Estimating Heterogeneous Corn Price Response in the United States

Nicholas J. Pates[1] and Nathan P. Hendricks[2]

Selected Paper prepared for presentation at the 2018 Agricultural & Applied Economics Association Annual Meeting, Washington, D.C., August 5-August 7

[1] PhD. Candidate, Department of Agricultural Economics, Kansas State University, Manhattan, KS 66502 njpates@ksu.edu

[2] Assistant Professor, Department of Agricultural Economics, Kansas State University, Manhattan, KS 66502, nph@ksu.edu

Estimating Heterogeneous Corn Price Response in the United States

Nicholas J. Pates and Nathan P. Hendricks

1 Introduction

As individuals, farmers often face different constraints, have different goals, and have different preferences to practices and risk and these differences can inform their production choices. This decision-making heterogeneity can impact farm-response on the aggregate making heterogeneity critically important for designing better-targeted policies and predicting outcomes of existing policy. Acreage response could also be spatially correlated with a host of policy relevant details including, environmentally sensitive areas and areas differing by potential yield and crop substitutability. With greater availability of rich, high density, and expansive geospatial agricultural datasets, incorporating spatial heterogeneity in production agriculture is becoming more feasible. In this paper, we estimate a series of corn supply response models over the contiguous United States (CONUS) between 2004 and 2016 while accounting for heterogeneity. In particular, we use contemporary frameworks used by (Hendricks, Smith, and Sumner, 2014; Hendricks et al., 2014) to estimate supply response heterogeneity over the contiguous United States. Our findings suggest that the corn acreage response to price changes is heterogeneous over the United States and, when aggregated is quite different from the estimated quantity response. These results indicate that supply response conforms to economic theory and highlight the importance of heterogeneous response models for producing accurate levels of quantity response and predicting response changes from area-specific shocks to production.

Accurate measures of supply response are important for understanding supply incentives in general and for producing better targeted policy. Between 2007 and 2012, domestic US corn prices have been highly volatile. Researchers attribute this volatility to a variety of factors including domestic ethanol policy, adverse weather in corn-growing regions including a droughts in Australia in 2006 and 2007 and a drought in the US Corn Belt in 2012, and growing international demand due to rising living standards abroad, particularly in China. Researchers have been especially critical of domestic ethanol policy, estimating that Renewable Fuel Standard (RFS) legislation led to around a 30% increase in the price of corn (Piesse and Thirtle, 2009; Carter, Rausser, and Smith, 2016; Roberts and Schlenker, 2013).

While heterogeneity is not completely characterized spatially, spatial relationships are likely important. The link between proximity and "relatedness" underlies much of the work in the field of geography and the analysis of spatial correlation is the cornerstone in geography studies (Miller, 2004; Tobler, 1970). Climate and soil characteristics vary spatially. Since these features in part dictate yield potential for different crops, understanding how farmer's planting response varies over space can aid in producing more accurate aggregate quantity responses to price for the country as a whole than utilizing averaged response estimates as studies have found correlations with supply response and yield potential (Roberts and Schlenker, 2013). While the response of individuals may be spatially correlated, the impact of responses is also spatially correlated. The expansion of corn acreage in the upper Mississippi River basin and the subsequent regional increase nitrogen-based fertilizers use is thought to be a major contributer to environmental problems in the Gulf of Mexico (Wu et al., 2004; Dale et al., 2010; Hendricks et al., 2014). Such problems can be partially avoided when corn is grown in dryer areas further away from major river systems.

Accurate measurements of spatial heterogeneous response are also important for conventional studies in supply response as a whole. Supply and demand elasticities are important measures for predicting the impact price changes have on producer and consumer surplus (e.g. (Lusk and Anderson, 2004)). Estimates of these elasticity measures are often performed at a point in time and assumed to be consistent in subsequent studies, sometimes over many years. If heterogeneous supply response exists, localized shocks to supply can impact the aggregate supply responsiveness price. Finally, policy impacts are not necessarily spatially homogeneous. From 2006 to 2011, at the height of RFS legislative activity, ethanol production capacity grew from 4.9 billion gallons to 13.9 billion gallons (Carter, Rausser, and Smith, 2016). Using 10 km$\times$10 km grid-level data from the US Corn Belt, Motamed, McPhail, and Williams found that increasing the available grid capacity available within 100 km by 1% resulted in a 1.5% increase in corn acreage planted and an increase of agricultural land by 1.7% (Motamed, McPhail, and Williams, 2016).

Contemporary supply response literature is placing more emphasis on heterogeneity. While many studies have incorporated heterogeneity using a fixed effect or additive separable heterogeneous effect frameworks (Lacroix and Thomas, 2011; Haile, Kalkuhl, and Braun, 2014; Motamed, McPhail, and Williams, 2016; Haile, Kalkuhl, and von Braun, 2016), models allowing for heterogeneous supply response coefficients are becoming more popular as more evidence for non-additive heterogeneity in supply response comes to light. Many studies allow for heterogeneous response while maintaining a multi-crop analysis using discrete choice models with individual coefficients including latent class and random parameter models(Koutchadé, Carpentier, and Femenia, 2018; Claassen, Langpap, and Wu, 2017). Others use pooled estimators, running separate models over different groups characterized by climatic, geological, biological, soil, and historic farm practice regions (Hendricks, Smith, and Sumner, 2014; Hendricks et al., 2014).

The contemporary supply response literature also stresses the importance of rotational effects in overall supply response. Using a multi-stage latent class model Claassen, Langpap, and Wu found that the crop insurance had a relatively small impact on broad land usage but exhibited large influence on rotational practices. There is a growing consensus that commodity price changes tend to significantly impact crop mixes grown in areas and, while limited evidence for extensification exists, that effect on rotations was a larger overall share of supply response from price changes(Hendricks, Smith, and Sumner, 2014; Claassen, Langpap, and Wu, 2017; Langpap and Wu, 2011). To allow for both rotational effects and individual heterogeneity we use the framework from Hendricks, Smith, and Sumner (2014) which use individual Markov transition probability models over a priori specifications of heterogeneity.

To incorporate heterogeneity in this study, we separately estimate models over a set of geographic boundaries established by the Natural Resource Conservation Service (NRCS) known as Major Land Resource Areas (MLRAs). These areas spatially classify areas of the country by features that correlate with agricultural productivity. Specifically MLRAs are a set of 278 spatial subregions within the country that control for such factors as physiography, geology, climate, water, soils, biological resources, and land use. MLRAs are actually subregions within larger regions known as Land Resource Regions (LRRs) and were drawn up as refinements over similar set of criteria. Physiography considers the general elevation of the area above sea level in feet and relief. These features are associated with drainage properties. The geology criteria refers to general geologic properties of the land such as rock age. Climate delineations were produced using Parameter-Elevation Regressions on Independent Slopes Model (PRISM) data based on ranges of annual precipitation, seasonal precipitation distribution, annual temperature ranges, and seasonal freeze statistics. The water criteria considers water resources available, its quality and quality, and water use within an MLRA. This includes seasonal effects and intertemporal usage such as drought-year water usage. The soil criteria characterizes the soil taxonomy to the "great group" as defined in the Soil Survey Geographic (SSURGO) and SSURGO2 database. Soils at the great group level are divided by such things as salinization, wetness, and other important soil properties such as fragipan which impacts water and root penetration (Soil Survey Staff, 2014). Biological resources involve the descriptions of the dominant flora and fauna in the area. The last category, land use, was produced using the 1997 National Resources Inventory (NRI) data on land use. The NRI consists of survey data collected at five year intervals over 800,000 sample sites in the 50 United States and in Puerto Rico and the US Virgin Islands. Land use categories used for MLRAs include cropland, grassland, forest, urban development, water, and other (NRCS, 2006; Natural Resources Conservation Service, 2001).

2 Conceptual Model

Following (Hendricks, Smith, and Sumner, 2014; Hendricks et al., 2014), we use a Markov transition probability framework for this analysis. The basis for these models arises from the results of Hennessy (2006), who generated a model where rotational choices influence either the yield or the input usage. Under his framework, crop rotations "memory" which represents the longest lag of crop choice that influences either the input usage or the yield or both. For instance if rotations have a one-year memory then only one-year lagged crop choices influence the contemporaneous yield or input usage (Hennessy, 2006).

In this paper we assume that rotations have a one-year memory and that farmers select between two crops, corn or some other competing crop. Under these assumptions, they select from a two-period rotational pattern among some set of rotations: $R = \{\{CC\}, \{OO\}, \{OC\}\}$. The rotation CC signifies a continuous corn rotation, OO is a continuous other-crop rotation, and OC is a rotation where the other crop is selected, followed by a corn rotation[1]. Farmers are assumed to be price takers and select a rotation to maximize profits conditional on prices. Also for the purposes of a parsimonious conceptual model, we assume that production follows constant economies of scale. Let c_{it} be the indicator function:

$$c_{it} = \begin{cases} 1 & | \text{ field } i \text{ planted corn at time } t \\ 0 & | \text{ otherwise.} \end{cases} \tag{1}$$

If rotational patterns influence the productivity and input requirements of planting a crop and rotations have one-year memory, we can write the profit functions of planting a crop in a given year is a function of the lagged crop indicator ($c_{i,t-1}$). Equations 2 and 3 show the profits from growing corn or the other selected crop. Both profit functions take on the classic revenue minus cost form where p_i^j is the j^{th}'s crop output price for field i for crop j which is, for simplicity, assumed constant over time. The yield function y_{it}^j is the yield for crop j at time t in field (i).[2] We assume that yield are functions of inputs for the respective crop $\left(x_{it}^j\right)$ with the respective input price vector w_i^j which, like output prices are assumed constant over time. One potential benefit of changing crops versus selecting a continuous crop rotation is that it may save input. For instance, planting a cover crop helps replenish nutritional content, saving fertilizer. If the farmer planted the other crop (corn) in $t-1$, planting corn (the other crop) in year t saves the farmer N^{OC} (N^{CO}) in input. Alternating crops between years also produces a yield effect, increasing yields by B^{OC} when planting corn and by B^{CO} when planting the other crop relative to mono-cropping.

$$\pi_{it}^C = \pi_{it}^C\left(x_{it}^C, c_{it-1}\right) = p_i^C\left[y_i^C\left(x_{it}^C + (1-c_{it-1})\, N_i^{OC}\right) + (1-c_{it-1})\, B_i^{OC}\right] - w_i x_{it}^C \tag{2}$$

$$\pi_{it}^O = \pi_{it}^O\left(x_{it}^O, c_{it-1}\right) = p_i^O\left[y_i^O\left(x_{it}^O + c_{it-1} N_i^{CO}\right) + c_{it-1} B_i^{CO}\right] - w_i x_{it}^O \tag{3}$$

Given a previous crop selection ($c_{i,t-1}$) the farmer will make a contemporaneous planting decision to maximize the conditional rotational profit function in equation 4.

$$\pi_{it} = \max_{\left(c_{it},\ x_{it}^C,\ x_{it}^O\right)\big|_{t=1}^{2}} \sum_{t=1}^{2} c_{it}\pi_{it}^C\left(x_{it}^C, c_{it-1}\right) + (1-c_{it})\,\pi_{it}^O\left(x_{it}^O, c_{it-1}\right) \tag{4}$$

By Hennessy (2006), the first order conditions for the optimal input vector x_{it}^j is a linear function of the rotational input savings according to equations 5 and 6. The functions $g^j(\cdot)$ are inverse functions of the j^{th}

[1] In this framework an OC rotation is identical to a CO rotation where an non-corn selected crop choice follows a corn choice.

[2] These are per-acre profit functions and so the only critical linkage between prices and yield is that price in dollars per "unit" have to correspond to the yield in "units" per acre. This will be a non-trivial detail since we represent the other crop as a composite crop with corresponding units besides bushels. In addition, these yields are indexed to the field-level, these yield functions can be scaled up to accommodate the size of the respective field.

crop's marginal yield equation

$$x_{it}^{c\star} = g^C\left(w_i, p_i^C\right) - \left(1 - c_{t-1}\right) N_i^{OC} \tag{5}$$

$$x_{it}^{o\star} = g^O\left(w_i, p_i^O\right) - c_{t-1} N_i^{CO}. \tag{6}$$

Inserting equations 5 and 6 into equations 2 and 3 and combining it all into 4, yields equation 7. This states that the profit of an OC rotation relative to a CC or OO rotation is determined by differences in the relative profits of one crop over another and linear functions of the yield premiums and input expense savings from rotations. Notice that y_{it}^j for $j \in \{O, C\}$ are not functions of the previous crop choices and are therefore independent of the rotation selection. This means that the profit function can be written more simply as equation 8. This and the concept of a single memory crop rotational pattern will be important in building our empirical model in the later sections.

$$\sum_{t=1}^{2} c_{it}\left[p_i^C\left[y_i^C\left(g^C\left(w_i, p_i^C\right)\right) + \left(1 - c_{it-1}\right) B_i^{OC}\right] - w_i\left(g^C\left(w_i, p_i^C\right) - \left(1 - c_{t-1}\right) N_i^{OC}\right)\right] + \tag{7}$$
$$\left(1 - c_{it}\right)\left[p_i^O\left[y_i^O\left(g^O\left(w_i, p_i^O\right)\right) + c_{it-1} B_i^{CO}\right] - w_i\left(g^O\left(w_i, p_i^O\right) - c_{t-1} N_i^{CO}\right)\right]$$

$$\sum_{t=1}^{2} c_{it}\left[p_i^C\left[y_i^C + \left(1 - c_{it-1}\right) B_i^{OC}\right] - w_i\left(g_i^C - \left(1 - c_{t-1}\right) N_i^{OC}\right)\right] + \tag{8}$$
$$\left(1 - c_{it}\right)\left[p_i^O\left[y_i^O + c_{it-1} B_i^{CO}\right] - w_i\left(g_i^O - c_{t-1} N_i^{CO}\right)\right]$$

This in turn means that profits for each rotation are:

$$\langle CC \rangle : 2\left(p_i^C y_i^C - w_i g_i^C\right) \tag{9}$$

$$\langle OO \rangle : 2\left(p_i^O y_i^O - w_i g_i^O\right) \tag{10}$$

$$\langle CO \rangle : p_i^C y_i^C + p_i^O y_i^O - w_i g_i^C - w_i g_i^O + p_i^C B_i^{OC} + p_i^O B_i^{CO} + w_i\left[N_i^{CO} + N_i^{OC}\right] \tag{11}$$

Comparing values of equations 9, 10, and 11, the each of the respective rotations will be optimal according to:

$$\langle CC^\star \rangle : p_i^C y_i^C - p_i^O y_i^O > p_i^C B_i^{OC} + p_i^O B_i^{CO} + w_i\left[N_i^{CO} + N_i^{OC}\right] \tag{12}$$

$$\langle OO^\star \rangle : p_i^O y_i^O - p_i^C y_i^C > p_i^C B_i^{OC} + p_i^O B_i^{CO} + w_i\left[N_i^{CO} + N_i^{OC}\right] \tag{13}$$

$$\langle CO^\star \rangle : p_i^C B_i^{OC} + p_i^O B_i^{CO} + w_i\left[N_i^{CO} + N_i^{OC}\right] \geqslant \left|p_i^C y_i^C - p_i^O y_i^O\right| \tag{14}$$

This implies that the choice among alternative rotation patterns will be a linear function of prices and individual yield and input benefits of alternating crops over two consecutive years.

3 Data

To study supply response at the farm-level we require two pieces of information. The first is the crop choices that were made and the second is the domain over which farmers makes their choices. The first component comes from the USDA's Cropland Data Layer (CDL) which provides a categorical raster file generated from satellite images. This raster file provides estimates of crop choice at the 30m angular resolution. We define the farmer's choice domain using the 2008 Common Land Unit (CLU) shapefile these data represent the smallest contiguous and bounded land units with a common owner and a common producer (Woodard, 2016).[3]

While the CDL dataset provides high resolution data over the contiguous United States, not all states entered the dataset in the same year. For example North Dakota entered the CDL in its inaugural year in 1997 while Texas entered in 11 years later in 2008. In the base model heterogeneity was incorporated using pooled regressions over the Natural Resource Conservation Service's (NRCS) MRLA boundaries. An attractive feature of MLRAs is that characterize similarities in growing environments as opposed to political boundaries. To ensure that MLRA-specific datasets are balanced panels, we restricted the years of analysis for each MLRA by the latest-earliest year for each of the CLU observations. For instance, North Dakota's CDL observations begin in 1997 while South Dakota's begin in 2006. A MLRA that overlaps with North and South Dakota will begin in 2006 as this is the latest year between 1997 and 2006. Figure 1 shows the MLRA map and earliest year of analysis for each MLRA[4].

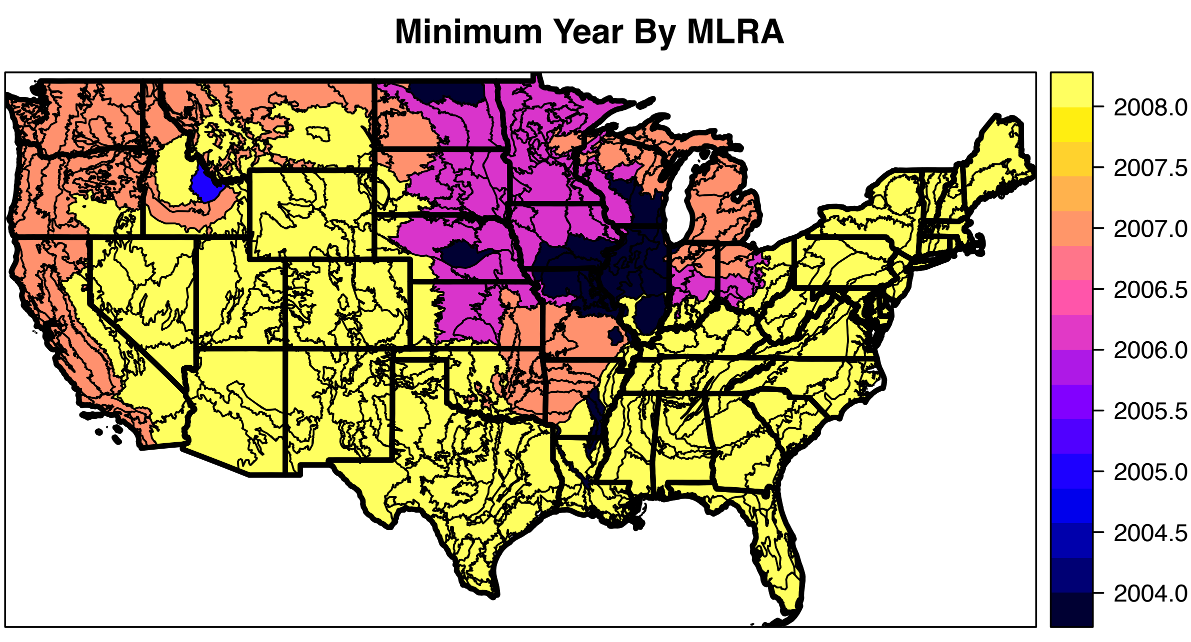

Figure 1: Minimum Observation Year by Major Land Resource Area

[3]Areas of the country with missing common land unit boundaries were filled in with boundaries from Yan and Roy (2016).
[4]The availability of price data was also a limiting factor. The panel dataset that offers the best spatial resolution over the longest time frame begins in 2004.

The importance of including expected prices and the debate how to properly account for expectations has been a consistent topic of discussion in agricultural supply response literature (Nerlove, 1956; Haile, Kalkuhl, and von Braun, 2016; Miao, Khanna, and Huang, 2016; Roberts and Schlenker, 2013; Gardner, 1976). A common theme in many of these discussions stresses the importance of using prices that reflect harvest-time expectations at or before the time planting when modeling planting decisions. The earliest and simplest forms of expected price were simply the previous years lagged harvest price. In Nerlove's famous supply response paper, he used an moving average model to produce expected prices (Nerlove, 1956). Others have used futures prices since, under the efficient market hypothesis, these prices should reflect information about expected price changes (Gardner, 1976; Haile, Kalkuhl, and Braun, 2014; Roberts and Schlenker, 2013). We employ a mixture of lagged prices and futures prices to form our expected price series. Much of the nation's corn acreage is planted in the month of April with planting starting in early April or late March. We assume that the planning process begins in the months of January and February as this gives time for required crop specific land preparation before planting begins. To construct expected prices, we first average local daily spot prices over the course of the months of January and February which we call the planting price $\left(C_{it}^{P}\right)$. Next, we average daily the nearby futures contract price and the harvest-time futures contract price for the respective commodities in January and February $\left(F_t^P \text{ and } F_t^H \text{ respectfully}\right)$. We construct the expected harvest-time spot prices according to equation 15. We first compute the average annual nearby basis at planting and add the harvest-time futures price. The pre-plant average nearby basis incorporates the local basis pattern of the individual market, providing an estimate for basis at harvest time. The average harvest-time price acts serves as the projected harvest time futures price. Adding the expected basis to to the expected harvest futures price gives the projected harvest-time local market spot price. Another way to think about this expected price is that the difference between the harvest-time contract price and the nearby contract price is the markets estimated cost of carry to harvest. If these prices are properly reflecting market transactions, this means that individuals with grain holdings would be willing to hold their grain in storage to receive this expected price at harvest.

$$E_{it}\left[P_{it}^{H}\right] = F_t^{PH} + \left[C_{it}^{P} - F_t^{PN}\right] = \underbrace{\left[F_t^{PH} - F_t^{PN}\right]}_{\text{Expected Cost of Carry}} + C_{it}^{P} \tag{15}$$

This study utilizes price data for corn, soybeans, hard red winter wheat, hard red spring wheat, soft red winter wheat, rice, and cotton. These prices are all quoted in dollars per bushel with the exceptions of rice and cotton which were quoted in dollars per pound. Daily values for futures prices for these commodities are readily available. We collected from these data from the Data Transfer Network (DTN). Spot price datasets were constructed using a combination of data from the Data Transfer Network (DTN) and Cash Grain Bids (CGB). To ensure that our estimates are not dictated by a small number of observed prices, we remove markets with less than 10 observations over January and February. Between these two sources data from 2004 to 2016 were collected from 1,367 corn price locations, 1,252 soybean price locations, 84 hard red spring wheat (HRSW) price locations, 96 hard red winter wheat (HRWW) price location, and 123 soft red winter wheat (SRWW) price locations. Rice and cotton prices were collected from the National Agricultural Statistics Service's (NASS) Quickstats at the national level from 2004 to 2016. Coverage for the continuously observed local markets was rather good, densely covering most of the major field crop production areas. Figure 3 shows the price coverage by crop.

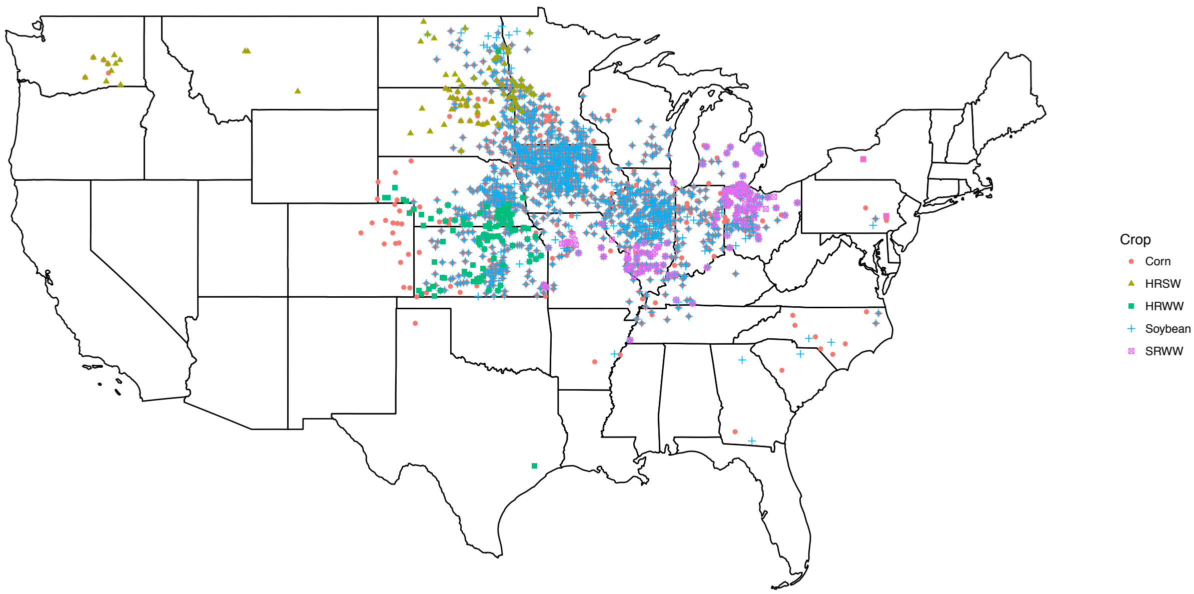

Figure 2: Commodity Price Locations Continuously Observed Between 2004 and 2016

Using the location of the market city as a reference, we used these prices to construct annual basis maps for each commodity over the contiguous United States. After estimating the expected commodity prices over the markets we interpolated these estimates using ordinary kriging. Ordinary kriging has advantages over other interpolation procedures such as inverse distance weighting since it takes the spatial correlation of the values observations into account to minimize the variance of the estimates. Basis map values were estimated over a raster image with a resolution of 0.01 square degrees (or approximately 36 square miles). Figures 3 shows an example of basis maps in 2009 and their respective standard errors for the localized markets of corn. To maintain consistency of the price expectations estimates over time, the original observation set contains only markets with continuously observed price averages in every year from 2004 to 2016.

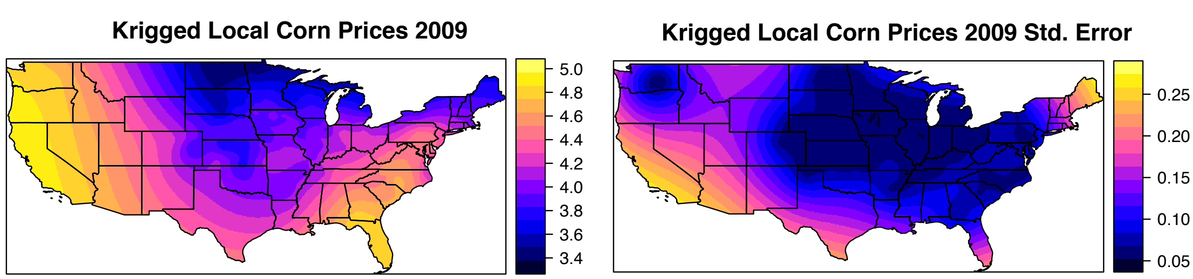

Figure 3: 2009 Corn Expected Price Map

Incorporating rotations in heterogeneous supply response studies is challenging since rotational choices are more likely to vary spatially. By focusing on the supply response of a single crop, corn, we can simplify the study somewhat by considering rotations between corn and some other crop. In this study we characterize the other crop price using a weighted average of soybeans, HRWW, SRWW, HRSW, cotton, and rice prices. Using a common weighting scheme for every field is problematic since the set of relevant alternative crops to corn production will differ by the area of the country. We construct the other crop price as a Laspeyres Index. This index is commonly used to compute the Consumer Price Index and creates a "basket" of commodities indexed from $k = 1, \ldots, K$ using the quantities in some base period (period 0) to track the basket's changing total product as a result of price changes. Equation 16 shows the functional form where p_{tk} is the price for commodity k at time t and q_{0k} is the total quantity of crop k produced in period 0. Using unique price indices for each MLRA ensures that the other crop price will largely consist of crops grown in the region. For instance, the dominant alternative crop in the state of Iowa is soybeans. If the alternative acreage only consists of soybeans, then $q_{k'} = 0 \; \forall k'$ where k' is not soybeans, which means that only soybeans prices would enter the price index.

$$P_t^O = \frac{\sum_{k=1}^{K} p_{tk} q_{0k}}{\sum_{k=1}^{K} p_{0k} q_{0k}} \tag{16}$$

There are several aspects that complicate the use of the Laspeyres index in this study. First, crop choices are subject to change over time so it is unclear that a single quantity measure q_0 would sensibly represent the typical crop choice basket over our study. Second, our analysis is at the CLU level which generally have a single observation each year. A particular crop planted at the beginning of the analysis does not preclude another crop from being considered in the future. This means we need to property define of the initial quantity q_{0k}. Lastly while we observe crop choices on some fields before 2008, the CDL dataset did not gain full coverage until 2008. Therefore we choose 2008 as the period to define the index to maintain consistency over space. Defining the index using production statistics before 2008 potentially introduces spatial differences in an independent variable that is exclusively due to data availability. This could be characterized as heterogeneously introducing degrees of measurement error over different areas of the country. To address these issues, we modified the Laspeyres index. In particular we define q_{0k} as the MLRA-specific total production of crop k from 2008 to 2016. We compute total production by each MLRA for each crop by merging the field level data with a county-level yield dataset, multiplying the yield with the field level acreage choices, then summing over the total MLRA from 2008 to 2016. Average annual yield data for each of these crops are available at the county-level through the National Agricultural Statistics Service (NASS) and we used agricultural district-level to fill in for missing county-level observations.

While the MLRAs are drawn up to reflect differences in soils and climate there can still be substantial intra-regional variation in soils and weather conditions. To control for such variation, we include the soil horizon composition and the slope, both provided by the SSURGO dataset. The soil horizon composition was broken up into three components clay, silt, and sand represented as percentages of the soil's horizon and add up to one. These are the three major features that define the texture of the soil. We also include the slope of the field which is a key component controlling runoff, and erosion. We also control for extreme precipitation events in the pre-plant stage. Extremely wet conditions can potentially delay planting and could cause farmers to plant alternative crops such as soybeans with later planting dates. Extremely dry conditions may have similar effects and may shift crops to more drought resistant crops. We describe extreme planting precipitation conditions include two indicator variables. The first variable is equal if the PRISM cell's annual total April-May precipitation was at or below the 25th percentile over the respective MLRA[5]. The other was equal to one if the PRISM grid cell April-May precipitation was above the 75th percentiles over the MLRA. These percentiles were calculated based off of data 34 years of historical data between 1983 to 2016.

4 Empirically Modeling Price Response

A direct result of the conceptual model is that the relative profit of selecting one rotation over another is a linear function of input and output for prices for each of the crops and rotational benefits with respect to inputs and yields. With the assumption that crop-rotations have a one-period memory, this means that rotations can be modeled as a first-order Markov decision process. This directly motivates the probabilistic structure used by Hendricks, Smith, and Sumner (2014) and Hendricks et al. (2014) which is the basis for this study. The choice framework in the empirical model can be characterized as two first-order Markov transition probability equations describing the probability of a farmer selecting corn as a crop conditional on the previous crop choice. These choices are characterized by the indicator variables y_{it}^{OC} and y_{it}^{CC} where:

$$y_{it}^{c_1 c_2} = \begin{cases} 1 \mid & \text{Crop type } c_1 \text{ planted on field } i \text{ in } t-1 \text{ and type } c_2 \text{ chosen in } t \\ 0 \mid & \text{Otherwise} \end{cases} \tag{17}$$

Equations 18 and 19 show the structure of the Markov transition probabilities of planting corn given corn or some other crop was planted in the previous year respectfully. The estimated coefficients β^C and β^O are of interest as they represent the marginal influence of crop prices have on the conditional probability of planting corn in the contemporary period. Only fields raising either corn or the other crop between the two consecutive periods are in the sample used to estimate these models. The probability of planting corn in period t takes the form of equation 20 and the probability of planting some other crop in period t takes the form of equation 21. Assuming long-run probabilities exist, they can be found by setting the lagged probabilities equal to the contemporaneous probabilities in equation 20[6]. This can be represented as equation 22 shown by plugging equation 21 into equation 20 and setting $\Pi_{it-1}^{C}\big|_{LR} = \Pi_{it}^{C} = \Pi_{it-1}^{C}$. Equations 23 provides the functional forms of the short-run counterpart which allows for fluctuations in the transition

[5]Corn planting for many of the largest corn producing states is most active in these months (NASS, USDA, 2010).
[6]The long-run probability of planting other crops is the complement of the long-run probability of planting corn.

probabilities.

$$y_{it}^{CC} = \beta_{10} + \beta_1^C P_{it}^C + \beta_1^O P_{it}^O + \gamma_1 \mathbf{X}_{it} + \varepsilon_{1it} \tag{18}$$

$$y_{it}^{OC} = \beta_{20} + \beta_2^C P_{it}^C + \beta_2^O P_{it}^O + \gamma_2 \mathbf{X}_{it} + \varepsilon_{2it} \tag{19}$$

$$\Pi_{it}^C = \Pi_{it-1}^C \hat{y}_{it}^{CC} + \Pi_{it-1}^O \hat{y}_{it}^{OC} \tag{20}$$

$$\Pi_{it}^O = 1 - \Pi_{it}^C \tag{21}$$

$$\Pi_{it}^C\big|_{LR} = \frac{\hat{y}_{it}^{OC}}{1 - \hat{y}_{it}^{CC} + \hat{y}_{it}^{OC}} \tag{22}$$

$$\Pi_{it}^C\big|_{SR} = \Pi_{it}^C\big|_{LR}\, \hat{y}_{it}^{CC} + \left(1 - \Pi_{it}^C\big|_{LR}\right) \hat{y}_{it}^{OC} \tag{23}$$

Using these equations, the marginal effect of price changes are computed using their first derivatives with respect to prices. The major differences between the marginal of short-run probabilities and their long-run counterparts is that the $\Pi_{it}^{(\cdot)}\Big|_{LR}$ terms are taken as given constants.

$$\left.\frac{\partial \Pi_{it}^C}{\partial P_{it}^k}\right|_{LR} = \frac{\left[1 - \hat{y}_{it}^{CC}\right] \hat{\beta}_2^k + \hat{y}_{it}^{OC} \hat{\beta}_1^k}{\left[1 - \hat{y}_{it}^{CC} + \hat{y}_{it}^{OC}\right]^2} \tag{24}$$

$$\left.\frac{\partial \Pi_{it}^C}{\partial P_{it}^k}\right|_{SR} = \Pi_{it}^C\big|_{LR}\, \hat{\beta}_1^k + \left(1 - \Pi_{it}^C\big|_{LR}\right) \hat{\beta}_2^k \tag{25}$$

We construct the elasticity terms by multiplying the respective marginal effects by the ratio of the field-specific price and the predicted probabilities. These elasticity terms, are useful for estimating the relative effect that a price change will have on production in different areas of the country. The Markov transition probability models described in this section took the form of linear probability models. While they provide a simpler functional form for exposition and estimation and can perform well, their estimated values are not confined to the $[0, 1]$ interval. Because the dataset is so large and varied and since we are interested in estimating MLRA-level elasticities, we use logit models instead since the elasticity measures rely on accurate probabilistic estimates. In this case the $\hat{\beta}$ coefficients in equations 24 and 25 take the form of the MLRA's respective average partial marginal effects.

5 Results

Performing heterogeneous supply response an area as large as the United States is challenging. While the Cropland Data Layer distinguishes a variety of crops, their respective prices are more difficult to find. In addition, many crops do not have associated futures contracts which makes expected prices difficult to construct. Since we account for heterogeneity using pooled regressions at the MLRA-level, we need to ensure each of our regressions have enough observations to produce reliable coefficient estimates. Our Markov transition probability regression modeling strategy requires that our training data consists of only observations with two consecutive corn or the other crop choices which can constrain the MLRA sample size.

For these reasons we divided the observation strata into 5 groups, corn, priced crops, other crops, cropland, and non-cropland. Table 1 shows the CDL observation designations. Corn consists of observations where the CDL signifies only a conventional corn observation. That is, "corn" does not include double-cropping observations involving corn (e.g. double-cropping corn and soybeans), and less conventional varieties such as sweet corn, or popcorn. Priced crops are crops with expected price observations that enter the "other" crop price index value (soybeans, rice, wheat varieties, cotton) and associated double-cropped observations with these crops (e.g. winter wheat-cotton double cropped observations). Other crops are crops that are considered substitutes in production to the priced crops. This category includes double-cropped observations containing these crops. The "Cropland" category contains other crops that are less substitutable to the priced crops. This category includes specialized fruit and vegetable crops and perennial crops such as alfalfa. The final category, non-cropland, contains land uses that are not immediately suitable for crop production including marshland, pasture, and developed lands.

Corn	Sweet Potatoes	Greens	Squash
Cotton	Triticale	Herbs	Strawberries
Rice	Alfalfa	Honeydew Melons	Sugarcane
Soybeans	Almonds	Lettuce	Sweet Corn
Spring Wheat	Apples	Mint	Switchgrass
Winter Wheat	Apricots	Misc Vegs	Tobacco
Barley	Aquaculture	Nectarines	Tomatoes
Buckwheat	Asparagus	Olives	Turnips
Camelina	Blueberries	Onions	Vetch
Canola	Broccoli	Oranges	Walnuts
Dry Beans	Cabbage	Other Crops	Watermelons
Durum Wheat	Caneberries	Other Hay/Non Alfalfa	Barren
Fallow/Idle Cropland	Cantaloupes	Other Tree Crops	Clouds/No Data
Flaxseed	Carrots	Peaches	Deciduous Forest
Hops	Cauliflower	Peanuts	Developed (All Levels)
Lentils	Celery	Pears	Evergreen Forest
Millet	Cherries	Peas	Forest
Mustard	Chick Peas	Pecans	Grassland/Pasture
Oats	Christmas Trees	Peppers	Herbaceous Wetlands
Other Small Grains	Citrus	Pistachios	Mixed Forest
Potatoes	Clover/Wildflowers	Plums	Nonag/Undefined
Rape Seed	Cranberries	Pomegranates	Open Water
Rye	Cucumbers	Pop or Orn Corn	Perennial Ice/Snow
Safflower	Eggplants	Prunes	Shrubland
Sorghum	Fruits	Pumpkins	Water
Speltz	Garlic	Radishes	Wetlands
Sugar Beets	Gourds	Shrubland	Woody Wetlands
Sunflower	Grapes	Sod/Grass Seed	
Legend			
Priced Crop	Other Crop	Cropland	Non-Cropland

Table 1: Cropland Data Layer Observation Designations

Using the CDL designations, we filter the data using three hurdles. First, we filter out MLRAs with less than 20% of the MLRA's total acreage applied to corn, priced crops, other crops, or cropland. This was done to filter out MLRAs with modest agricultural activity such as desert, mountainous, and developed areas. We ultimately use the price index made up of the prices of the priced crops as the price for alternative other crops. To ensure that the other price index can reasonably serve as a measurement for prices for alternative crops to corn, we filter out MLRAs where the corn and priced crops make up less than 50% of their other crops. Third, since our goal is to understand how corn plantings respond to price, we also need a minimum threshold for corn observations within each MLRA. We therefore filtered out MLRAs with less than 10% of their total other crop acreage in corn. In addition to these three hurdles, we also remove MLRAs with less than 50,000 total observations and MLRAs where less than 35,000 observations enter either of their respective Markov transition regressions[7]. After filtering, our data includes 57 MLRAs and over 29.2 million individual observations. According to the cropland data layer and county-level NASS estimated yields, these MLRAs accounted for 92% percent of the nation's corn production, 95% of its soybean production, 46% of its cotton production over 47% of its wheat production, and 77% of its rice production between 2008 and 2016.

There was a fair degree of heterogeneity in the supply response between MLRAs. Table 5 shows the summary statistics over the 57 MLRAs. Between these MLRAs the average estimated long-run probability of planting corn was around 42% over the years of the analysis. The MLRA with the smallest long-run probability had just under 10% probability of planting corn. Based on our three hurdle approach, suggests that our models are working properly. The corn supply response to price were as expected on average. The probability of planting corn tends to increase when the price of corn increases and decrease when the price of the other composite commodity increases. On average, a one dollar increase in the price of corn correlates with a 5% increase in the probability of planting corn in the short-run but, as the table suggests, there is a lot of heterogeneity in corn supply response to price. Some MLRAs had a negative average own-price response. In others, a dollar increase in the price of corn is estimated to increase the probability of planting corn by over 20%. On the whole, the average values agree with Hendricks, Smith, and Sumner (2014) in that the short-run effects are larger than the long-run effects. This is a consequence of prices impacting temporary rotational patterns. While on average results are as expected, in several MLRAs, the estimated signs were not as expected. Figure 4 shows the distribution of the marginal effects.

[7] Observations also exclude CLUs that are 15 acres or larger.

Statistic	min	mean	max
Long-Run Probability of Planting Corn	0.0952	0.424	0.738
Short-Run Own-Price Marginal Effect	-0.0527	0.0549	0.228
Long-Run Own-Price Marginal Effect	-0.0916	0.0427	0.231
Short-Run Cross-Price Marginal Effect	-0.496	-0.109	0.231
Long-Run Cross-Price Marginal Effect	-0.508	-0.0818	0.335
Short-Run Own-Price Elasticity	-0.501	0.704	4.22
Long-Run Own-Price Elasticity	-0.871	0.606	4.67
Short-Run Cross-Price Elasticity	-4.02	-0.563	0.758
Long-Run Cross-Price Elasticity	-4.46	-0.486	1.10

Table 2: Corn Supply Response to Price Summary Statistics

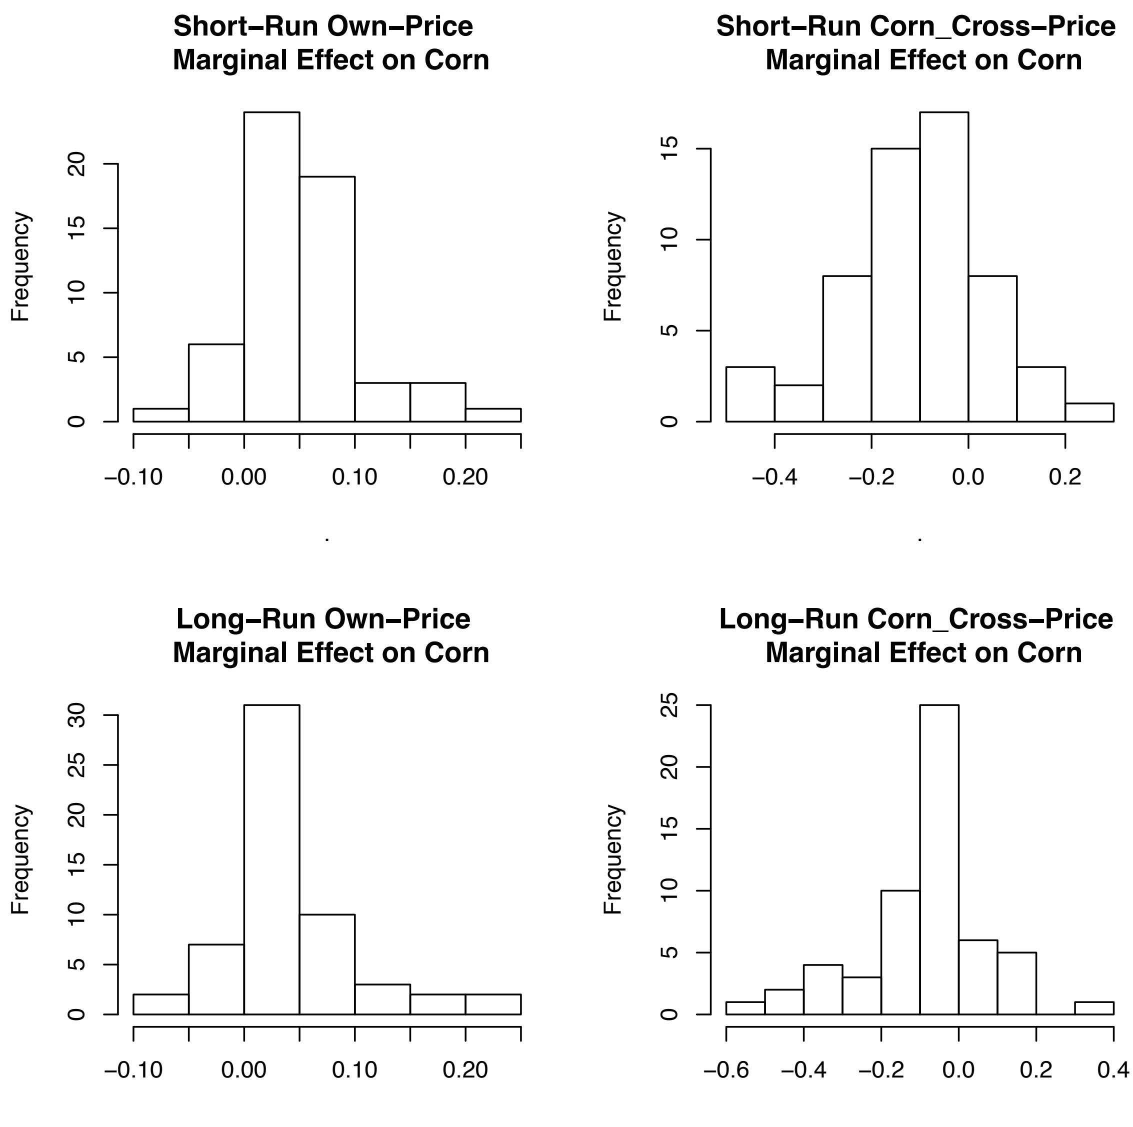

Figure 4: MLRA-Specific Marginal Effects on the Probability of Planting Corn

Approximately 10% of the MLRA results were not as expected. The MLRAs with negative own-price marginal effects tended to have positive coefficient values in the OC Markov transition probability model. These MLRAs also tended to have smaller corn coefficient value in the corn-given-other Markov transition probability model. This suggests that negative values are likely due to high coefficients estimator- and corn estimated probability-variance. To further check our elasticity estimates, we compared them with simple county-level NASS elasticity estimates. To produce our NASS estimates, we added up the planted corn acreage reported by NASS for every county that touches each MLRA over each year from 2008 to 2016. We then used the MLRA average corn prices and the average other price index to compute our NASS elasticity estimates in each year. We finally estimate MLRA specific elasticities with the following equation 26. Here $\overline{Corn\ Price}_t$ is the mean-MLRA corn price in a particular year and $\overline{Other\ Price\ Index}_t$ is the other price index counterpart. $NASS\ Corn\ Acres_t$ is the sum of county-level NASS acreage in year t for counties that touch a given MLRA. While the NASS estimate will not be identical to our CDL estimates, they provide a general comparison to our MLRA estimates using CDL data. Figure 5 plots the NASS estimated elasticities relative to our CDL estimated elasticities. The NASS and CDL estimates do not correspond perfectly but do have relationship with a near unit slope and a near zero intercept. This lends further credibility to the CDL estimates. The acreage response elasticities estimated with the NASS data were estimated with a short time-series of only 9 years (observations). Despite the small sample size, many of this of these simple models produced statistically significant elasticity terms. To further address the elasticity negative values from the CDL, we estimated the standard error of $\hat{\beta}_1$ in our NASS equations and divided the points by standard error quantiles. This further suggests that high marginal effect variance may be to blame for the negative own-price elasticity of supply estimates.

$$\ln\left(NASS\ Corn\ Acres_t\right) = \beta_0 + \beta_1 \ln\left(\overline{Corn\ Price}_t\right) + \beta_1 \ln\left(\overline{Other\ Price\ Index}_t\right) + \beta_3 t + \eta_t \tag{26}$$

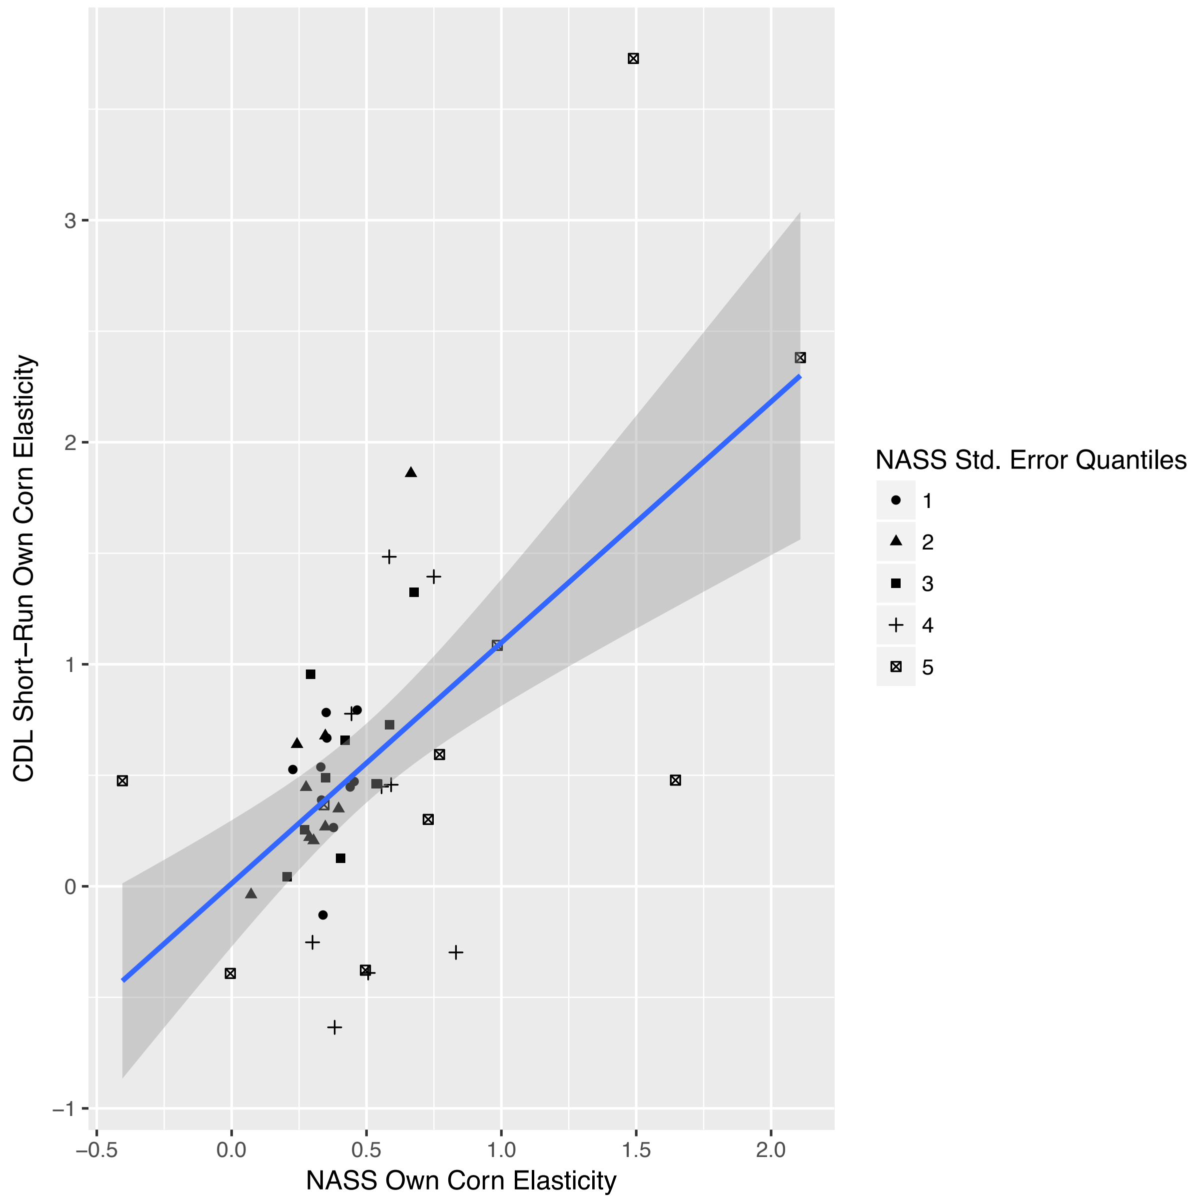

Figure 5: Short-Run Corn Price Elasticity NASS Comparison

We now present the elasticity results of the supply response study. The spatial distribution of the supply response is similar in the long-run so we present the short-run elasticity maps in figures 6 and 7 to save space. The marginal effects correspond closely with the elasticity values and so we focus our attention on the elasticity measures. These maps show that there is a high degree of heterogeneity of supply response between MLRAs within the traditional corn belt and those outside of the corn belt. The probability of planting corn is particular sensitive in the east and central areas of North Dakota and along Mississippi delta. The model shows that corn plantings in the eastern portion of North Dakota and the western portion of Minnesota had an especially high price sensitivity with an exceptional own-price elasticity of over 4.2. While this value is extreme, the MLRA's corresponding NASS value was around 1.5 which also indicates that there is an elastic supply response in the region (see figure 5). This area of the country is known for its flooding. Furthermore, lagged snowfall over the winter and rapid snow melt are two primary factors in spring flooding in this region. If adverse weather events were correlated with the low corn price years the elasticity measurements could be misleading. This may be the case since, over our coverage of the study, especially severe floods impacted east-central North Dakota in 2009 and 2011 which coincide with low corn price years.

MLRAs with high own-price elasticity also tended to have deep negative cross price elasticities. While it can be difficult to read on the maps, the sensitivity to prices tends to be higher in the Northern central United States, for instance MLRAs in Iowa tended to have higher price sensitivity than MLRAs in Illinois and Indiana.

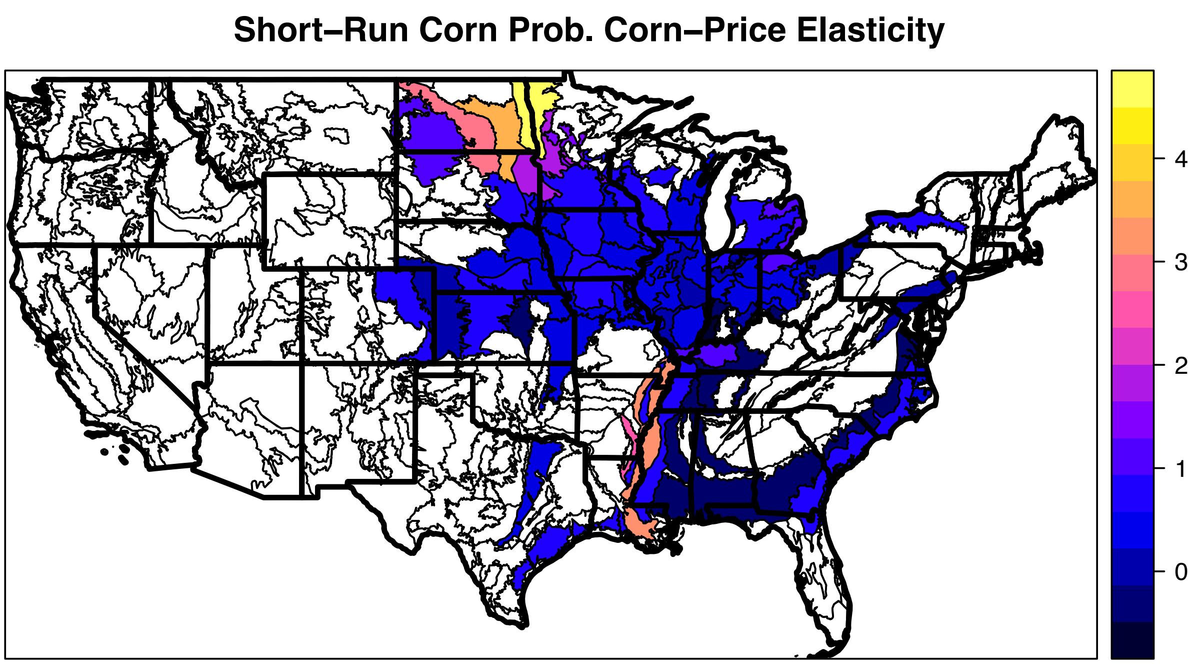

Figure 6: Short-Run Corn Price Elasticity of the Probability of Planting Corn

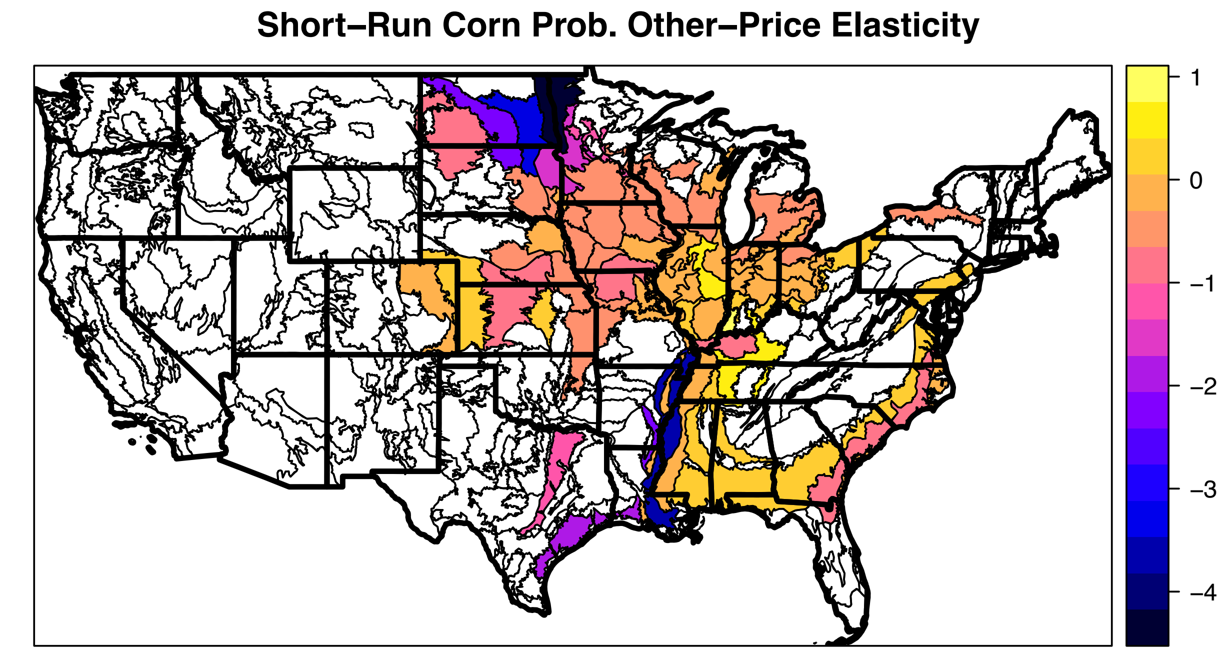

Figure 7: Short-Run Other Price Elasticity of the Probability of Planting Corn

6 Acreage Supply Response Versus Quantity Supply Response

To conclude the analysis, we compare the estimated quantity response to the acreage response to a dollar increase in the price of corn. We do this by combining the field level own-price marginal effect the NASS yield values. Over these 57 MLRAs, a $1 increase in the price of corn increases the total acreage planted to corn by 16.7% while total production would only increase by 6%. This suggests that there is a negative correlation between yield potential and cropland conversion. This is expected if, for a given price farmers initially allocate lands to their most profitable use. If the price of corn increases while holding the price of all other crops constant, we would expect that land with lower corn yield potential be brought into corn production. It therefore makes sense that the expected acreage increase is larger than the expected quantity increase. Figures 8 and 9 show the expected land-use percentage change by MLRA and the overall additional expected bushels from each MLRA respectfully using 2016 yield estimates. These maps show that a dollar increase in the corn price would induce large acreage in west-central North Dakota, more supply response would occur in the eastern part of North Dakota, central Minnesota and central Iowa. It also shows that the Mississippi Delta has high acreage and quantity responses to an increase in the price of corn.

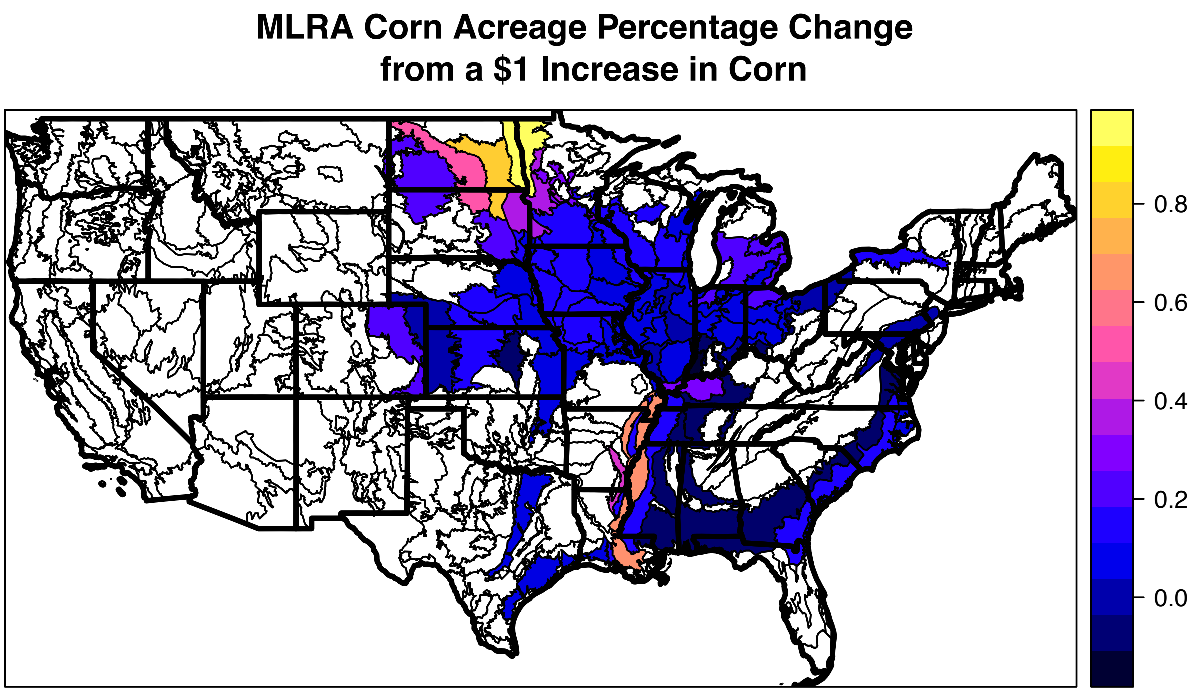

Figure 8: Expected MLRA Percentage Change in Corn Acreage from a $1 Increase in Corn Price

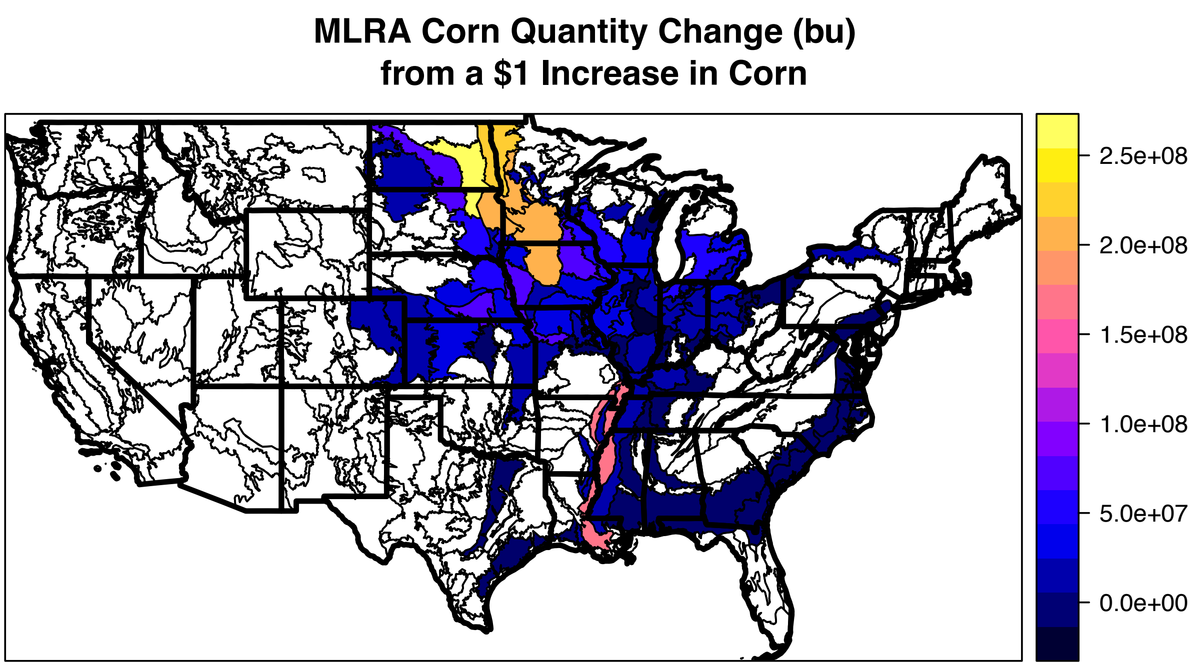

Figure 9: Expected Corn Quantity Change in Bushels from a $1 Increase in Corn Price

7 Conclusions

The initial results in this paper indicate that there is a considerable difference in the degree to which corn producers respond to corn prices. Planted acreage response to price ranged from near zero to highly elastic. These results tended to be consistent with estimates using NASS planted acreage estimates. Corn planting decisions in the northern United States particularly in east-central North Dakota and western Minnesota were quite sensitive to corn price fluctuations. There could be many reasons for high price sensitivity in this area. First, prices in this part of the country are relatively low compare to other areas inside of the corn belt. In general, the lowest corn prices in the country are in central Dakotas. The spatial distribution of prices of the 2009 corn basis map in figure 3 is quite stable over time. This could mean that, in most years, corn prices are not high enough relative to other crops grown in the area. The eastern-most MLRA is particularly known for sugar beet production. While the corn-soybean price ratio rose by nearly 30% between 2005 and 2007, the corn-sugar beet price ratio nearly doubled over this same time frame. Relative to other crops, sugar beet prices also tended to be more stable from 2008 to 2016 which may explain the high sensitivity to crop price volatility.

Corn planting decisions in the Mississippi Delta was also sensitive to price changes. The corn-to-rice price ratio also nearly doubled between 2010 and 2011 in Arkansas and Mississippi. The prevalence of rice production in the region coupled with the fact that rice-soybean rotations are common could have lead farmers to transition to the corn-soybean rotations usually seen in traditional Corn Belt states like Iowa. High corn prices may have also had rotational impacts. While planting cotton to monoculture is popular, researchers studying the delta suggest that this practice introduces the crop to pest risk and suggest a corn-cotton rotation as an alternative (Snipes, 2005).

Accounting for heterogeneous acreage response, we find that in percentage terms, the aggregate acreage response is above the yield response. This suggests that that there is a negative spatial correlation between yield potential and acreage response. This conforms with economic theory where farmers will utilize the land to maximize profits for a given price ratio. When the price of corn increases relative to price of other crops, farmers the marginal product of inputs for corn tends to fall. This highlights the importance of heterogeneous supply response when yield response at or near zero.

There are many areas for future research. In this analysis we use a Markov-transition probability framework. This framework relies on assumptions of the rotational memory of the crops in question. In this paper, we assume that corn and the composite other crop both have a single period memory which greatly reduces the number of equations we need to estimate (Hendricks, Smith, and Sumner, 2014; Hennessy, 2006). While there is evidence that corn rotations have a single period memory, this may not the be the case for the other composite crop. Furthermore, there is no research that we are aware of suggesting this memory holds generally over the entire country. We plan on extending this idea by testing rotation memory within each MLRA and updating the modeling framework when appropriate.

We also plan on extending the paper by replicating the study using modern machine learning methods and comparing the results. While our average planted acreage response statistics conform with our expectations, a number of MLRAs had negative own-priced and positive cross-price elasticities of acreage supply. These results persist despite the fact that each of the included MLRAs had over 30,000 observations that we used to estimate the models. While MLRAs plausibly control for a great degree spatial variation, these results suggest that there could be a degree of within-MLRA heterogeneity yet to be controlled for. Many of these MLRAs are quite large, some spanning dozens of counties and handfuls of states. In addition many of the values delineating the MLRAs, particularly those involving weather, land-use, and biological resources are from the 1990s and may no longer be relevant for agricultural production. While other panel estimation procedures are available, due to the size of the dataset, we plan on implementing the framework of Athey and Imbens (2016) and Prest (2017) who have developed a tree-based approaches for estimating causal heterogeneous treatment effects. The potential advantage of these methods is that they are well suited to estimating models with highly non-linear relationships over large datasets.

References

Athey, S., and G. Imbens. 2016. “Recursive Partitioning for Heterogeneous Causal Effects.” *Proceedings of the National Academy of Sciences* 113:7353–7360.

Carter, C.A., G.C. Rausser, and A. Smith. 2016. “Commodity Storage and the Market Effects of Biofuel Policies.” *American Journal of Agricultural Economics*, pp. aaw010.

Claassen, R., C. Langpap, and J. Wu. 2017. “Impacts of Federal Crop Insurance on Land Use and Environmental Quality.” *American Journal of Agricultural Economics* 99:592–613.

Dale, V.H., K.L. Kline, J. Wiens, and J. Fargione. 2010. *Biofuels: Implications for Land Use and Biodiversity*. Ecological Society of America Washington, DC.

Gardner, B.L. 1976. “Futures Prices in Supply Analysis.” *American Journal of Agricultural Economics* 58:81–84.

Haile, M.G., M. Kalkuhl, and J. Braun. 2014. “Inter-and Intra-Seasonal Crop Acreage Response to International Food Prices and Implications of Volatility.” *Agricultural Economics* 45:693–710.

Haile, M.G., M. Kalkuhl, and J. von Braun. 2016. “Worldwide Acreage and Yield Response to International Price Change and Volatility: A Dynamic Panel Data Analysis for Wheat, Rice, Corn, and Soybeans.” In *Food Price Volatility and Its Implications for Food Security and Policy*. Springer, pp. 139–165.

Hendricks, N.P., S. Sinnathamby, K. Douglas-Mankin, A. Smith, D.A. Sumner, and D.H. Earnhart. 2014. “The Environmental Effects of Crop Price Increases: Nitrogen Losses in the US Corn Belt.” *Journal of Environmental Economics and Management* 68(3):507–526.

Hendricks, N.P., A. Smith, and D.A. Sumner. 2014. "Crop Supply Dynamics and the Illusion of Partial Adjustment." *American Journal of Agricultural Economics*, pp. aau024.

Hennessy, D.A. 2006. "On Monoculture and the Structure of Crop Rotations." *American Journal of Agricultural Economics* 88:900–914.

Koutchadé, O.P., A. Carpentier, and F. Femenia. 2018. "Modeling Heterogeneous Farm Responses to European Union Biofuel Support with a Random Parameter Multicrop Model." *American Journal of Agricultural Economics*, pp. .

Lacroix, A., and A. Thomas. 2011. "Estimating the Environmental Impact of Land and Production Decisions with Multivariate Selection Rules and Panel Data." *American Journal of Agricultural Economics* 93:784–802.

Langpap, C., and J. Wu. 2011. "Potential Environmental Impacts of Increased Reliance on Corn-Based Bioenergy." *Environmental and Resource Economics* 49:147–171.

Lusk, J.L., and J.D. Anderson. 2004. "Effects of country-of-origin labeling on meat producers and consumers." *Journal of Agricultural and Resource Economics*, pp. 185–205.

Miao, R., M. Khanna, and H. Huang. 2016. "Responsiveness of Crop Yield and Acreage to Prices and Climate." *American Journal of Agricultural Economics* 98:191–211.

Miller, H.J. 2004. "Tobler's First Law and Spatial Analysis." *Annals of the Association of American Geographers* 94:284–289.

Motamed, M., L. McPhail, and R. Williams. 2016. "Corn Area Response to Local Ethanol Markets in the United States: A Grid Cell Level Analysis." *American Journal of Agricultural Economics*, pp. aav095.

NASS, USDA. 2010. "Field Crops: Usual Planting and Harvesting Dates." *USDA National Agricultural Statistics Service, Agriculural Handbook* 628.

Natural Resources Conservation Service. 2001. "National Resources Inventory: Background."

Nerlove, M. 1956. "Estimates of the Elasticities of Supply of Selected Agricultural Commodities." *Journal of Farm Economics* 38:496–509.

NRCS, U. 2006. "Land Resource Regions and Major Land Resource Areas of the United States, the Caribbean, and the Pacific Basin." *US Department of Agriculture Handbook* 296.

Piesse, J., and C. Thirtle. 2009. "Three Bubbles and a Panic: An Explanatory Review of Recent Food Commodity Price Events." *Food policy* 34:119–129.

Prest, B.C. 2017. “Peaking Interest: How Awareness Drives the Effectiveness of Time-of-Use Electricity Pricing.” In *Riding the Energy Cycles, 35th USAEE/IAEE North American Conference, Nov 12-15, 2017*. International Association for Energy Economics.

Roberts, M.J., and W. Schlenker. 2013. “Identifying Supply and Demand Elasticities of Agricultural Commodities: Implications for the US Ethanol Mandate.” *The American Economic Review* 103:2265–2295.

Snipes, C.E. 2005. “Current Agricultural Practices of the Mississippi Delta.”, pp. 1–18.

Soil Survey Staff. 2014. *Keys to Soil Taxonomy*, USDA-Natural Resources Conservation Service, Washington, DC. 12th ed.

Tobler, W.R. 1970. “A Computer Movie Simulating Urban Growth in the Detroit Region.” *Economic geography* 46:234–240.

Woodard, J. 2016. “Ag-Analytics.org Data Warehouse.” Available online at: ag-analytics.org.

Wu, J., R.M. Adams, C.L. Kling, and K. Tanaka. 2004. “From Microlevel Decisions to Landscape Changes: An Assessment of Agricultural Conservation Policies.” *American Journal of Agricultural Economics* 86:26–41.

Yan, L., and D. Roy. 2016. “Conterminous United States crop field size quantification from multi-temporal Landsat data.” *Remote Sensing of Environment* 172:67–86.