associated points in the value function. The vector $\mathbf{s}$ thus spans all combinations of evaluation points across all possible states.

Computational cost is tied to the size of $\mathbf{s}$. In a model with one year of crop history, the crop-choice state variable, $\mathbf{I}_t$, may assume one of two values: zero if soybeans were planted in year $t-1$ and one if corn was planted. In the model with two years of crop history, $\mathbf{I}_t$ may assume one of four values, one for continuous soybean (SS), two for soybean followed by corn (SC), three for continuous corn (CC), and four for corn followed by soybean (CS). Similarly, in the model with three years of crop history, the crop-choice state variable may assume one of eight feasible crop histories.

We divide the three continuous state variables (the vector $\mathbf{p}$) into 10 to 20 equally-spaced evaluation points that range between \$100 and \$1,600 per acre, \$100 and \$1,700 per acre, and \$0.05 and \$1.60 per pound, for state-level corn revenues, soybean revenues and fertilizer prices, respectively. These ranges are somewhat larger than those observed in the data (described below). When the crop history includes a single year, the previous year's crop choice can assume only two values. Multiplying out the set of feasible evaluation points, we have a minimum of $2 \times 10 \times 10 \times 10 = 2,000$ evaluation points in vector $\mathbf{s}$. The number of evaluation points then doubles for each additional year of crop history considered.

Increasing the number of evaluation points for continuous state variable involves a delicate balance between accuracy and computational expense. If we choose 20 points rather than 10 for each of the three continuous state variables, the dimension of the state vector increases from 2,000 to 16,000 (2×20^3) for a single-year crop history; with 30 points the dimension would be 54,000 (2×30^3). If we solve the model using 2-years of history, then we need to allocate 4 evaluation points for the first state variable which again doubles the memory requirements. Though expensive, in some applications the number of evaluation points can significantly influence how well the overall model approximates the true solution.[2] Our

[2]For example, a recent article by Cafiero et al. (2009) reconciled a long-standing puzzle about the amount of autocorrelation in storable commodity prices by simply increasing the number of evaluation points in the stochastic dynamic programming model. Approximating the inverse demand function with just 20 evaluation points smoothes over a critical kink in the policy function where stockout occurs. This seemingly minor

computational constraint makes it difficult to solve models with more than 10 evaluation points when we use all four state variables. If we fix fertilizer prices and thereby reduce the number of continuous state variables to two, we are able to increase the number of evaluation points to 20. We therefore use the model with fixed fertilizer price to test the sensitivity of the solution to the number of evaluation points.

The dynamic program is solved recursively. It begins by solving the terminal condition of the dynamic programming model, which is a static expected profit function (equation 4) conditional on each of the 2,000+ evaluation points in the state vector $\mathbf{s}$. This decision involves finding the optimal crop to plant and, if corn is planted, how much nitrogen fertilizer to apply. In this initial step the longer-run consequences of these decisions are ignored. We store the expected profit-maximizing crop choices and per-acre profit levels and then compute the coefficients of a four-dimensional linear spline so that we may evaluate expected profits in-between the evaluation points. This spline function is the first estimate of the value function (or terminal condition) which we denote $V_T(\mathbf{I}, \mathbf{p})$.

The next step involves substituting the first estimate of the value function into the right-hand-side of the Bellman equation (5) and solving the two-period problem: maximize current profits plus the discounted expected value of $V_T(\mathbf{I}, \mathbf{p})$. This optimization depends on the current state, $(\mathbf{I}_{T-1}, \mathbf{p}_{T-1})$. To trace out the whole value function we therefore perform this optimization for each evaluation point in the state vector $\mathbf{s}$. Optimization requires that we integrate $V_T(\mathbf{I}, \mathbf{p})$ over the probability distribution of future states, conditional on each $s_j \in \mathbf{s}$. While crop choice affects the future state space in deterministic way, the distribution of the vector $\mathbf{p}_T$ conditional on s_j is given by the vector autoregressive process described above (2) and its assumed trivariate normal distribution of innovations, which have mean zero mean and an estimated covariance matrix $\mathbf{\Omega}$ (reported below). [3]

Having stored optimal values and optimal choices for each $s_j \in \mathbf{s}$), we again compute the coefficients of a four-dimensional linear spline so that we may approximate expected values in-between the evaluation points in $\mathbf{s}$. This spline function gives the next iteration's estimate of the value function, which we denote $V_{T-1}(\mathbf{I}, \mathbf{p})$.

computational issue leads to a spurious implication that commodity prices should have little autocorrelation (Deaton & Laroque 1992, Deaton & Laroque 1996). In a model with 1000 evaluation points, the numerical solution accurately recovers the kink in the demand curve and predicts commodity prices with high degree of autocorrelation, which is what we observe in commodity price data. Since we do not expect to there to be non-smooth kinks in our policy functions, the relative coarseness of our policy function approximation should not cause a severe problem.

[3] In solving the model, the error terms and weights are drawn following the multivariate normal distribution using a Gaussian quadrature rule (Miranda & Fackler 2004, p. 97). We provide evidence in an appendix to show that the normal-distribution assumption appears reasonable in this case.

By substituting $V_{T-1}(\mathbf{I}, \mathbf{p})$ into the right-hand side of the Bellman equation, we can repeat this process until either a given planning horizon has been considered or until the value function converges to very similar values between iterations. Full convergence satisfies the Bellman equation and gives the solution to the infinite-horizon problem. We therefore repeat the procedure until the maximum absolute value of the change in the value function at each state vector was less than \$0.01, which generally occurred after slightly over 200 iterations. We also considered one and two-year-horizon problems as a basis for comparison.

4 Data and Parameter Estimates

The main parameters are the agronomic revenue adjustment functions $a^i(\mathbf{I}, n_t)$ and the matrix of autoregressive coefficients $\mathbf{B'}$s from equation 2. Estimation of the adjustment cost function comes from analysis of experimental plot data from Northeastern Iowa that was generously provided by David Hennessy. Estimation of the autoregressive coefficients comes from analysis of historical data on Iowa corn and soybean prices and yields that are publicly available from USDA's National Agricultural Statistical Service (NASS) and from USDA data on nitrogen fertilizer prices.

4.1 Revenue Adjustment Functions

The experimental plots are comprised of 7 large plots with 12 sub-plots in each plot. One of the four fertilizer application rates (0, 80, 160, and 240 lbs./acre) are applied and held constant over time for a specific sub-plot.[4]

The data from the experimental plots, including yields and controlled fertilizer usage, is summarized in table 1. Note that fertilizer application rates are balanced via the experimental design across rotation treatments, so that differences in average yields, conditional sampling error, can be interpreted as causal.

The data show that corn yields following soybean plantings average about 28 percent greater than corn yields following corn plantings. Planting decisions in years prior to the immediately preceding year appear to have little influence on yield. For soybeans, yields following corn plantings average about 25 percent greater than yields cultivated in soybean

[4]The fact that fertilizer application rates are held constant over time in the experimental plot data is one reason why our adjustment function does not include past application rates. Another reason is computational feasibility: including past application rates in the production function would add yet another state variable to the dynamic programming model.

Table 1: Summary statistics from experimental field plots

	(**C** ≡ Corn and **S** ≡ Soybeans)							
Past Plantings:	**C-C-C**	**C-C-S**	**C-S-C**	**C-S-S**	**S-C-C**	**S-C-S**	**S-S-C**	**S-S-S**
	Corn Yields (bushels/acre)							
Mean	106.9	105.5	107.9	na	135.5	137.6	na	na
SD	42.8	42.0	43.8	na	36.2	36.6	na	na
N	300	300	600	na	600	300	na	na
	Nitrogen Applications on Corn (lbs./acre)							
Mean	120.0	120.0	120.0	na	120.0	120.0	na	na
SD	89.5	89.5	89.5	na	89.5	89.5	na	na
N	300	300	600	na	600	300	na	na
	Soybean Yields (bushels/acre)							
Mean	49.8	48.3	45.2	na	na	na	na	38.5
SD	11.5	11.0	11.4	na	na	na	na	9.8
N	300	300	300	na	na	na	na	300
	Nitrogen Applications on Soybean (lbs./acre)							
Mean	120.0	120.0	120.0	na	na	nc	na	120.0
SD	89.5	89.5	89.5	na	na	na	na	89.5
N	300	300	300	na	na	na	na	300

Notes: The table reports average yields and fertilizer application rates for different rotations in the experimental data. Column headings give different rotation histories. For example, **C-S-S** means corn was planted in the immediately preceding year and soybeans were planted both years preceding the last. Not all feasible rotation histories are present in the data, as indicated by "na." Nitrogen fertilizer was applied to soybeans but will be shown insignificant below.

monoculture. In comparison to corn, average soybean yields are more sensitive to plantings two and three years past. The greater the frequency and more recent prior corn plantings, the greater the current soybean yield. Average yield for soybeans following three years of corn is about 9 percent greater than average yield in continuous corn-soybean rotation.

We use regression analysis and the data summarized in table 2 to calibrate the adjustment cost functions $a^i(\mathbf{I}, n)$. We report several specifications so readers can judge the tradeoff between parsimony and goodness of fit. Using a longer crop history improves the fit slightly, but accounting for it requires a larger state space in our dynamic programing model. We also experiment with different kinds of interactions between fertilizer applications and rotation history. The regressions show that the fit improves little after an account of the previous year's plantings and fertilizer use. We also find a strong interaction effect between past plantings and nitrogen use for corn: the marginal productivity of fertilizer is uniformly greater when corn is planted after corn rather than after soybeans.

The regression models have the form:

$$log(Y_{pt}) = +a_t + \mathbf{I}_{pt}\beta + \gamma_1 \log(n_p + 1) + \gamma_2(log(n_p + 1))^2 + \epsilon_{pt} \tag{6}$$

where Y_{pt} is yield on plot p in year t, a_t is a year fixed effect, n_p is nitrogen application rate in pounds per acre (fixed across years), $\mathbf{I}_{pt}$ is a vector of 0-1 dummy variables indicating the rotation history and ϵ_{pt} is the model error. Note that while the experimental data include only four discrete levels of nitrogen application, we treat it as a continuous variable in the regression analysis. This treatment allows us to infer yield and revenue effects over a continuum of application levels in our dynamic programing model. Separate models were estimated for corn and soybeans. For soybean yields, nitrogen is not applied and the applicable terms in equation 6 are dropped.

Because the regressions account for rotation history using dummy variables, and because some feasible crop histories are not present in the data, we must interpolate to make yield predictions for the unobserved histories. We do this by combining regression coefficients from similar non-missing histories. For example, an SSC history is not in our data of soybean yields, but SSS is in our data, as are CCC and CCS. We therefore approximate the soybean yield for SSC with that for SSS, plus the difference between yields for CCC and CCS, because this difference is akin to the difference between SSC and SSS. We make similar interpolations for all missing rotation histories. Results from these interpolations are shown in table 3. Holding fertilizer applications fixed at 130 pounds for corn yields, and taking the average of year fixed effects, the table reports average predicted yield for each feasible rotation history.

To parameterize the revenue adjustment functions ($a^i(\mathbf{I}, n)$) from the yield regressions, we divide the regression-predicted yield (with necessary interpolations described above) by the average state-level yield observed in Iowa, which is 127.96 bushels per acre of corn and 31.36 bushes per acre (see table 3) of soybeans in the one year history case. Thus, assuming farmers optimize, the typical value of the adjustment function should be approximately one. The resulting estimated adjustment functions for corn are shown in figure 1.

4.2 State-Level Revenues and Fertilizer Prices

We use a vector autoregressive model to characterize the stochastic evolution of state-level revenues per acre and fertilizer prices, which are key state variables the stochastic dynamic programming model. The revenue data come from USDA's National Agricultural Statistics

Table 2: Regressions Predicting Crop Yield Conditional on Planting History and Fertilizer Use.

	Log Corn Yield			**Log Soybean Yield**			
Crop History:	1	2	3	1	2	3	3
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Estimate/(Standard Error)						
C, CC or CCC	3.79 (0.02)	3.79 (0.03)	3.79 (0.03)	3.72 (0.02)	3.75 (0.02)	3.77 (0.02)	3.77 (0.02)
CCS			3.78 (0.03)			3.74 (0.02)	3.74 (0.02)
CS or CSC		3.80 (0.03)	3.80 (0.03)		3.67 (0.02)	3.67 (0.02)	3.67 (0.02)
S, SC or SCC	4.47 (0.03)	4.47 (0.03)	4.46 (0.03)	3.50 (0.02)	na	na	na
SCS			4.48 (0.03)			na	na
SS or SSS		na	na		3.50 (0.02)	3.50 (0.02)	3.48 (0.02)
$log(n+1)\times$C	0.10 (0.02)	0.10 (0.02)	0.10 (0.02)				$6.97^{\#}$ (10.9)
$log(n+1)^2\times$C	$14.8^{\#}$ (3.14)	$14.8^{\#}$ (3.14)	$14.8^{\#}$ (3.14)				$-1.30^{\#}$ (2.08)
$log(n+1)\times$S	0.07 (0.02)	0.07 (0.02)	0.07 (0.02)				$-6.10^{\#}$ (18.9)
$log(n+1)^2\times$C	$1.37^{\#}$ (3.62)	$1.37^{\#}$ (3.62)	$1.37^{\#}$ (3.62)				$2.47^{\#}$ (3.61)
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Sample size	2100	2100	2100	1200	1200	1200	1200
Adjusted R^2	0.81	0.81	0.81	0.82	0.84	0.84	0.84

Notes: (#) indicates numbers multiplied by 1000. Crop history indicates how many years of past plantings used in the regression. For one year, the history is either C or S; for two years, the history is CC, CS, SC or SS; for three years, the history is indicated by a three-letter sequence. Histories "SSC" and "CSS" are not present in the data. Missing rotations are not shown in the table (e.g., SSC) or are indicated by "na" if the rotation exists for corn but not for soybeans. n represents fertilizer applications, in pounds per acre, and takes on one of four distinct values: 0, 80, 160, and 240. We use $n+1$ to make the log operator well defined at n =0.

Service for the state of Iowa and are adjusted to year 2000 dollars using the consumer price index. The fertilizer price data were obtained from USDA's Economic Research Service (http://www.ers.usda.gov/data/fertilizerUse/, table 7). These data are plotted in figure 2.

Table 3: Yield predictions for different rotation histories.

History :	**1 Year**		**2 Years**		**3 Years**	
Current Crop:	S	C	S	C	S	C
<u>Rotation</u>						
S, SS or SSS	33.4	131.7	33.4	132.6	33.4*	133.7*
SSC					34.8	130.7*
SC or SCS			37.0*	131.7	37.7*	133.0
SCC					39.0*	131.1
C, CS or CSS	41.8	106.8	39.4	107.3	40.7*	106.9*
CSC					39.4	107.3
CC or CCS			43.0	106.4	43.6	106.2
CCC					42.3	106.6

Notes: The table reports average predicted yields based on the regression models reported in table 2, fixing fertilizer at 130 pounds per acre and using the average of year fixed effects. A (*) indicates a rotation history that is missing in the data and a prediction that has therefore been interpolated. Interpolation is done by combining similar non-missing rotations. For example, a SSC history is not in our data of soybean yields, but SSS is in our data. We therefore approximate the soybean yield for SSC with that for SSS, plus the difference in yields for CCC and CCS.

In tables 4 and 5 we report results from first- and second-order autoregressive models. The first table reports regression models without fertilizer prices, which are used for models with fixed fertilizer price.[5] The second table reports models with fertilizer prices. Although a second-order model would be computationally infeasible if used in our dynamic program, we report the regression results to show that earlier lags have little or no predictive power, which suggests the 1^{st}-order model is sufficient. AIC and BIC selection criteria also prefer a first-order process. We chose the first-order model for these reasons plus the fact that adding additional continuous state variables greatly increases computational expense. One hundred years of simulated revenues and fertilizer prices, based on the first-order model in table 5, are plotted in the appendix. The appendix also reports various residual plots and Kolmolgorov-Smirnoff test statistics, which fail to reject the null hypothesis that each error is distributed normal.

[5] We use models with fixed fertilizer price to explore the sensitivity of results to the number of grid points used to approximate the value function.

Figure 1: *Estimated Revenue Adjustment Functions for Corn*

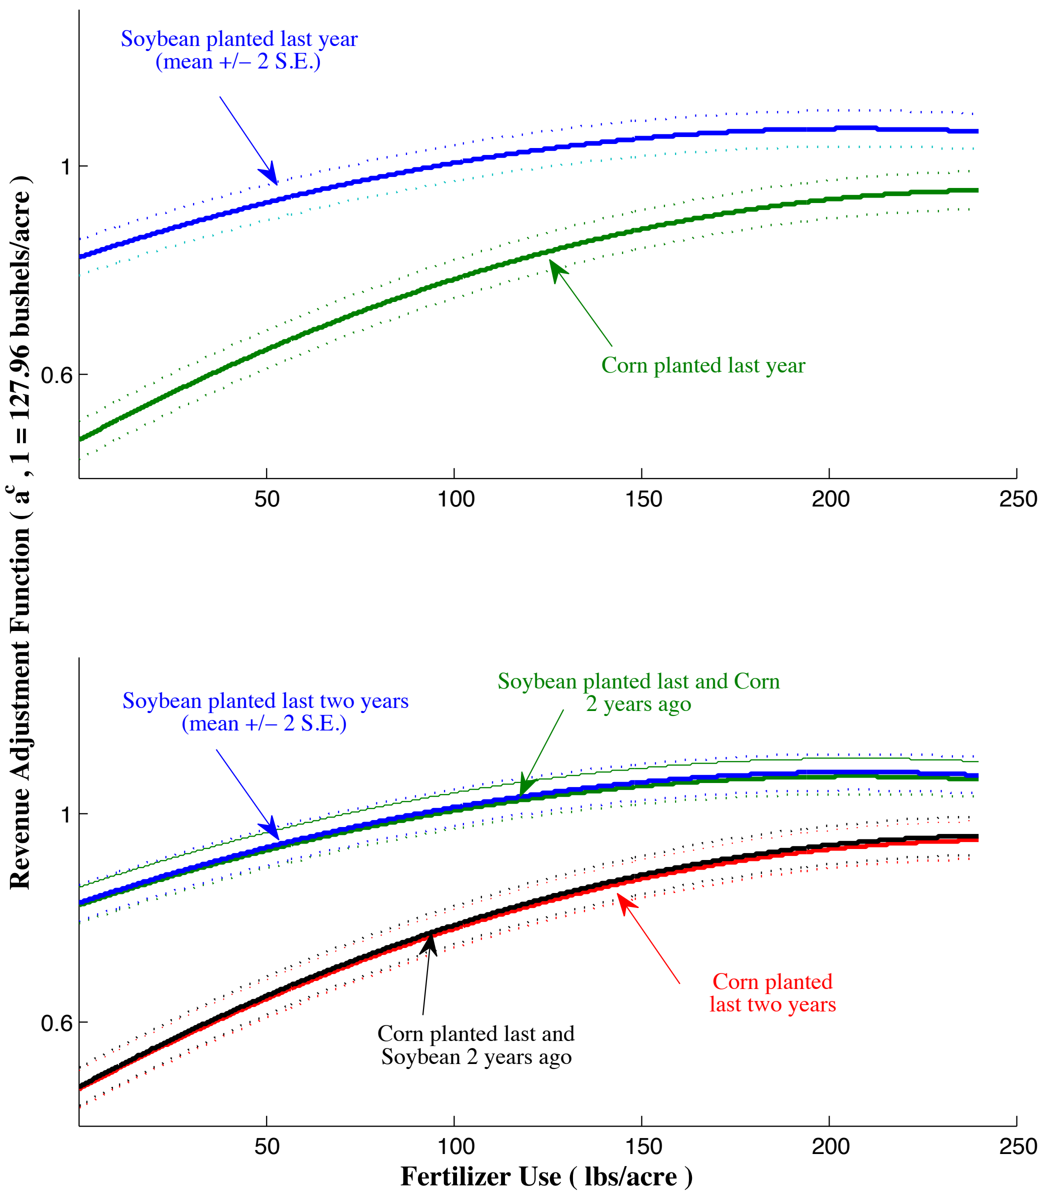

4.3 National Resources Inventory

We use data from National Resources Inventory (NRI), a survey administered by USDA's Natural Resources Conservation Service (NRCS), to obtain data on actual rotations for comparison with model predictions. The NRI is a survey that repeatedly samples approximately

Table 4: Vector Autoregression Models of State-Level Revenues

	Corn Revenue Log r_t^c (1)	Corn Revenue Log r_t^c (2)	Soybean Revenue Log r_t^s (1)	Soybean Revenue Log r_t^s (2)
	Estimates / (Standard Errors)			
Intercept	1.09	1.21	0.75	0.70
	(0.52)	(0.54)	(0.48)	(0.51)
$\log(r_{t-1}^c)$	0.48	-0.31	0.16	-0.14
	(0.19)	(0.22)	(0.17)	(0.20)
$\log(r_{t-2}^c)$		0.61		0.20
		(0.21)		(0.19)
$\log(r_{t-1}^s)$	0.35	0.22	0.71	0.22
	(0.19)	(0.22)	(0.17)	(0.21)
$\log(r_{t-2}^s)$		0.30		0.59
		(0.22)		(0.21)
Sample size	49	48	49	48
Adjusted R^2	0.66	0.71	0.71	0.71

Variance-Covariance of Innovations (AR1)

	u^c	u^s
u^c	3.61	2.34
u^s	2.34	3.13

Notes: The table reports estimates of first-order and second-order vector regressive models. The first-order coefficients (columns 1 and 3) comprise our estimate of the matrix $\mathbf{B}$ in equation 2 and the variance-covariance matrix the estimate of $\mathbf{\Omega}$, all for the case when fertilizer price is assumed fixed. The models were estimated using OLS.

a million points across the United States in order to tack parcel-specific land use change. The key advantages of the NRI in comparison to other surveys is that land units rather than farms are sampled and the land units are fixed over time. Until 1997, the survey was conducted every five years; however, on cultivated cropland, the survey also obtained crop choices for the four years prior to each sampled year. Combining surveys from 1982, 1987, 1992 and 1997 gives a panel of crop choices running from 1979 through 1997 with only the years 1983, 1988, and 1993 missing. Since 1997 the survey has become annual. Although the point data are not publicly available since 1997, we were able to analyze these data on site at USDA's Economic Research Service.

For our analysis, we limit the data set to 6,513 parcels in Iowa from 1979-2007 that

Table 5: Vector Autoregression Models of State-Level Revenues and Fertilizer Prices

	Corn Revenue Log r_t^c		Soybeans Revenue Log r_t^s		Fertilizer Price f_{t+1}	
	(1)	(2)	(3)	(4)	(5)	(6)
	Estimates/ (Standard Errors)					
Intercept	1.71	2.22	1.65	1.84	-2.21	-2.30
	(0.76)	(0.92)	(0.69)	(0.88)	(0.60)	(0.74)
$\log(r_{t-1}^c)$	0.42	-0.39	0.07	-0.23	0.30	-0.03
	(0.20)	(0.21)	(0.18)	(0.20)	(0.16)	(0.17)
$\log(r_{t-2}^c)$		0.46		0.11		0.34
		(0.22)		(0.21)		(0.17)
$\log(r_{t-1}^s)$	0.33	0.19	0.67	0.21	0.14	-0.13
	(0.19)	(0.22)	(0.17)	(0.21)	(0.21)	(0.17)
$\log(r_{t-2}^s)$		0.42		0.64		0.14
		(0.23)		(0.22)		(0.19)
$\log(n_t)$	0.11	-0.23	0.16	-0.07	0.68	-0.07
	(0.10)	(0.17)	(0.09)	(0.16)	(0.08)	(0.14)
$\log(n_{t-1})$		0.37		0.25		0.72
		(0.20)		(0.19)		(0.16)
Sample size	49	48	49	48	49	48
Adjusted R^2	0.66	0.68	0.73	0.72	0.82	0.82

Variance-Covariance Matrix of Innovations (AR1)

	u^c	u^s	u^f
u^c	3.51	2.20	0.60
u^s	2.20	2.92	-0.15
u^f	0.60	-0.15	2.21

Notes: The table reports estimates of first-order and second-order vector regressive models of fertilizer prices and state-level corn and soybean revenues per acre. The first-order coefficients (columns 1, 3 and 5) comprise the estimate of the matrix **B** in equation 2, and the variance-covariance matrix gives the estimate of $\boldsymbol{\Omega}$. The models were estimated using OLS.

planted either corn or soybeans. Observations are not available for 1983, 1988 and 1993. These data are summarized in the appendix (table 11). Looking across the years in the sample, the frequency of corn ranges 50% to 60% of sampled parcels. Corn and soybeans are planted in rotation (corn after soybean or vice versa) on 67% to 91% of the parcels. Corn is planted after corn 3% to 15%, and the frequency of soybeans after soybeans ranges from 0.1% to 1.6%.

Figure 2: *Historical Iowa Corn and Soybean Revenues and Fertilizer Prices*

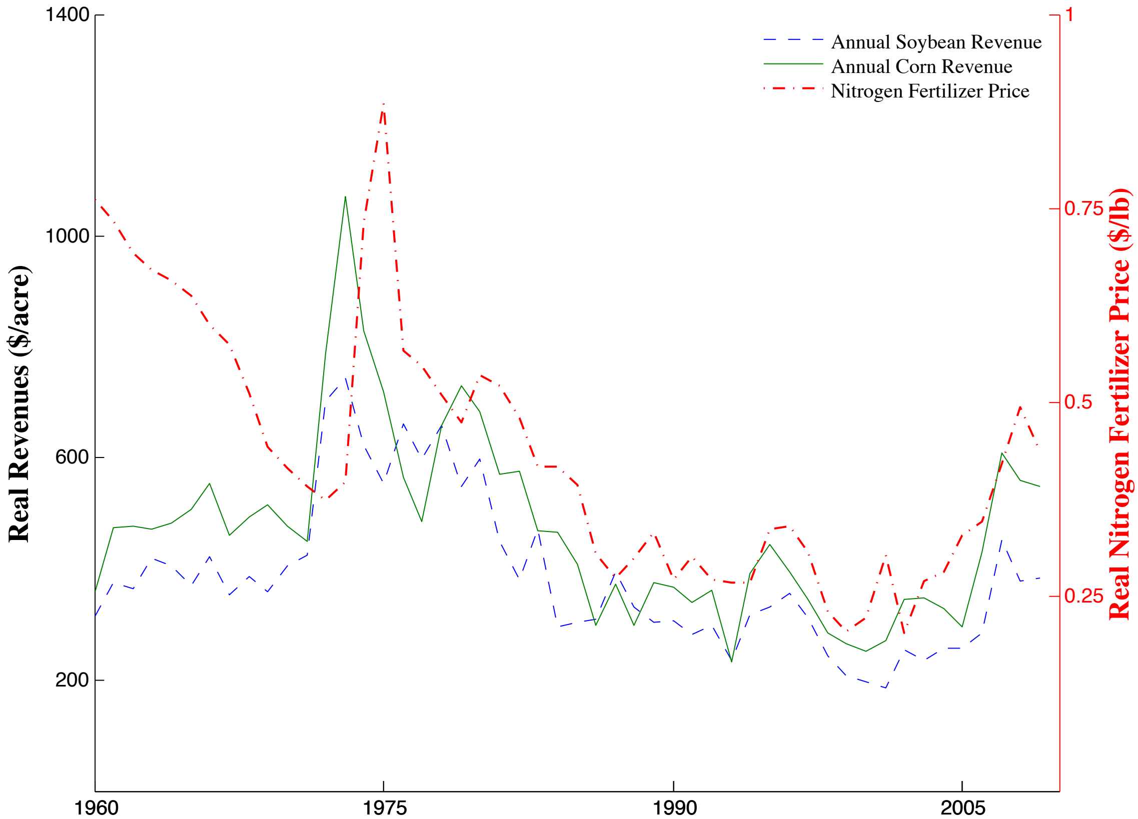

5 Results

5.1 Scenarios Considered

Here we report rotation decisions, fertilizer applications, average profits, and profit variability under a series of different modeling assumptions and decision rules based on four planning horizons:

(i) *One-Year Horizon*: A farmer that optimizes current expected profits conditional on past plantings.

(ii) *Two-Year Horizon*: A farmer that maximizes the expected sum of current and subsequent year's profits conditional on past plantings.

(iii) *Infinite Horizon*: A fully optimizing farmer that maximizes the present discounted value over an infinite horizon.

(iv) *Always Rotate*: A rule-of-thumb farmer that rotates corn after soybeans and vice-versa, regardless of prices, but applies fertilizer optimally conditional on planting decisions.

Comparing results over these different objectives allows us to evaluate the economic costs associated with less-than-optimal or rule-of-thumb decision criteria. These cost margins are an important consideration because real option values, which could be implicit in rotational decisions of forward-looking farmers, tend to be small in size even when they have a large influence on decisions. It would not be surprising to find farmers using simpler decision rules. Thus, in addition to simulated outcomes, we also examine which decision criteria best predict actual plot-level decisions.

We also consider optimal decisions of a "risk-averse" farmers that maximize the present value of log profits rather than raw profits.[6] Our goal is simply to examine whether decisions are particularly sensitive to tastes about profit variability. Because farm-level profits will sum outcomes across many parcels of land, and because consumption is likely to be much smoother than profits, maximizing the expected present value of log profit for a single acre of production implicitly assumes both risk aversion and large degree uninsurable risk.

Thirdly, we compare solutions with observed price and revenue processes with one in which corn plantings receive an additional premium of $20 acre. This premium approximately equals the premium received by farmers selling corn to nearby ethanol plants, which now consume a significant share of Iowa corn production.[7] This also gives some indication of planting response to a permanent change in relative prices of corn relative to soybeans.

We make all of these comparisons using a range of modeling assumptions. In some models we consider a three-year crop history and others only a one-year crop history. In some models the fertilizer price is fixed at $ 0.42/lb and others the fertilizer price evolves stochastically according to the vector-autoregression model reported above. We chose a fixed fertilizer price of $ 0.42/lb because it is close to the recent high prices of nitrogen and it best predicts actual planting decisions reported in the National Resources Inventory. For the model with a fixed fertilizer price, we have one less continuous state variable, which greatly reduces computation expense. We therefore use this model to examine how the granularity of the value function approximation (20 nodes instead of 10 nodes) influences the results. These different modeling strategies are used to evaluate the robustness of predicted choices

[6] We put "risk-averse" in quotes because this is a crude approximation of diminishing marginal utility of wealth or consumption.

[7] McNew & Griffith (2005) estimate a corn price premium of 12.5 cents per bushel near ethanol plants, which amounts to $20/acre if yield is expected to be 160 bushels per acre, which is typical in recent years in Iowa.

to various modeling assumptions.

After solving the infinite-horizon and shorter-horizon objectives, we simulate 100 years of the estimated stochastic generating process of prices and revenues and evaluate decisions and profits for each solved model. We then replicate this 100-year simulation 1000 times using the same pseudo-random outcomes across all models so that differences across models cannot be attributed to sampling error. For each simulation we evaluate average profits, the standard deviation of profits, the present value of profits, average fertilizer applied to corn, the average frequency of corn plantings, and the average frequency of corn after corn. For each measure, we report the average over the 1000 simulations.

5.2 Policy Functions

In figure 3 we show solved policy functions for risk-neutral objectives with fixed fertilizer prices, one-year of planting history and no corn price premium. We can only illustrate policy functions with one year of planting history and fixed fertilizer prices; other versions of the model have too many dimensions for graphical presentation. Each line gives an indifference curve for planting corn or soybeans under a given planning horizon (naive, two year, or infinite) and planting history (**C** last year or **S** last year). If a revenue pair falls below the indifference curve, the policy function indicates soybeans should be planted; if a revenue pair falls above the indifference curve, the policy function indicates corn should be planted. The points on each line indicate node locations. Dots indicate actual historical revenue pairs.

There are three notable features to the policy-function indifference curves: (1) The longer the planning horizon, the wider the region wherein it is optimal to rotate (plant corn after soybeans and vice versa), which reflects the real option values embodied in rotations described above; (2) nearly all observed revenue pairs lie between the indifference curves, regardless of the planting horizon, the one exception being the year 1971, when only the naive farmer plants corn after corn; and (3) observed historical revenue pairs generally appear closer to the corn-soybean indifference following corn as compared to following soybeans, especially for shorter planning horizons. Note that we indicate revenue pairs from recent years so they can be placed in historical context. Thus, while our model would predict the nearly identical land use choices (i.e., always rotate) for all planning horizons, the model suggests that corn after corn plantings may be more likely than soybean after soybean plantings.

In the next section considers simulations of these and all the other models.

Figure 3: Indifference Curves Between Corn and Soybean Plantings Conditional on Past State-Level Renenues

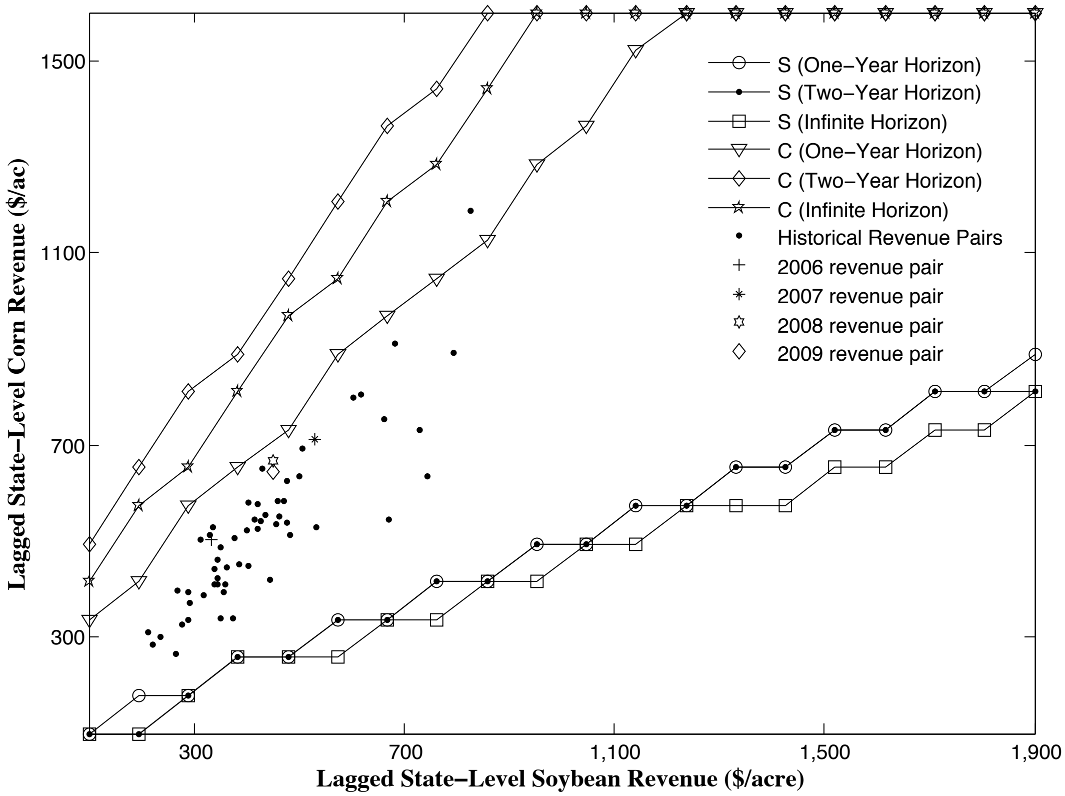

Notes: The figure illustrates solved policy functions for risk-neutral objectives with fixed fertilizer prices, one-year of planting history and no corn price premium. Each line gives an indifference curve for planting corn or soybeans under a given planning horizon (naive or infinite) and planting history (**C** last year or **S** last year). If a revenue pair falls below the indifference curve, the policy function indicates soybeans should be planted; if a revenue pair falls above the indifference curve, the policy function indicates corn should be planted. The points on each line indicate node locations. Dots indicate historical revenue pairs.

5.3 Simulation Results

Results for models with fixed fertilizer price (tables 6 and 7) show remarkably similar decisions and present values across modeling assumptions. In the great majority of states it is optimal to rotate regardless of the planning horizon, whether farmers are risk averse, or whether or not corn plantings receive a $20 per acre premium. This result is especially striking when only one year of crop history is considered (table 6). In this case, regardless of prices or state-level revenues, only naive current-year maximizing farmers ever plant corn after corn, with corn making up an average of 51.1 to 56.4% of plantings depending on the corn premium and risk aversion. Furthermore, there is little economic cost to being short-sighted or using the always-rotate rule, amounting to an average of \$1 to \$4 per acre and an average net present value loss of \$67 to \$112 per acre. Unsurprisingly, regardless of whether or not farmers receive a \$20 premium for corn, they have higher average current profits when planting corn than when planting soybeans.

Somewhat larger differences arise when a three-year crop history is considered. Here corn plantings can be as little as 48.4% (rare soybean after soybean plantings) for farmers with two-year planning horizons and up to 61.6% for risk-neutral naive farmers near an ethanol plant that receive a corn price premium. It is notable that farmers with a two-year planning horizon rotate more than infinite-horizon maximizers. While not surprising, it is also notable that more forward looking or alway-rotating objectives tend to apply less fertilizer. Results for a 2-year crop history, reported in the appendix, are similar to those for a three-year history.

Risk-averse farmers give up little in average or present value of profits relative to risk neutral farmers. Risk-averse farmers also apply less nitrogen fertilizer, but otherwise manage crop choices very much as risk neutral farmers would. Somewhat surprisingly, risk-averse farmers with a short time horizon have higher average profits than risk-neutral farmers maximizing current profits. Thus, risk aversion partially compensates for lack of foresight.

A key advantage to holding fertilizer price fixed is that the reduced dimensionality of the problem makes it feasible to consider a finer approximation of the state space by using 20 nodes in each continuous dimension rather than 10 nodes. This change gave little to no discernible difference in any the values reported. At least in this application, 10 nodes in each dimension appears to provide an adequate approximation of the state space, which gives us some confidence about the model with a stochastic fertilizer price, which was constrained to 10 nodes due to computational limitations.

The results for models with stochastic fertilizer price (tables 8 and 9) differ only slightly from those where fertilizer price is fixed. The most noticeable difference from models with fixed fertilizer price is that fertilizer applications are now greater on average and more variable. This result is not surprising given the usual convexity of profit functions in prices and the fact that other decisions change little. Although results with a stochastic fertilizer price have a lower expected present value, this difference follows mainly from subtle differences in the estimated vector autoregressive processes when fertilizer is and is not included. When fertilizer prices are included in the, mean state-level revenues are 497 and 406 dollars per acre for corn and soybeans, respectively. When fertilizer prices are excluded, the means are 503 and 411.

5.4 Comparison of Models to Observed Rotations

To test how different models are close to the real planting decisions of individual farmers, we matched a predictions conditional on historical state variables to actual rotations observed in the NRI data. For the models with fixed fertilizer price and one-year of crop history, all but one of the observations are in the rotation region of the state-space that lies between the two sets of indifference curves. The one exception (1975) is a year not observed in the NRI data, so all models predict the same planting behavior and have the same prediction accuracy (figure 4). Prediction accuracy is no better, and sometimes worse, in the model

Figure 4: Prediction Accuracy of fixed nitrogen fertilizer price models compared to NRI data

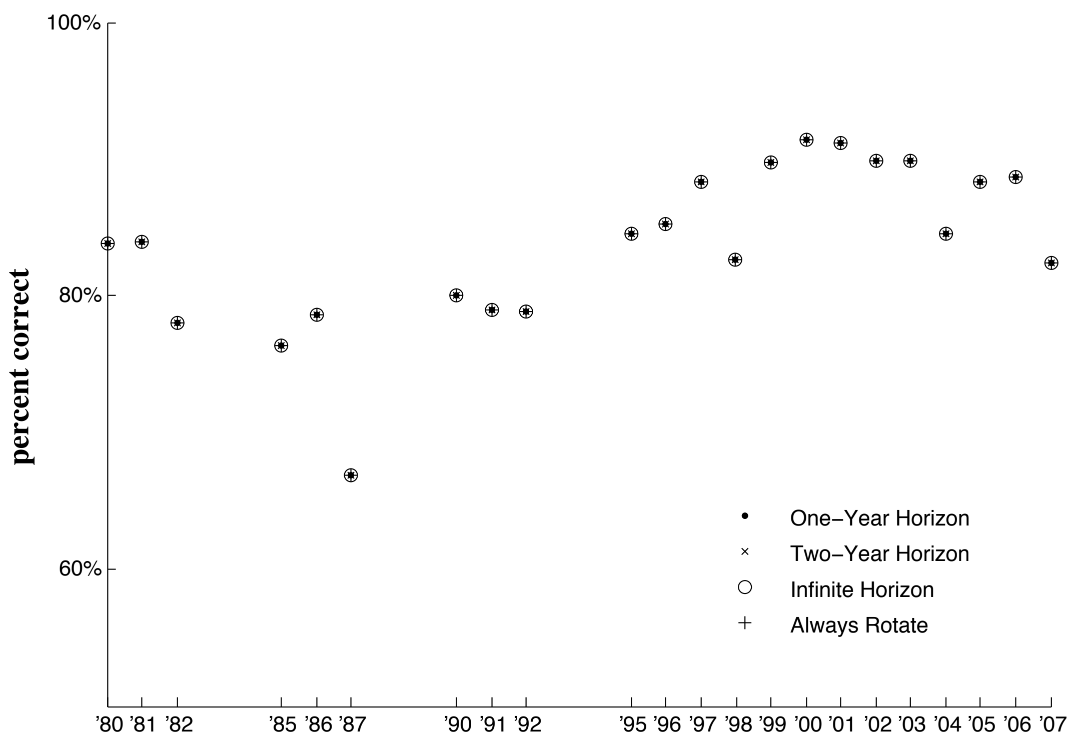

This graph summarized the prediction accuracy of different revenue maximization models given one year rotation history with no premium on corn price and fixed nitrogen fertilizer price. The prediction accuracy are almost the same for every model.

with stochastic fertilizer price (figure 5). Note that, for this field-level model, always rotate turns out to have the best possible prediction accuracy, since this is the most frequent land use choice in every year and we assume the same, field-experiment-based productivity for all parcels. Prediction accuracy for two-year and three-year histories are similar to the one-year model.

6 Conclusion

This paper develops a new dynamic model of crop planting decisions built around the agronomic benefits of crop rotations and price uncertainty. Where the effects of price uncertainty on planting decisions has long been a focus in agricultural economics, that traditional focus has been built primarily around a static model with risk aversion, wherein planting decisions

Figure 5: Prediction Accuracy of stochastic nitrogen fertilizer price models compared to NRI data

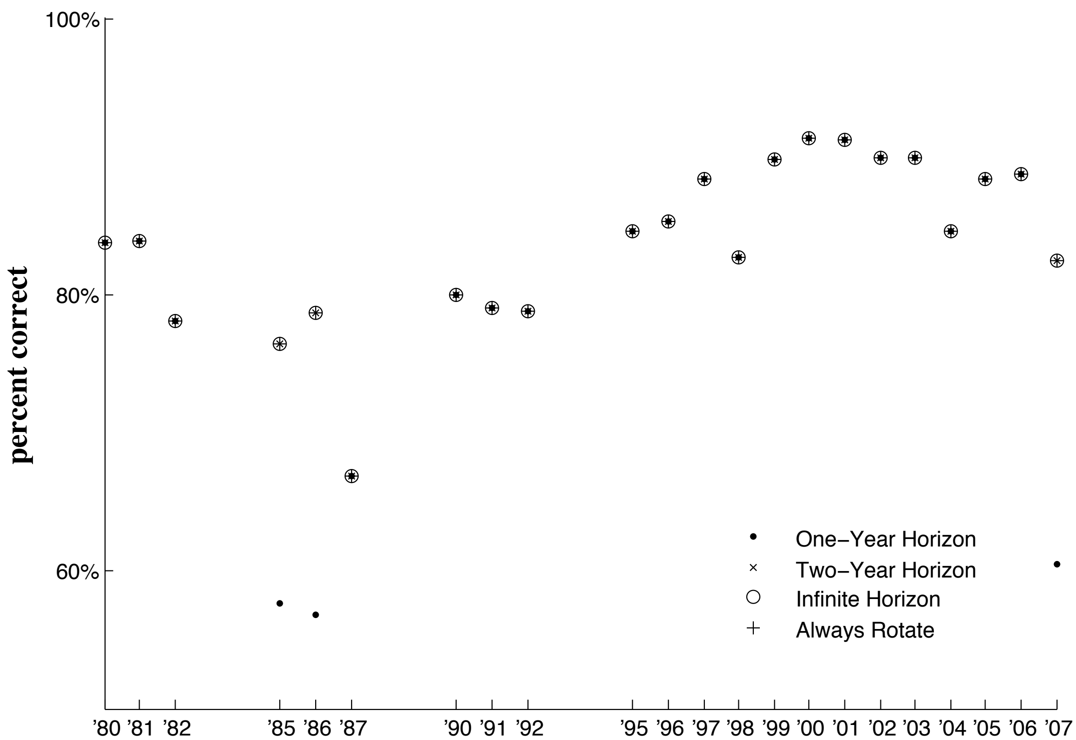

This graph summarized the prediction accuracy of different revenue maximization models given one year rotation history with no premium on corn price and stochastic fertilizer price. The prediction results are the same for two years horizon models and always rotating models.

are treated as a portfolio problem (Sandmo 1971, Feder 1980, Just & Zilberman 1983, Chavas & Holt 1990, Pope & Just 1991). In a static model with risk-aversion, a farmer plants a mix of corn and soybeans to reduce profit risk, not because they are complements in the production process. In this traditional approach, the effect of uncertainty is driven by curvature of a utility function, not by option values associated with longer-run price uncertainty. While both the static portfolio approach and stochastic-dynamic views are likely important, there has been relatively little attention paid to the latter, and this paper develops a first attempt to address that gap.

It is important to note a fundamental the difference between the portfolio/allocation and the rotation/option-value approach that we take here: where effects of risk aversion in agriculture follow mainly from imperfect insurance markets, real option values exist even in perfect markets. A corollary of this observation is that an absence of monoculture does not constitute prima facie evidence of market failure.

Calibration of our model using data from experimental field trials as well as historical price and yield data from Iowa indicates a powerful incentive to rotate. Even in locations near ethanol plants, which are assumed to receive a $20 per acre premium for producing corn, the model shows farmers rotating nearly as often as farmers without such a premium. Indeed, always rotating with or without the premium is imperceptibly different from the optimal decision rule. This suggests near-zero supply response to both temporary (stochastically evolving) and permanent price shocks.

Still, the market does not harshly punish suboptimal actors. Naive farmers, which have planting decisions that are more responsive to temporary and permanent price changes, lose little for their suboptimal behavior. The reason is that the value function is nearly flat with respect to planting decisions on the margin of corn-after-corn versus corn-after-soybeans when corn and soybean prices are generally high. These observations, combined with the fact that Iowa has a high concentration of corn and soybean production, suggests it may be dubious to use our model for large-scale prediction of supply response. Because a range of planting decisions are plausible (i.e., optimal or nearly so), and decisions aggregated to the state level would feed back to influence prices, the shape of the demand curve would likely pin down realized supply response. We leave for future work consideration of an equilibrium model to account for price feedback.

The flat value function near the corn-corn, corn-soybean margin at high price levels suggests risk aversion *might* influence planting decisions in a multiple-field setting. Because diversification would presumably be maximized or nearly so with an even split of corn and soybean plantings, the always-rotating rule-of-thumb would seem to nearly maximize both average profits and diversification, so long as the initial condition of the land was half corn and half soybeans.[8] Nevertheless, it may be interesting for future work to explicitly consider the portfolio problem simultaneously with rotations.

Our approach to modeling crop planting decisions could otherwise be extended to consider a number of potential policy implications. For example, subsidized crop insurance, by limiting downside risk, could discourage rotation in favor of monoculture, a potentially acute form of dynamic moral hazard (Vercammen & van Kooten 1994). By discouraging rotations, insurance might also encourage greater chemical use and lead to worse environmental outcomes. Horowitz & Lichtenberg (1993) found insurance increased chemical use if inputs were "risk-increasing." In our view, a more natural explanation of this phenomenon is substitutability of inputs (fertilizer and/or pesticides) with rotations that insurance might discourage. There may also be implications for introduction of genetically modified seed with pest resistance, like Bt corn. Given adoption of genetically modified crops has grown simultaneously with corn ethanol production, and corn expansion has come mainly at the expense of reduced soybean plantings, future research might investigate the degree to which these new seed varieties substitute for rotation benefits.

References

Burt, O.R., and J.R. Allison. 1963. "Farm Management Decisions With Dynamic Programming." *Journal Of Farm Economics*, 45(1): 121–136.

Burt, O.R., and V.E. Worthington. 1988. "Wheat Acreage Supply Response In The United States." *Western Journal Of Agricultural Economics*, 13(1): 100–111.

Cafiero, C., H. Bobenrieth, SA Eugenio, RA Juan, and B.D. Wright. 2009. "The Empirical Relevance Of The Competitive Storage Model." *Journal Of Econometrics*, 162(1): 44–54.

Chavas, J.P., and M.T. Holt. 1990. "Acreage Decisions Under Risk: The Case Of Corn And Soybeans." *American Journal Of Agricultural Economics*, 72(3): 529–538.

Deaton, A., and G. Laroque. 1992. "On The Behaviour Of Commodity Prices." *The Review Of Economic Studies*, 59(1): 123.

[8] The risk-minimizing portfolio would depend on the variances and covariances of corn and soybean returns.

Deaton, A., and G. Laroque. 1996. "Competitive Storage And Commodity Price Dynamics." *Journal Of Political Economy*, 104(5): 896–923.

Eckstein, Z. 1984. "A Rational Expectations Model Of Agricultural Supply." *The Journal Of Political Economy*, 92(1): 1–19.

El-Nazer, T., and B.A. McCarl. 1986. "The Choice Of Crop Rotation: A Modeling Approach And Case Study." *American Journal Of Agricultural Economics*, 68(1): 127–136.

Feder, G. 1980. "Farm Size, Risk Aversion And The Adoption Of New Technology Under Uncertainty." *Oxford Economic Papers*, 32(2): 263–283.

Hennessy, D.A. 2006. "On Monoculture And The Structure Of Crop Rotations." *American Journal Of Agricultural Economics*, 88(4): 900914.

Horowitz, J.K., and E. Lichtenberg. 1993. "Insurance, Moral Hazard, And Chemical Use In Agriculture." *American Journal Of Agricultural Economics*, 75(4): 926–935.

Johansson, R., M. Peters, and R. House. 2007. "Regional Environment And Agriculture Programming Model (REAP)."

Just, R.E., and D. Zilberman. 1983. "Stochastic Structure, Farm Size And Technology Adoption In Developing Agriculture." *Oxford Economic Papers*, 35(2): 307–328.

Kennedy, J.O.S. 1986. *Dynamic Programming Applications To Agriculture And Natural Resources.* Elsevier Science Pub. Co., Inc.,New York, NY.

Mallarino, A.P., E. Ortiz-Torres, and K. Pecinovsky. 2005. "Effects Of Crop Rotation And Nitrogen Fertilization On Crop Production." Iowa State University.

McNew, K., and D. Griffith. 2005. "Measuring The Impact Of Ethanol Plants On Local Grain Prices." *Applied Economic Perspectives And Policy*, 27(2): 164–180.

Miranda, M.J., and P.L. Fackler. 2004. *Applied Computational Economics And Finance.* Vol. 1, The MIT Press.

Musser, W.N., V.J. Alexander, B.V. Tew, and D.A. Smittle. 1985. "A Mathematical Programming Model For Vegetable Rotations." *Southern Journal Of Agricultural Economics*, 17(01): 169–176.

Pope, R.D., and R.E. Just. 1991. "On Testing The Structure Of Risk Preferences In Agricultural Supply Analysis." *American Journal Of Agricultural Economics*, 73(3): 743–748.

Sandmo, A. 1971. "On The Theory Of The Competitive Firm Under Price Uncertainty." *The American Economic Review*, 61(1): 65–73.

Vercammen, J., and G.C. van Kooten. 1994. "Moral Hazard Cycles In Individual-Coverage Crop Insurance." *American Journal Of Agricultural Economics*, 76(2): 250–261.

Williams, J.C., and B. Wright. 1991. *Storage And Commodity Markets.* Cambridge Univ Pr.

Table 6: Simulation Results for Fixed Fertilizer Price and One-Year Crop History

Planning Horizon	Present Value ($/ac)	Corn (%)	Double Corn (%)	Annual Profit ($/ac)	Corn Profit ($/ac)	Soybean Profit ($/ac)	Fertilizer Mean (lb/ac)	Fertilizer SD (lb/ac)
			risk neutral[a], no premium on corn					
One-Year	8826	52	4.5	442	465	416	135	31
Two-Year	8893	50	0	443	464	423	129	22
Infinite	8893	50	0	443	464	423	129	22
Always Rotate	8893	50	0	443	464	423	129	22
			risk neutral, premium on corn is $20					
One Year	8973	56.4	13.5	450	477	414	146	36
Two Years	9081	50	0	454	485	423	131	21
Infinite	9083	50	0.2	454	485	422	131	21
Always Rotate	9081	50	0	454	485	423	131	21
			risk averse, no premium on corn					
One Year	8880	51.1	2.7	443	465	419	129	29
Two Years	8892	50	0	443	464	423	126	23
Infinite	8893	50	0.1	443	464	423	126	23
Always Rotate	8892	50	0	443	464	423	126	23
			risk averse, premium on corn is $20					
One Year	8980	55.8	12.2	450	480	412	142	37
Two Years	9081	50	0	454	485	423	128	22
Infinite	9082	50	0.2	454	486	422	128	22
Always Rotate	9081	50	0	454	485	423	128	22

Notes: The table summarizes results from 1,000 replicates of 100-year simulations, wherein a one-acre field is managed using each of four different management strategies: (1) Maximization of current profits conditional on the state variables; (2) Maximization of the current plus discounted subsequent year's profits; (3) Maximization of the expected net present value using the solution to the full, infinite-horizon SDP with discount factor $\beta = 0.95$; (4) Maximization of current profits subject to always rotating. All values are means taken across of the 1,000 series. Risk-neutral farmers maximize the expected present value of profits and risk-averse farmers maximize the expected present value of log profits. The $20 premium accords with evidence on the price premium that corn receives near ethanol plants. All simulations assume the same stochastic evolution of state variables (the first-order autoregressive model reported in table 4) and use same sequence of pseudo random draws in order to minimize differences stemming from chance error.

Table 7: Simulation Results for Fixed Fertilizer Price and Three-Year Crop History

Planning Horizon	Present Value ($/ac)	Corn (%)	Double Corn (%)	Annual Profit ($/ac)	Corn Profit ($/ac)	Soybean Profit ($/ac)	Fertilizer Mean (lb/ac)	Fertilizer SD (lb/ac)
			risk neutral, no premium on corn					
One Year	8691	54.9	10.4	434	465	395	142	36
Two Years	8695	49.9	0.1	434	470	398	130	22
Infinite	8697	52.3	5.1	434	468	397	133	32
Always Rotate	8695	50	0	434	469	398	130	22
			risk neutral, premium on corn is $20					
One Year	8868	62.2	25	443	466	405	156	39
Two Years	8888	49.9	0.2	444	491	398	132	21
Infinite	8905	54.8	10.2	445	484	397	140	36
Always Rotate	8885	50	0	445	490	398	131	21
			risk averse, no premium on corn					
One Year	8691	54.9	10.4	434	465	395	139	37
Two Years	8696	49.9	0.2	434	470	398	127	23
Infinite	8697	52.3	5.1	434	468	397	133	32
Always Rotate	8695	50	0	434	469	398	126	23
			risk averse, premium on corn is $20					
One Year	8873	61.5	23.7	443	467	405	152	40
Two Years	8891	50.1	0.5	445	491	398	129	23
Infinite	8905	54.8	10.1	445	484	397	140	36
Always Rotate	8885	50	0	444	490	398	128	22

Notes: The table summarizes results from 1,000 replicates of 100-year simulations, wherein a one-acre field is managed using each of four different management strategies: (1) Maximization of current profits conditional on the state variables; (2) Maximization of the current plus discounted subsequent year's profits; (3) Maximization of the expected net present value using the solution to the full, infinite-horizon SDP with discount factor $\beta = 0.95$; (4) Maximization of current profits subject to always rotating. All values are means taken across of the 1,000 series. Risk-neutral farmers maximize the expected present value of profits and risk-averse farmers maximize the expected present value of log profits. The $20 premium accords with evidence on the price premium that corn receives near ethanol plants. All simulations assume the same stochastic evolution of state variables (the first-order autoregressive model reported in table 4) and use same sequence of pseudo random draws in order to minimize differences stemming from chance error.

Table 8: Simulation Results for Stochastic Fertilizer Price and One-Year Crop History

Planning Horizon	Present Value ($/ac)	Corn (%)	Double Corn (%)	Annual Profit ($/ac)	Corn Profit ($/ac)	Soybean Profit ($/ac)	Fertilizer Mean (lb/ac)	Fertilizer SD (lb/ac)
			risk neutral, no premium on corn					
One Year	8676	53.2	6.9	432	443	419	139	27
Two Years	8737	50.1	0.2	435	453	416	131	16
Infinite	8736	50.2	0.7	434	452	417	131	18
Always Rotate	8737	50	0	435	454	416	130	16
			risk neutral, $20 premium on corn					
One Year	8798	62.9	26.5	437	428	452	159	38
Two Years	8924	50.2	0.7	445	473	416	133	18
Infinite	8924	50.3	0.9	445	472	417	133	18
Always Rotate	8925	50	0	445	474	416	132	15
			risk averse, no premium on corn					
One Year	8692	51.1	2.8	433	451	415	131	23
Two Years	8735	50.2	0.6	434	452	416	128	18
Infinite	8735	50.2	0.7	454	452	416	128	18
Always Rotate	8737	50	0	435	454	416	127	16
			risk averse, $20 premium on corn					
One Year	8848	55.3	11.3	441	457	422	142	32
Two Years	8923	50.3	0.9	445	472	417	130	19
Infinite	8923	50.3	0.9	445	472	417	130	19
Always Rotate	8925	50	0	445	474	416	129	16

Notes: The table summarizes results from 1,000 replicates of 100-year simulations, wherein a one-acre field is managed using each of four different management strategies: (1) Maximization of current profits conditional on the state variables; (2) Maximization of the current plus discounted subsequent year's profits; (3) Maximization of the expected net present value using the solution to the full, infinite-horizon SDP with discount factor $\beta = 0.95$; (4) Maximization of current profits subject to always rotating. All values are means taken across of the 1,000 series. Risk-neutral farmers maximize the expected present value of profits and risk-averse farmers maximize the expected present value of log profits. The $20 premium accords with evidence on the price premium that corn receives near ethanol plants. All simulations assume the same stochastic evolution of state variables (the first-order autoregressive model reported in table 5) and use same sequence of pseudo random draws in order to minimize differences stemming from chance error.

Table 9: Simulation Results for Stochastic Fertilizer Price and Three-Year Crop History

Planning Horizon	Present Value ($/ac)	Corn (%)	Double Corn (%)	Annual Profit ($/ac)	Corn Profit ($/ac)	Soybean Profit ($/ac)	Fertilizer Mean (lb/ac)	Fertilizer SD (lb/ac)
			risk neutral, no premium on corn					
One Year	8496	55.7	12.2	423	442	400	145	32
Two Years	8535	49.6	0.7	425	457	394	132	18
Infinite	8520	52.3	5.5	425	450	397	138	26
Always Rotate	8542	50	0	425	459	392	131	15
			risk neutral, $20 premium on corn					
One Year	8680	61.4	23.4	433	444	416	156	36
Two Years	8729	50.1	0.9	436	477	394	134	18
Infinite	8716	54.9	10.7	435	463	400	145	31
Always Rotate	8732	50	0	436	480	392	133	15
			risk averse, no premium on corn					
One Year	8509	54.5	9.6	424	445	399	140	31
Two Years	8535	49.6	0.7	425	457	394	129	18
Infinite	8512	53.5	7.9	424	445	399	138	30
Always Rotate	8541	50	0	425	459	392	128	16
			risk averse, $20 premium on corn					
One Year	8682	61.2	23.1	433	444	416	153	36
Two Years	8728	50.1	0.9	436	477	394	131	19
Infinite	8713	55	10.8	435	463	400	142	31
Always Rotate	8731	50	0	436	490	392	130	16

Notes: The table summarizes results from 1,000 replicates of 100-year simulations, wherein a one-acre field is managed using each of four different management strategies: (1) Maximization of current profits conditional on the state variables; (2) Maximization of the current plus discounted subsequent year's profits; (3) Maximization of the expected net present value using the solution to the full, infinite-horizon SDP with discount factor $\beta = 0.95$; (4) Maximization of current profits subject to always rotating. All values are means taken across of the 1,000 series. Risk-neutral farmers maximize the expected present value of profits and risk-averse farmers maximize the expected present value of log profits. The $20 premium accords with evidence on the price premium that corn receives near ethanol plants. All simulations assume the same stochastic evolution of state variables (the first-order autoregressive model reported in table 5) and use same sequence of pseudo random draws in order to minimize differences stemming from chance error.

Appendices

A Comparison of State-Level Prices and State-Level Revenues

Our model uses state-level revenues (price × yield) instead of prices. There are at least two key benefits to this approach: (1) it accounts for covariance between price outcomes and yield outcomes, which affect both expected profits and profit variances; (2) state-level revenues appear stationary where yields and (possibly) prices are non-stationary. Here we show the relationship between state-level revenues and prices.

Figure 6: Iowa State-Level Prices Received and Revenues Per Acre

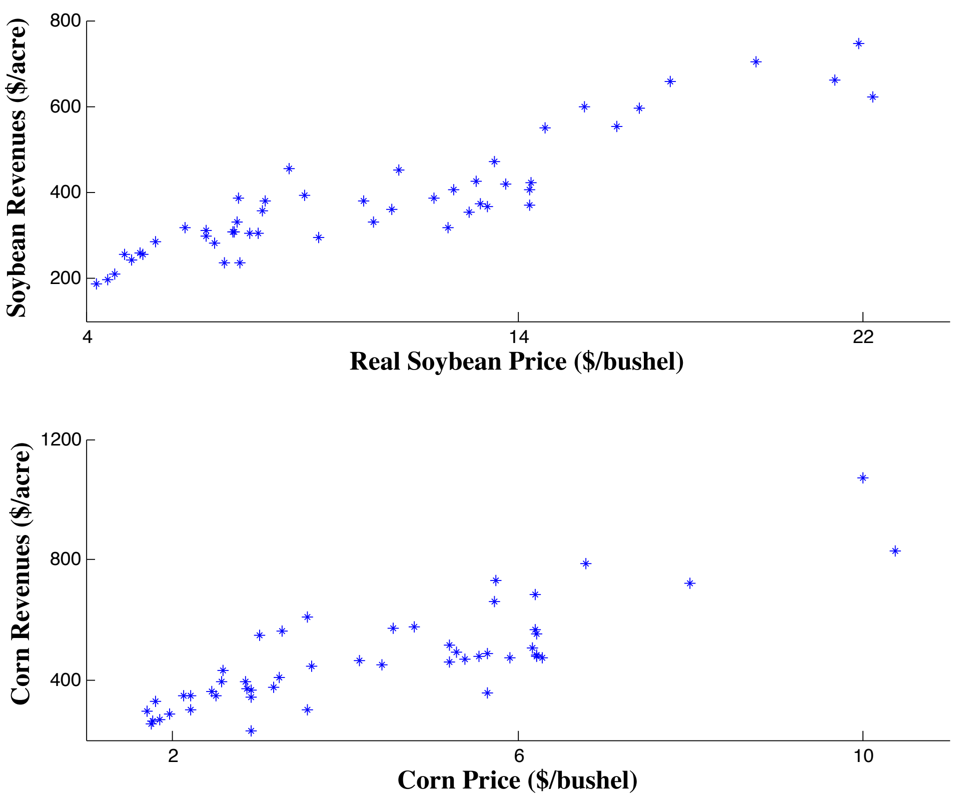

B Simulation of VAR Models

Figure 7: One-Hundred Years of Simulated Corn and Soybean Revenues and Fertilizer Prices

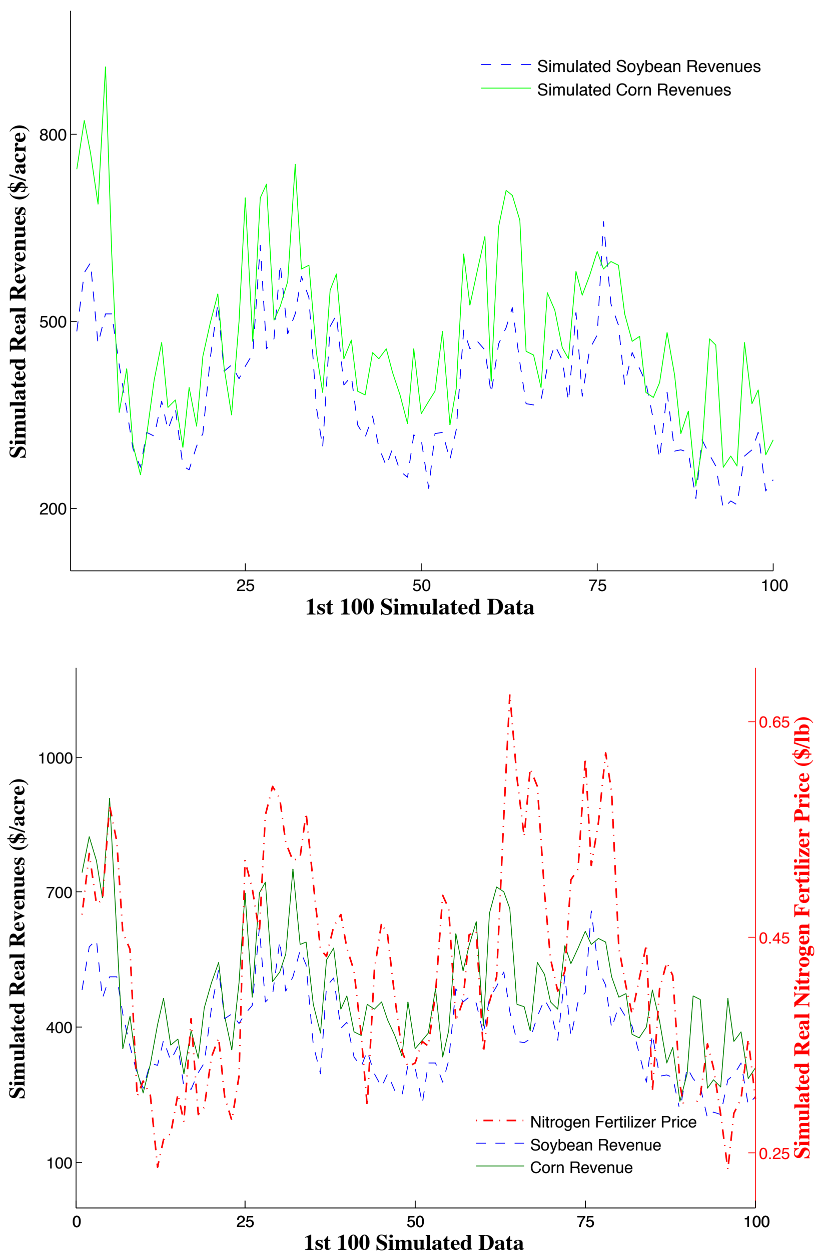

A single 100-year simulation of the estimated VAR processes without (top) and with (bottom) fertilizer prices.

C Analysis of VAR Model Residuals

Figure 8: Quantile-Quantile Plots of Residuals from VAR Model Without Fertilizer Prices

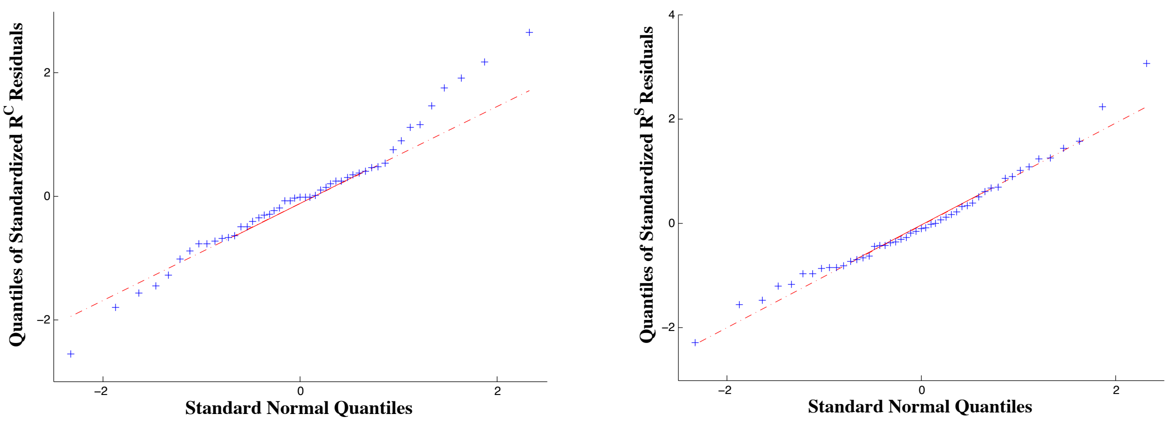

The figure plots quantiles of a standard normal distribution against standardized residuals from estimation of equation 2 when fertilizer prices are *ex*clouded. Because the regression is run in logs, a normal distributed error implies a revenue distribution that is positively skewed.

Table 10: Kolmogorov-Smirnov Tests for Normality

Variable	KS Statistic	P-Value
VAR model without fertilizer price		
Soybean Revenue	0.07	0.93
Corn Revenue	0.11	0.52
VAR model with fertilizer price		
Soybean Revenue	0.13	0.38
Corn Revenue	0.13	0.33
Fertilizer Price	0.06	0.98

Notes: The table reports two-sided Kolmogorov-Smirnov tests against the null hypothesis that the standardized residuals in the estimated VAR models are distributed Normal. The test statistic if the largest absolute difference between the empirical distribution and the standard normal distribution.

Figure 9: Quantile-Quantile Plots of Residuals from VAR Model With Fertilizer Prices

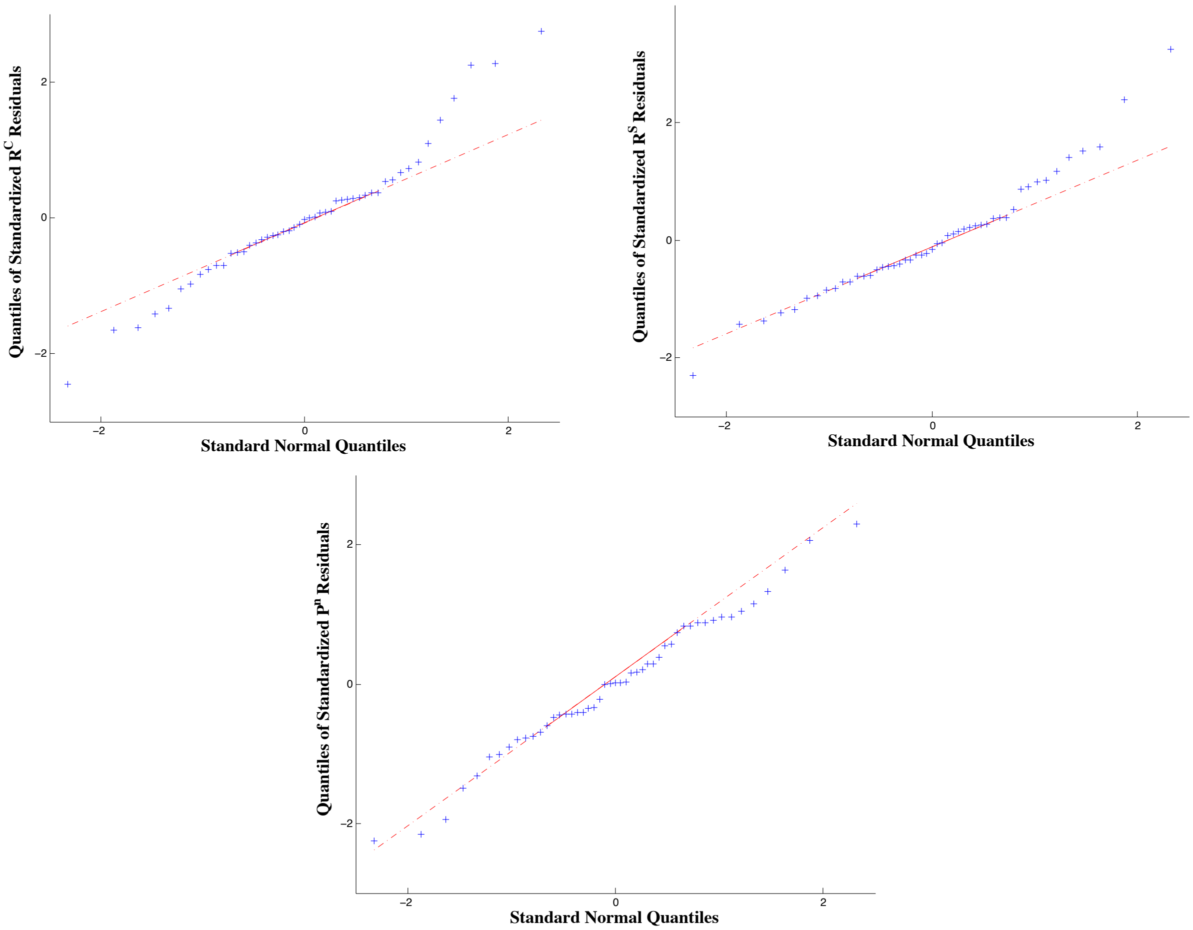

The figure plots quantiles of a standard normal distribution against standardized residuals from estimation of equation 2 when fertilizer prices are *in*clouded. Because the regression is run in logs, a normal distributed error implies a revenue distribution that is positively skewed.

Figure 10: Residual Paris from VAR Model With Fertilizer Prices

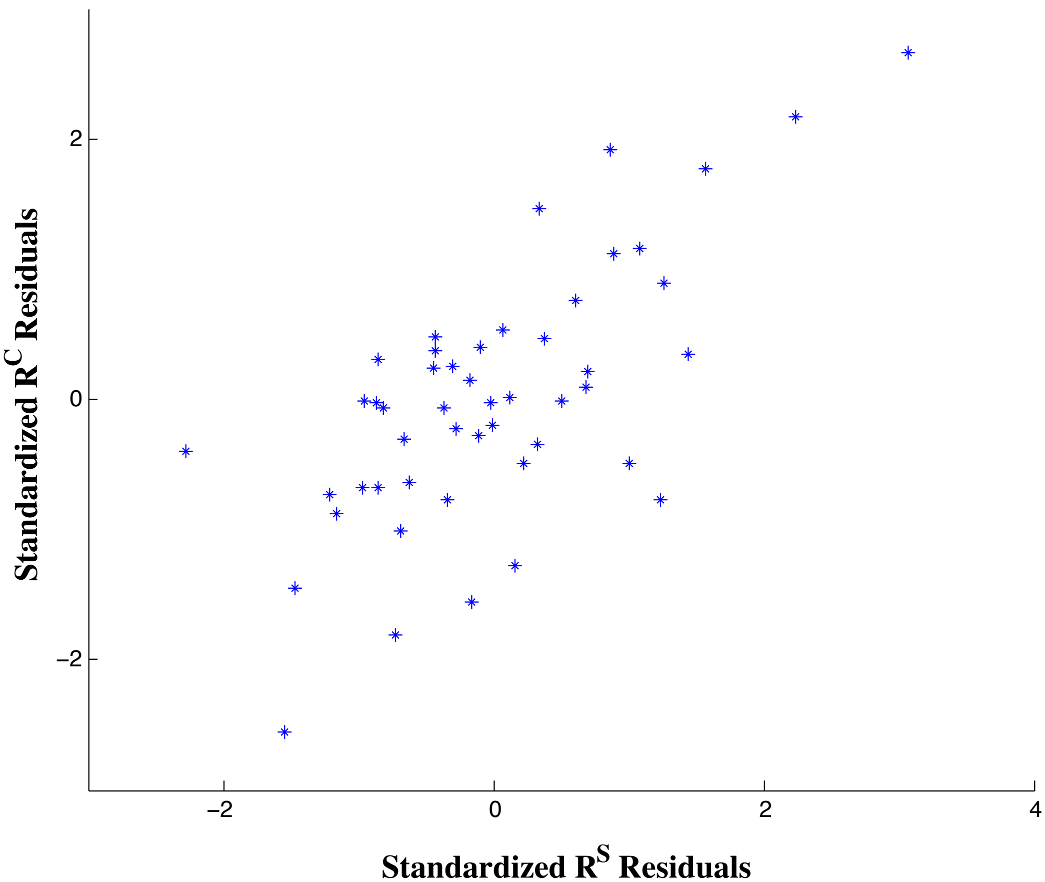

The figure shows scatterplots of the residuals from the VAR model without fertilizer prices. The associations are captured by the estimated error covariance matrix reported in the main paper.

Figure 11: Residual Paris from VAR Model With Fertilizer Prices

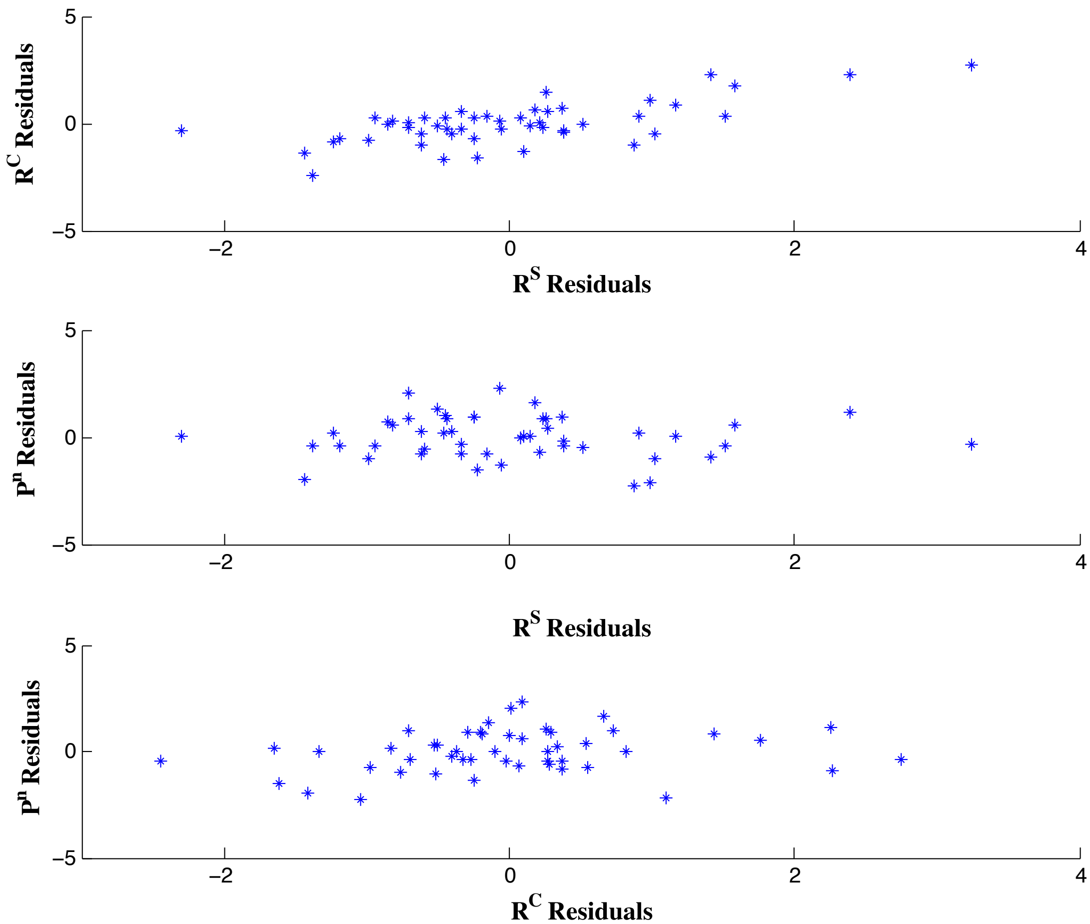

The figure shows scatterplots of the residuals from the VAR model with fertilizer prices. The associations are captured by the estimated error covariance matrix reported in the main paper.

D Summary Statistics and Predictions for Actual Rotations

Table 11: Rotation Frequencies for Iowa Fields in the National Resources Inventory

	(**C** ≡ Corn and **S** ≡ Soybeans)							
	C-C-C	**C-C-S**	**C-S-C**	**C-S-S**	**S-C-C**	**S-C-S**	**S-S-C**	**S-S-S**
	(Percent of Fields)							
1981	13.1	1.9	41.0	0.8	2.0	0.8	40.2	0.3
1982	13.0	6.6	40.1	0.8	2.0	2.1	35.1	0.3
1986	14.5	3.4	36.8	2.1	6.1	2.6	33.6	0.9
1987	12.0	9.8	30.1	1.9	5.9	9.7	29.1	1.6
1991	11.6	5.7	36.2	2.0	5.6	2.8	35.2	0.8
1992	11.7	5.9	38.1	2.8	5.7	2.7	32.3	0.9
1996	7.8	4.5	41.8	1.9	4.8	1.6	36.8	0.9
1997	6.1	3.8	40.5	1.8	6.2	1.1	39.9	0.7
1998	4.6	7.8	41.7	1.3	5.3	4.4	34.5	0.5
1999	4.7	3.5	38.7	4.0	7.8	1.1	39.4	1.0
2000	4.0	2.5	45.7	1.3	4.2	1.5	40.1	0.7
2001	2.6	3.3	41.8	1.7	3.9	2.5	43.8	0.5
2002	2.7	3.7	44.6	2.4	3.1	3.1	39.8	0.6
2003	2.8	5.0	41.3	3.2	3.6	1.7	41.9	0.6
2004	3.8	8.2	42.8	1.4	4.0	2.7	36.3	0.9
2005	6.0	5.1	39.7	3.5	6.0	0.6	39.2	0.1
2006	6.3	4.3	44.7	0.3	4.8	0.4	38.9	0.3
2007	7.6	9.2	43.0	0.7	2.9	0.7	35.9	0.1

Notes: Each row sums to 100%. The table reports frequencies for all 6,513 sampled Iowa fields planted with either corn or soybeans in all years. The first letter represents the crop choice of current year; the middle letter represents the crop choice from the previous year; and the last letter represents the crop choice two years prior.

E Results for Two-Year Crop Histories

The main paper reports results for one-year and three-year histories. Results for two-year histories, which are similar to three-year histories, are reported here.

Table 12: Simulation Results for Fixed Fertilizer Price and Two-Year Crop History

Planning Horizon	Present Value ($/ac)	Corn (%)	Double Corn (%)	Annual Profit ($/ac)	Corn Profit ($/ac)	Soybean Profit ($/ac)	Fertilizer Mean (lb/ac)	Fertilizer SD (lb/ac)
				risk neutral, no premium on corn				
One Year	8689	54.9	10.4	433	462	397	142	37
Two Years	8640	48.4	0.6	431	469	396	140	23
Infinite	8690	54.6	9.8	433	462	398	141	36
Always Rotate	8645	50	0	431	464	398	129	22
				risk neutral, premium on corn is $20				
One Year	8896	61.6	23.8	444	466	408	155	39
Two Years	8831	48.7	0.8	441	490	395	133	22
Infinite	8904	58.8	18.3	444	468	411	150	27
Always Rotate	8833	50	0	442	485	398	131	21
				risk averse, no premium on corn				
One Year	8690	54.8	10.2	433	462	397	138	37
Two Years	8643	48.4	0.8	431	468	396	128	24
Infinite	8690	54.6	9.9	433	462	398	138	37
Always Rotate	8645	50	0	431	464	398	126	23
				risk averse, premium on corn is $20				
One Year	8898	61.4	23.5	444	466	409	152	40
Two Years	8834	48.9	1.2	441	489	396	130	24
Infinite	8904	58.9	18.3	444	467	411	147	37
Always Rotate	10371	50	0	433	478	389	127	22

Notes: The table summarizes results from 1,000 replicates of 100-year simulations, wherein a one-acre field is managed using each of four different management strategies: (1) Maximization of current profits conditional on the state variables; (2) Maximization of the current plus discounted subsequent year's profits; (3) Maximization of the expected net present value using the solution to the full, infinite-horizon SDP with discount factor $\beta = 0.95$; (4) Maximization of current profits subject to always rotating. All values are means taken across of the 1,000 series. Risk-neutral farmers maximize the expected present value of profits and risk-averse farmers maximize the expected present value of log profits. The $20 premium accords with evidence on the price premium that corn receives near ethanol plants. All simulations assume the same stochastic evolution of state variables (the first-order autoregressive model reported in table 5) and use same sequence of pseudo random draws in order to minimize differences stemming from chance error.

Table 13: Simulation Results for Stochastic Fertilizer Price and Two-Year Crop History

Planning Horizon	Present Value ($/ac)	Corn (%)	Double Corn (%)	Annual Profit ($/ac)	Corn Profit ($/ac)	Soybean Profit ($/ac)	Fertilizer Mean (lb/ac)	Fertilizer SD (lb/ac)
risk neutral, no premium on corn								
One Year	8490	55.5	11.7	423	440	401	144	31
Two Years	8477	49.5	1.2	422	452	393	132	19
Infinite	8492	55.4	11.7	423	440	401	144	31
Always Rotate	8492	50	0	423	454	392	130	16
risk neutral, $20 premium on corn								
One Year	8678	61.7	24.1	433	442	418	157	36
Two Years	8674	50.5	2.6	433	471	394	135	22
Infinite	8695	58.8	18.4	434	449	412	152	35
Always Rotate	8680	50	0	433	474	392	132	15
risk averse, no premium on corn								
One Year	8495	54.4	9.6	423	442	400	139	31
Two Years	8479	49.7	1.2	422	452	393	129	20
Infinite	8495	54.4	9.6	423	442	400	139	31
Always Rotate	8492	50	0	423	454	392	127	16
risk averse, $20 premium on corn								
One Year	8678	61.6	23.7	432	442	418	153	37
Two Years	8679	52.1	5.5	433	465	399	136	36
Infinite	8695	58.8	18.4	434	449	412	138	35
Always Rotate	8680	50	0	433	474	392	129	16

Notes: The table summarizes results from 1,000 replicates of 100-year simulations, wherein a one-acre field is managed using each of four different management strategies: (1) Maximization of current profits conditional on the state variables; (2) Maximization of the current plus discounted subsequent year's profits; (3) Maximization of the expected net present value using the solution to the full, infinite-horizon SDP with discount factor $\beta = 0.95$; (4) Maximization of current profits subject to always rotating. All values are means taken across of the 1,000 series. Risk-neutral farmers maximize the expected present value of profits and risk-averse farmers maximize the expected present value of log profits. The $20 premium accords with evidence on the price premium that corn receives near ethanol plants. All simulations assume the same stochastic evolution of state variables (the first-order autoregressive model reported in table 5) and use same sequence of pseudo random draws in order to minimize differences stemming from chance error.