ESSAYS ON AGRICULTURAL AND ENVIRONMENTAL ECONOMICS

A DISSERTATION SUBMITTED TO THE GRADUATE DIVISION OF THE
UNIVERSITY OF HAWAI‘I AT MĀNOA IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

IN

ECONOMICS

DECEMBER 2018

By

A. John Woodill

Dissertation Committee:

Michael J. Roberts, Chairperson
PingSun Leung
Stuart T. Nakamoto
Kimberly Burnett
Teresa Molina
Thomas Giambelluca

Keywords: agriculture; climate change; adaptation; crop-switching; pest management; decision analysis

To my parents, who taught me to always keep learning.
To Jenna, who keeps my heart guided while my head is afloat.

Like most others, I was a seeker, a mover, a malcontent, and at times a stupid hell-raiser. I was never idle long enough to do much thinking, but I felt somehow that some of us were making real progress, that we had taken an honest road, and that the best of us would inevitably make it over the top.

At the same time, I shared a dark suspicion that the life we were leading was a lost cause, that we were all actors, kidding ourselves along on a senseless odyssey. It was the tension between these two poles - a restless idealism on one hand and a sense of impending doom on the other - that kept me going.

Hunter S. Thompson, The Rum Diary

ACKNOWLEDGMENTS

I would like to first thank PingSun Leung and Stuart T. Nakamoto for their support, mentorship, and guidance. I never thought I would know so much about coffee in Hawaii and I have enjoyed every moment working with you both. I would also like to sincerely thank Michael J. Roberts for his patience and guidance in helping me navigate your many contributions to climate change research. It has been an enriching experience. Many thanks to my dissertation committee Kimberly Burnett, Teresa Molina, and Thomas Giambelluca. Your comments and suggestions have allowed me to see the bigger picture. Thank you to all my friends and family who have supported and listened through all the ups and downs that this journey was. And last, but certainly not least, to my fellow graduate students, and friends, Jonathan Sweeney and Sean Doyle. The many hours spent thinking about questions, working on problems, and the continued encouragement through it all has been an experience like no other.

ABSTRACT

This dissertation offers three essays on environmental and agricultural economics. The first chapter examines how climate change adaptation through crop-choice decisions affect agricultural outcomes through a well-known microeconomic principle, the envelope theorem. In short, the envelope theorem suggests that a first-order approximation without adaptation provides a reasonable estimate of nonlinear effects with adaptation. That is, decisions by farmers in the long-run are already optimized, so accounting for adaptation can be safely ignored; thus, adaptation is a second-order effect of extreme events and nonlinearities. A simple model is provided to show that the envelope theorem holds across continuous and discrete crop switching decisions. An empirical analysis is then carried out to show that nonlinear effects without adaptation have a negative impact on revenue-per-acre as temperatures increase, but the effects are slightly reduced when accounting for adaptation. The analysis also shows that a marginal approximation without adaptation provides a reasonable estimate of nonlinear effects with adaptation. The results suggest the effects of climate change on agriculture can be simplified by eliminating the necessity of accounting for adaptation in order to get useful estimates.

The second chapter extends beyond the first chapter by providing a more rigorous analysis of crop choices and agricultural outcomes. A long history of crop choice and productivity outcomes are used to estimate the effects of both weather and climate change on major field crops in the United States. Climates are defined by backward-looking, rolling means of weather measures, with the lag length selected through cross-validation. The effect of climate change on crop revenue-per-acre is estimated from a base level to uniform increases in temperature of +0 to +5 Celsius. The results show that adaptation by crop-switching slightly reduces the impact of climate change relative to estimates that consider weather alone. Therefore, nonlinear responses due to climate change adaptation provide a slight improvement to outcomes but are unable to completely mitigate the effects of climate change.

The final chapter diverges from the theme of the first two chapters and considers optimal decisions Hawaii coffee farmers make to combat damage from a new invasive species, the coffee berry borer. The effect of a decision to spray or not spray a biological insecticide, *Beauveria bassiana*, is estimated based on the expected damage from not spraying versus the cost to spray. If damages are greater than the cost to spray, then it is beneficial to spray in order to mitigate damage to coffee. A Markov-chain is used to estimate economic damage in each month based on a spray decision. The Markov-chain is incorporated into a dynamic programming model to optimize the net-benefit during a coffee growing season. Results provide an optimal decision path for a coffee growing season and an optimal final net-benefit. Next, the economic model is compared against alternative pest management strategies: IPM, decisions to always spraying or to never spray. Results show our economic model performs best when optimizing net-benefit for a typical farm in Kona Hawaii.

CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1
ADAPTATION AND THE ENVELOPE THEOREM†

1.1 Introduction

A common theme in the economics of climate change is that farmers will adapt by planting different crops that are less sensitive to extreme heat or by planting a variety of the same crops to offset the impact on yields — this adaptation is termed crop-switching. Estimating the effectiveness of crop-switching involves understanding how farmers' crop choice adjustments are able to mitigate damages from climate change. However, farmers observe climate through long-run changes in weather, so it is reasonable to assume they make small and gradual adjustments that will have marginal effects on the value of farm-level activity and crop yields. I show in this paper that the marginal effect of adaptation in the long-run equals zero – a result provided by the envelope theorem – and that we can accurately approximate climate change impacts by ignoring adaptation altogether.

Identifying the effect of adaptation is difficult because other factors can also affect farmers' decisions, such as changes in prices, government regulations, and technological improvements. Much of the literature on climate change adaptation has focused on understanding how farmers have adjusted to climate historically while holding all other factors constant. One approach, known as the hedonic approach, utilizes cross-sectional associations to account for adaptation implicitly in farm-level outcomes. A seminal study that utilizes this approach suggests adaptation can completely mitigate climate change impacts on farmland values (Mendelsohn et al. 1994), however, other studies that utilize the same approach suggest large negative effects (Schlenker et al. 2005, 2006). Another approach ignores adaptation implicit in the hedonic approach by estimating the direct impact of year-to-year weather variation and also finds large negative declines in agriculture outcomes (Deschênes and Greenstone 2007; Schlenker and Roberts 2009; Fisher et al. 2012; Deschênes and Greenstone 2011; Dell et al. 2012). A more recent approach explicitly estimates adaptation by comparing cross-sectional differences in climate over time (long-differences) to year-to-year weather variation on crop yields and finds no evidence of adaptation (Burke and Emerick 2016). The existing literature provides a variety of claims about the degree to which adaptation is likely to mitigate damages, but taken together, the effectiveness of adaptation appears to be limited.

In this paper, I use the envelope theorem to show that accounting for adaptation, to a first-order approximation, is not necessary for estimating the impact of climate change on agriculture. The envelope theorem states that exogenous changes – from weather or climate, for example – to an optimal objective function, such as profits, return a direct effect (without adaptation) and an indirect effect (with adaptation) where the indirect effect equals zero along the margins. In short,

†This chapter is a result of collaboration with Michael J. Roberts.

beliefs about climate have no effect on the decisions that farmers make because their decisions are already optimized. The theorem suggests that, to a first-order approximation, adaptation (indirect effect) can be safely ignored. As a result, an approximation of the direct effect from climate change is only needed to estimate the impacts of climate change.

Cross-sectional studies that claim adaptation will completely mitigate the effects of climate change on agriculture are in direct contrast to the envelope theorem. These studies account for adaptation implicitly, so there is no direct measure of adaptation on outcomes. Further, for their claims to hold, crop choice decisions (indirect effect) would have to be large enough to overcome the large negative direct effects that have been reported (Schlenker and Roberts 2009). Thus, crop choice decisions would have to be very sensitive to climate in order to counter the direct effects.

This paper makes a number of contributions to the literature. First, I resolve an issue that suggests the envelope theorem does not hold for discrete choices at pivotal climate points – like switching entirely to less heat sensitive crops – that provide benefits under adaptation (Guo and Costello 2013). I show that it is important to consider a continuum of climates across outcomes and that the envelope theorem holds at those pivotal climate points. Integrating marginal aggregate effects under the distribution of climates brings this work back to the original envelope theorem.

Next, I use a long history of agriculture outcomes and weather data in the U.S. ranging from 1950 to 2010 to provide an empirical exercise that tests the envelope theorem. The empirical exercise uses a "binning" method to estimate revenue-per-acre – a measure used to estimate the value of activity – for five main crops in the U.S.: corn, cotton, hay, wheat, and soybean. The measure is applied across three weather variables: average temperatures, degree days between 10°C and 30°C, and degree days greater than 30°C. Adaptation is modeled as crop choice responses across a continuum of weather variables. To simulate climate change, uniform increases in temperature from +1°C to +5°C are simulated for each weather variable. Next, nonlinear effects are estimated from climate change with and without adaptation by crop-switching. I then explicitly apply the envelope theorem to estimate the first-order approximations without adaptation by crop-switching.

The main results show that nonlinear responses without adaptation are most impacted by increases in temperature, but nonlinear effects with adaptation only slightly improve outcomes due to climate change. Nonlinear indirect effects can overcome some of the damages from climate change, but not all of them. Applying the envelope theorem shows that the first-order approximation of direct effects provides a similar estimate of nonlinear effects with adaptation. Therefore, if nonlinear effects are small and countervailing, then a first-order approximation of direct effects provides a reasonable assessment of climate change impacts, thus adaptation can be safely ignored.

The remainder of this paper is organized as follows. Section 1.2 outlines a theoretical framework and clarifies the envelope theorem. Section 1.3 describes the empirical analysis. Section 1.4 provides the results. Section 1.5 offers a discussion and Section 1.6 concludes the study.

1.2 Theoretical Framework

Adaptation is thought to be an important component in dealing with climate change and mitigating future damages. The envelope theorem provides some useful insight: it implies that, on the margin, we can safely ignore changes in behavior in response to changes in climate. Adaptation, like a nonlinear response, matters only for substantial changes. Thus, adaptation is a second-order effect. In this section, I show how this result also generalizes to discrete and continuous choices of crops.

Assuming continuous choice sets, differential calculus and the chain rule imply that marginal changes in exogenous factors like climate, affect an objective through direct and indirect mechanisms. The indirect mechanisms account for changes in farm-level decisions as a result of climate. However, the indirect effect, or adaptation, is unable to affect the objective because the objective is already optimized. Since decisions are already optimized, changes in farm-level decisions are unable to improve production activity. As a result, changes to the objective only occur through the direct effects of climate change, such as short-run effects through weather realizations.

As a concrete example, suppose a farmer wishes to maximize profit by selecting a vector of continuous actions, x, which may include crop choice, seed varieties, planting dates, and other activities. Profit depends on both these choices and the climate, summarized by a vector of exogenous measures c.

An optimized objective function can be defined as the value of activity as a function of climate, $y(c)$, equal to optimal crop choice decisions, x, and climate, c:

$$y(c) = \max_x f(x, c)$$

Maximized optimal crop decisions are defined as $x^*(c)$, which solves the implicit functions given by $\frac{\partial f(x^*,c)}{\partial x} = 0$. These exist under the usual assumptions about convex production sets. The objective can then be written as:

$$y(c) = \max_x f(x^*(c), c) \tag{1.1}$$

A small change in climate will have a marginal effect on the crop choice and has both a direct effect ($\frac{\partial f}{\partial c}$) and an indirect effect ($\frac{\partial f}{\partial x}\frac{\partial x}{\partial c}$). Adaptation, as defined by much of the literature, concerns the indirect effect. Deriving the first-order conditions involves simply finding the effect of climate on the value of activity, $\frac{\partial y}{\partial c}$.

$$\frac{\partial y}{\partial c} = \underbrace{\frac{\partial f}{\partial x}\frac{\partial x}{\partial c}}_{\text{Adaptation}} + \underbrace{\frac{\partial f}{\partial c}}_{\text{w/o Adaptation}} \tag{1.2}$$

The envelope theorem states that, at the margin, we need only consider direct effects because the indirect adaptive effect is zero due to the first order condition ($\frac{\partial f(x*,c)}{\partial x} = 0$). Indirect effects – the effect of a behavioral change in response to climate – drop out because, at the margin, the behavior is already optimized. There is zero marginal gain from slightly adjusting choices in response to a small change in climate. In other words, in a first approximation, adaptation can be ignored.

It is useful to note that the envelope theorem result applies to outcomes besides firm profits, such as productivity, total output, or other consumer-side outcomes since decisions underlying those outcomes are presumably optimized with regards to climate. It can also be observed that empirical studies that link outcomes to weather do account for some kind of implicit effects that filter through decisions – decisions that respond to weather after the fact or to meaningful forecasts of weather. The conventional notion of adaptation, therefore, mirrors the Le Chatelier principle, or the idea that long-run responses to price or other exogenous factor tend to be greater than short-run responses since more decisions can be adjusted (Silberberg 1971). In this context, the Le Chatelier principle clarifies the difference between weather and climate, not as a metric – a prevailing average or deviation from the norm – but as an issue that pertains to expectations about the weather and the timing of when different decisions are made.

Another corollary not clearly addressed within the literature concerns discrete choice and aggregation. Guo and Costello (2013) argue that the envelope theorem holds for continuous choices but not for discrete choices. Discrete choices involve completely switching to an alternative crop at pivotal climate points, while continuous choices allow a variety of crop choices. At points where the climate is not pivotal to a discrete choice, the usual envelope theorem follows (see Figure 1.1). The issue concerns pivotal climates demarking one discrete alternative from another. An example of such a choice may concern whether to use a parcel of land for forest or crops, or whether to plant soybean or wheat. At the pivotal climate, the first order condition ($\frac{\partial f(x*,t)}{\partial x} = 0$) is not defined. All choices, even continuous ones, can change discontinuously at this pivotal climate. As a result, Guo and Costello show that the envelope theorem provides little guidance in this situation.

There are two counterpoints to Guo and Costello that are elaborated upon here. First, the profitability of the two discrete alternatives is equal at the pivotal climate, and the decision, therefore, has zero impact on profit at the margin (it will, however, likely have discontinuous effects on input use, output, and productivity.) Thus, at least concerning the profit outcomes emphasized by many, the envelope theorem is still valid. Second, and more generally, it is important to recognize that there is a continuum of climate that must be considered and aggregated, represented by the smooth distribution of $\phi(t)$ in Figure 1.1 adapted from Guo and Costello 2013. Pivotal climates – a couple of points in Figure 1.1 that account for discrete changes in crop choice – account for the zero measure of this continuum. The marginal aggregate effect would integrate marginal effects under the whole distribution of climates, and since climate pivotal for discrete choices has a zero measure, these can be ignored. Thus, bringing back the original envelope theorem. If one defines the distribution of climate as $\phi(c)$ and $c \in c^*$ as the finite set of pivotal climates, the impact of

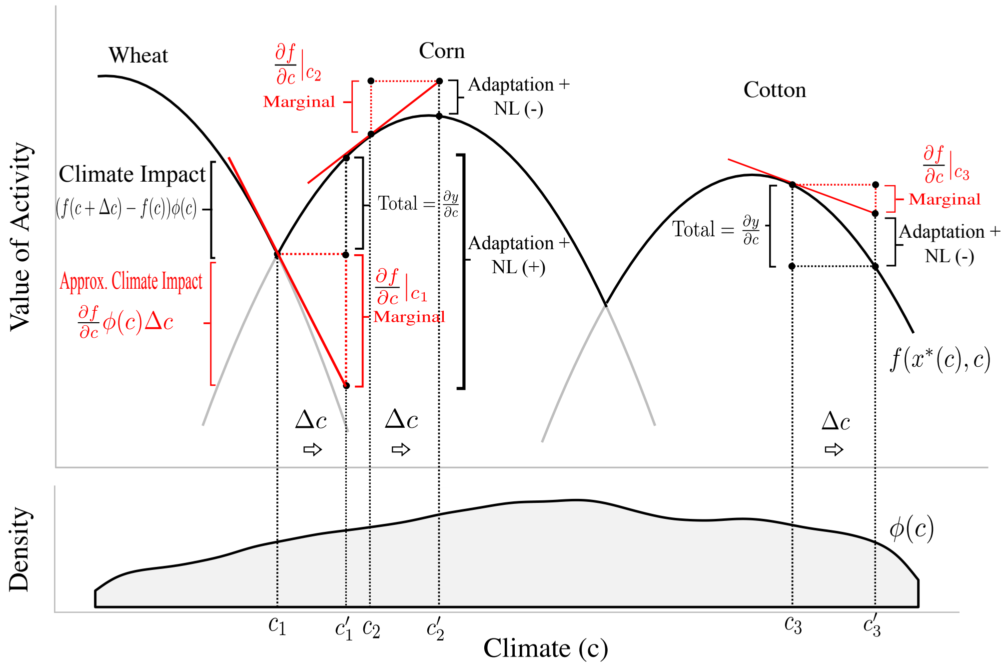

Figure 1.1: **Total Climate and Approximate Climate Impacts**

Notes: Figure illustrates measuring adaptation and the envelope theorem through a change in climate. The figure provides production functions for wheat, corn, and cotton and the value of activity (revenue-per-acre). Pivotal climate points are where production functions cross. A distribution of climate, $\phi(c)$, is provided in the bottom panel to account for climate across a variety of crop outcomes. Climate impacts are assessed with a change in climate Δc. The total effect from Δc is calculated from nonlinear responses. Approximate climate impacts (red) are estimated using a first-order approximation (marginal). The envelope theorem says that approximate climate impacts with adaptation can provide a reasonable estimate of climate impacts with adaptation.

climate change can be formalized as:

$$\int_{c \notin c*} (y(c + \Delta c) - y(c))\phi(c)dc + \sum_{c \in c*} Prob(c)(y(c + \Delta c) - y(c))$$

Under a continuous distribution, the probability of an exact point equals zero[1]. Therefore, at pivotal climate points $\sum_{c \in c*} Prob(c) = 0$, one can approximate the impact of climate change as:

$$\approx \int_{c \notin c*} \frac{\partial y(c)}{\partial c}\phi(c)dc = \int_{c \notin c*} \frac{\partial f(x,c)}{\partial c}\phi(c)dc \tag{1.3}$$

Finally, note that second-order adaptive effects can be negative as well as positive. If profit, production, or other outcomes of interest have a concave relationship with the weather, as is typically illustrated and empirically validated, then extrapolation of marginal relationships will tend to give optimistic projections concerning discrete climate changes, as is illustrated for climates c_2 and c_3 in Figure 1.1. In reality, profit functions may have both convex and concave regions, and a first-order aggregated response may over or understate impacts inclusive of adaptation. Accounting for nonlinearity may matter as much as or more than adaptation.

Therefore, even under discrete choices, the envelope theorem holds. Changes in profit at pivotal climate points is equal to zero on the margin because the changes in profit are equal at these points. As Guo and Costello (2013) show, continuous choices equal zero along the margins, thus establishing the envelope theorem. Here, it is further shown that when a distribution of climate across a variety of crops is considered, the aggregate marginal effect also equals zero. Because, at discrete choices or pivotal climate points, the marginal change in profit equals zero, the envelope theorem still holds. This result suggests that crop switching provides zero adaptation benefit to a first-order approximation when considering ways to combat changes in climate.

1.3 Empirical Analysis

To illustrate the envelope theorem and adaptation from crop switching, an empirical exercise is provided that utilizes long-term data on crop outcomes and weather ranging from 1950 to 2010. A "binning" method is used to estimate crop revenue and crop choices for the five main crops in the US (corn, cotton, hay, wheat, and soybean) across a variety of weather variables. Three equations are estimated from the data due to a change in climate: (1) nonlinear responses with adaptation, (2) nonlinear responses without adaptation, and (3) marginal effect without adaptation that applies the envelope theorem. These three are compared to estimate the nonlinear responses with and without adaptation through crop switching and the effects of climate change. Marginal

[1]With a continuous random variable, a probability $\phi(c)$ exists between intervals, but at a single climate point, the probability is equal to zero. In other words, as the exact point on each side of the density function is approached, the probability decreases to zero. In the context of climate, the probability of a climate event happening along the envelope at those pivotal climate events is equal to zero, which holds with the envelope theorem

effects without adaptation are also estimated in order to apply the envelope theorem and show that it provides a close approximation of nonlinear effects with adaptation.

1.3.1 Estimation

The empirical analysis in this paper relies on estimating the objective function in equations 1.1 and 1.2 explicitly. In particular, the goal is to optimize an objective function where $y(c)$ is revenue-per-acre as a function, $f()$, that includes a measure of climate, c, and optimal crop choices, $x^*(c)$. A continuum of climate, $\phi(c)$, is also accounted for using average temperatures and degree days between 10°C and 30°C, and degree days above 30°C. Utilizing all parts of the equation will allow the nonlinear and approximate effects of a change in climate to be estimated.

Figure 1.1 provides a visual representation of each part of the empirical analysis and shows how to estimate nonlinear climate impacts and approximate climate impacts. The figure shows production functions, $f()$, for wheat, corn, and cotton and their value of activity, y. Along the horizontal axis is a measure of climate, c, with the density, $\phi(c)$, provided in the bottom portion of the plot. To account for crop choices, crop acre choices across a continuum of climates are represented as $x(c)$ (not shown in the figure).

To estimate changes in climate, assume there is a shift in climate temperatures from c_1 to c_1', also represented by Δc. The nonlinear climate impact is calculated as

$$\frac{\partial y}{\partial c} = (f(c + \Delta c)\phi(c + \Delta c) - f(c))\phi(c)),$$

The total nonlinear effect, $\frac{\partial y}{\partial c}$, accounts for the nonlinear effect on revenue-per-acre due to changes in climate, holding all else constant. The nonlinear effect can also account for adaptation by including the choice of changing crop acres, $x(c + \Delta c)$, or without adaptation by keeping acres fixed, $x(c)$. The total aggregate effect integrated at each climate point c across all climates C with and without adaptation is calculated as

Nonlinear effects with adaptation

$$\int_{c \in C} f(c + \Delta c, x(c + \Delta c)) \cdot \phi(c + \Delta c)dc - f(c, x(c)) \cdot \phi(c)dc \tag{1.4}$$

Nonlinear effects without Adaptation

$$\int_{c \in C} f(c + \Delta c, x(c)) \cdot \phi(c + \Delta c)dc - f(c, x(c)) \cdot \phi(c)dc. \tag{1.5}$$

If the envelope theorem holds, then a first-order approximation should provide a reasonable estimate of nonlinear effects with adaptation. To calculate the approximate climate impact, the first-order approximation is calculated at c_1 as,

$$\frac{\partial f}{\partial c}|_{c_1} = \frac{\partial f}{\partial c}\phi(c)\Delta c,$$

where $\frac{\partial f}{\partial c}$ represents a first-order approximation. The total aggregate effect integrated across all climates C is calculated as

First-order approximation without adaptation

$$\int_{c \in C} f'(c, x(c))dc \cdot \phi(c) \cdot \Delta c. \tag{1.6}$$

The empirical analysis in this paper directly estimates equations 1.4, 1.5, and 1.6 using a long history of agricultural data. To estimate each of the equations, revenue-per-acre, y, and crop choice, $x(c)$, functions are derived for all five crops: corn, cotton, hay, wheat, and soybean. The goal for deriving the equations is to produce functions similar to Figure 1.1 for each of the components needed to estimate the equations.

For each crop, individual functions are approximated for revenue-per-acre and crop choices. A "binning" method is used to approximate the functional forms, in which the measure of climate is put into a number of bins, averaged in each bin, and then linearly interpolated between points. For example, to estimate corn revenue-per-acre across average temperatures from 0-30°C, temperatures are split into bins of three degrees, for a total of 10 bins. Within each bin, the average revenue-per-acre is calculated and then connected using linear interpolation. The revenue-per-acre curve for corn should be an inverted-U with lower and higher temperatures producing less revenue-per-acre.

Figures 1.2, 1.3, and 1.4 provide the results of the binning method for average temperatures, degree days 10°C-30°C, and degree days greater than 30°C. The top panel provides revenue-per-acre and the bottom panel provides the proportion of crop acres. Ten bins were used in the main analysis to ensure fewer kinks when linearly interpolating. A smoother functional form ensures that results are not prone to significant kinks in revenue-per-acre and proportion of crop acres. Robustness checks with different bins are provided in the discussion below.

1.3.2 Data

Data for the empirical analysis comes from the National Agricultural Statistics Service (NASS) released by the United States Department of Agriculture (USDA). The data reports county-level estimates for production and acreage for 1,923 counties from 1950-2010. The crops used in this analysis include corn, cotton, hay, wheat, and soybean, which make up the majority of crops in the U.S. (57% of total U.S. agricultural production). The sample retains counties east of the 100th-degree meridian that relies on precipitation rather than irrigation. Counties that report at least one observation for crop acres in 1950 are also included in the sample. To simplify the "binning" method, I aggregate to county-level observations, thus removing the time component. The data, as with previous cross-sectional association problems, cannot account for unobserved heterogeneity.

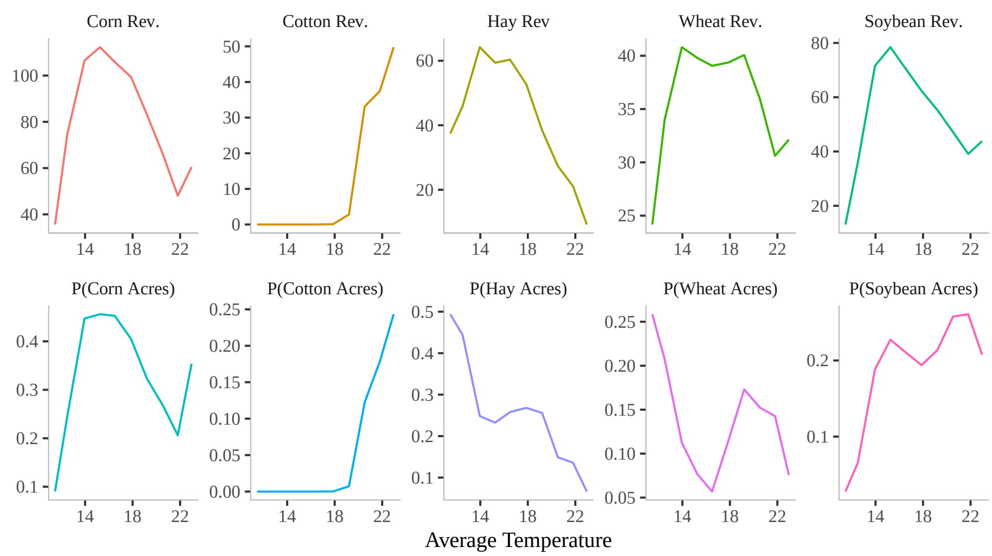

Figure 1.2: **Approximate Functions for Revenue-per-acre and Proportion of Acres under Average Temperatures**

Notes: Figure provides revenue-per-acre and proportion of crop acres for five main crops corn, cotton, hay, wheat, and soybean. Functions are approximated by "binning" average temperatures and averaging revenue and crop acres in each bin. Points are linearly interpolated to estimate an approximate functional form for each crop. Revenue-per-acre is calculated using a fixed-price times yield. The proportion of crop acres is calculated by dividing crop acres by total crop acres.

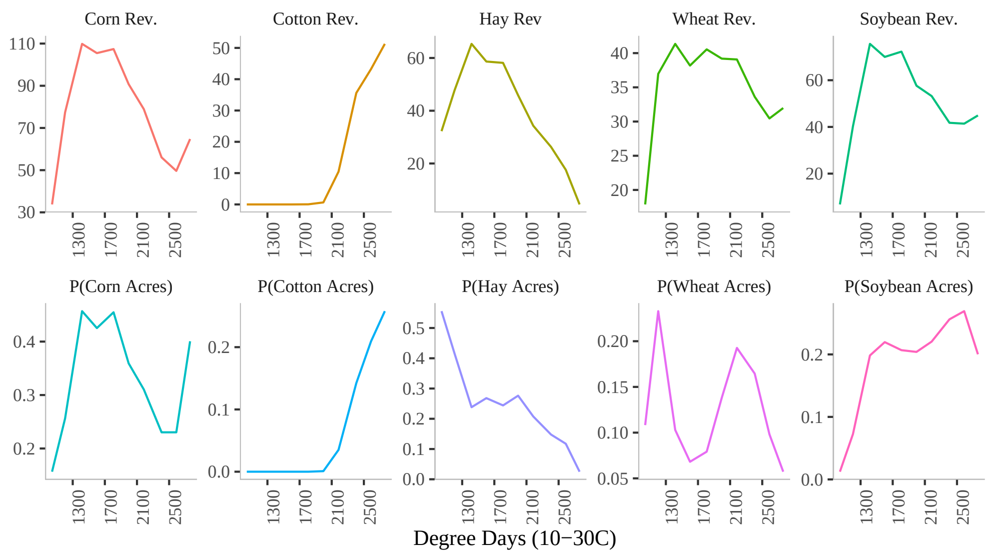

Figure 1.3: **Approximate Functions for Revenue-per-acre and Proportion of Acres under Degree Days 10-30°C**

Notes: Figure provides revenue-per-acre and proportion of crop acres for five main crops corn, cotton, hay, wheat, and soybean. Functions are approximated by "binning" degree days between 10 and 30°C and averaging revenue and crop acres in each bin. Points are linearly interpolated to estimate an approximate functional form for each crop. Revenue-per-acre is calculated using a fixed-price times yield. The proportion of crop acres is calculated by dividing crop acres by total crop acres.

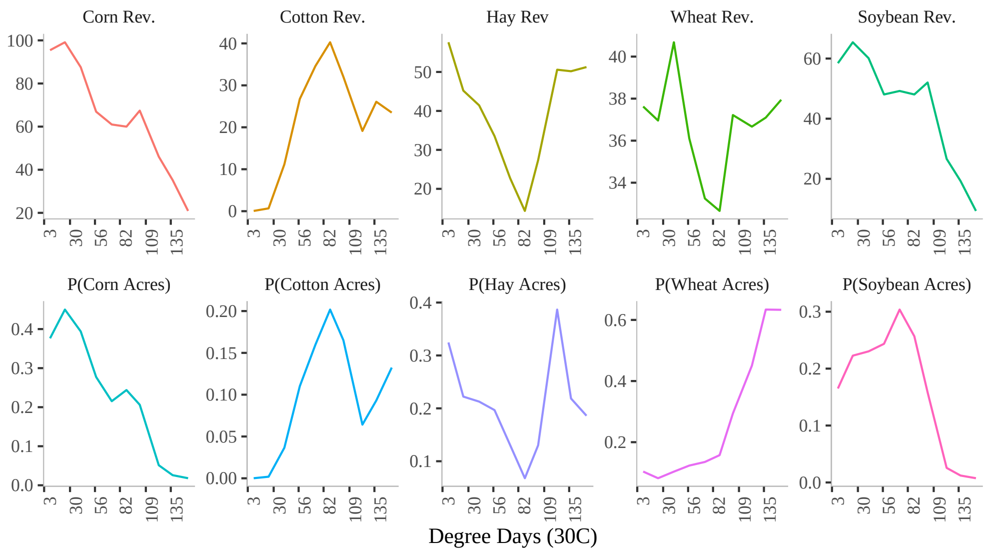

Figure 1.4: **Approximate Functions for Revenue-per-acre and Proportion of Acres under Degree Days 30°C**

Notes: Figure provides revenue-per-acre and proportion of crop acres for five main crops corn, cotton, hay, wheat, and soybean. Functions are approximated by "binning" degree days greater than 30°C and averaging revenue and crop acres in each bin. Points are linearly interpolated to estimate an approximate functional form for each crop. Revenue-per-acre is calculated using a fixed-price times yield. The proportion of crop acres is calculated by dividing crop acres by total crop acres.

As derived, the envelope theorem pertains to profits. However, it can also pertain to other outcomes, including revenue-per-acre. Presumably, optimal decisions underlie profits in response to climate, so optimal decisions also underlie revenue-per-acre. The envelope theorem implies decisions are fixed at the margin because changes in optimized decisions are small relative to changes in the exogenous variable. If decisions are assumed fixed at the margin, then the envelope theorem can apply to revenue-per-acre just as it does to profit.

Crop revenue-per-acre is used as the primary outcome variable because it is a simple way to value crops. When aggregating the value of crops, this value provides a way to allow complementarities between crops to affect productivity. For example, crop rotations can affect pest management and fertilizer costs or improve time management since different crops can have different planting and harvesting times (Cai et al. 2013; Livingston et al. 2014).

To construct revenue-per-acre, I fix state-level prices of each crop at the average over all years. Prices are adjusted for inflation using the GDP deflator with the base year of 2010. Fixing prices resolves problems with endogeneity, such as a storage and weather-induced price shocks (Fisher et al. 2012). Prices are also volatile from year-to-year, so one year may not reflect the appropriate terms of trade between two crops. Averaging prices over the time series resolves the volatility in prices and follows that crop prices tend to move together over the long-run (Sumner 2009; Roberts and Schlenker 2013).

To account for changes in acreage choices, I estimate the proportion of crop shares in each county. Crop shares are estimated using individual crop acres divided by total crop acreage in that county. If a county does not report a crop acre, a value of zero is reported. Crop shares are reported as a proportion of total shares between zero and one.

Historical temperature data comes from interpolation techniques using a relative anomaly spline approach[2]. I utilize three weather variables for this analysis: (1) average temperatures, (2) degree days between 10°C and 30°C and (3) degree days greater than 30°C. Average temperatures are simply the average temperature during the growing season from March until October. To include nonlinear temperature effects, I also include degree days that account for growing degree days (10°C-30°C) and extreme heat degree days (30°C). Degree days account for the rise and fall of temperatures during the day. Degree days has been shown to provide a better estimate of the temperature effects on agricultural outcomes.

1.4 Results

Figure 1.5 provides the main results for average temperatures. Equations 1.4, 1.5, and 1.6 are estimated with a uniform change in temperatures from +1°C to +5°C. The figure plots the percentage change in revenue-per-acre from the baseline of a zero change in temperature (+0°C).

[2] Documentation for developing fine-scale weather data and calculating degree days is available at https://github.com/johnwoodill/Fine-Scale-Weather-Interpolation

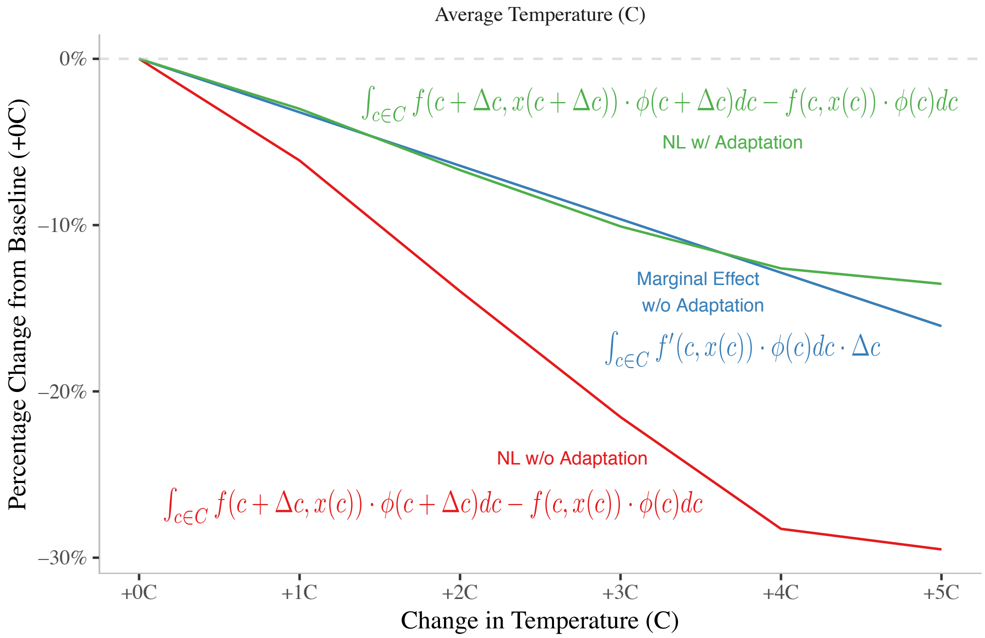

Figure 1.5: **Nonlinear and First-order Approximation for Average Temperatures**

Notes: Figure provides nonlinear (NL) with and without adaptation and first-order approximations (marginal) without adaptation for average temperatures. Nonlinear without adaptation are worse than with adaptation but is unable to completely mitigate the effects of climate change. Marginal effects without adaptation provide a close approximation of nonlinear with adaptation.

These results show that nonlinear effects without adaptation are greater as temperatures increase. If farmers are not able to adapt to increases in temperature of $+2$°C – which the IPCC suggests could be reached by 2050 if emissions remain unabated – then effects on revenue-per-acre could see as much as a 13% decline. If farmers can adapt through crop switching, the nonlinear effects improve slightly, but declines in revenue-per-acre will still occur at a 9% loss. This loss suggests that adaptation through crop switching is unable to completely mitigate the effects of climate change on agriculture.

Additionally, as the envelope theorem describes, the marginal effect without adaptation provides a reasonable approximation of the nonlinear effects with adaptation. Revenue-per-acre to a first order approximation at a $+2$°C temperature change is expected to see declines of 7% (compared to declines of 9% for nonlinear effects with adaptation). Therefore, when estimating the effects of climate change, adaptation can be safely ignored to a first-order approximation.

Figure 1.6 provides results for the other measures of climate in this analysis: degree days of 10°C-30°C and degree days greater than 30°C. Similar results are produced when applied to additional measures of weather – to a first-order approximation, marginal effects without adaptation provide

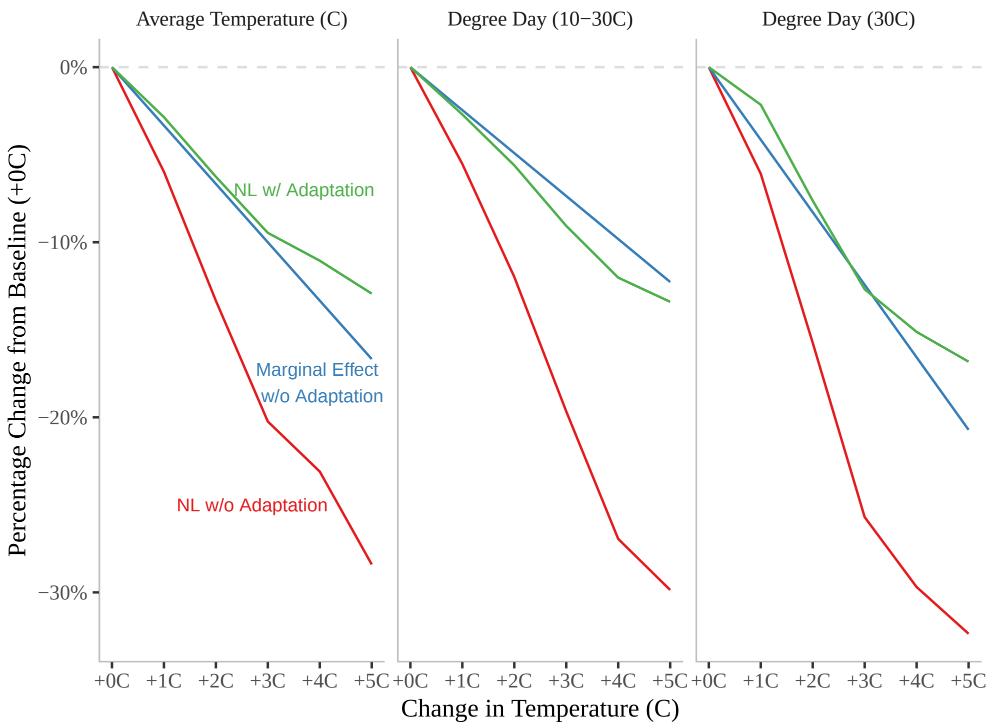

Figure 1.6: **Nonlinear and First-order Approximation for All Measures of Climate**

Notes: Figure provides nonlinear (NL) with and without adaptation and first-order approximations (marginal) without adaptation for average temperatures, degree days 10-30°C, and degree days greater than 30°C. Nonlinear without adaptation are worse than with adaptation, but is unable to completely mitigate the effects of climate change. Marginal effects without adaptation provide a close approximation of nonlinear with adaptation.

a close approximation of nonlinear effects with adaptation.

However, declines in revenue-per-acre with a +2°C increase in temperature are greater for degree day estimates. Degree days have been shown to provide more robust estimates for measuring climate change impacts, especially degree days greater than 30°C (Schlenker et al. 2006; Schlenker and Roberts 2009). Climate change impacts for a +2°C change in temperature for degree days greater than 30°C without adaptation is expected to include a loss of 32% in revenue-per-acre. When accounting for adaptation, to a first-order approximation, losses are expected to be 8%.

These main results suggest significant losses in revenue-per-acre without adaptation, but when adapting by crop switching, positive benefits from adaptation exist. Further, a first-order approximation provides a reasonable assessment of climate change impacts. Thus, it is not necessary to account for adaptation to approximate climate change impacts.

1.5 Discussion

The envelope theorem provides important insight into the effects of climate change. If environmental change is small and gradual, then we expect the net effects of adaptation to be zero along the envelope. At a first-order approximation, marginal benefits from adaptation are met with equal marginal costs of adaptation. As a result, measuring adaptation is not necessary for understanding climate change damages. However, at a second-order approximation, or sudden and drastic changes in the environment, we should be concerned with measuring both adaptation and nonlinearities to understand the full implications of those damages (Nordhaus 2010). Second-order effects matter, but by how much is an empirical question. Unfortunately, we don't yet know how big second-order effects from adaptation are, so first-order approximations without adaptation can provide accurate estimates.

The analysis in this paper focuses on this insight by estimating and comparing nonlinearities with adaptation and a first-order approximation to show that the envelope theorem holds. Further, benefits exist due to nonlinearities and adaptation. However, it is important to note that this analysis only focuses on changes in temperature while holding all else equal. Climate change is expected to impact other factors in the environment, such as precipitation, humidity, vapor pressure deficit, and carbon dioxide concentrations. Agriculture policy may also change. Therefore, while this paper does not take these factors into account, it is important to consider them when accounting for the full impacts of climate change on agriculture.

As mentioned previously, when using the "binning" method, the number of bins matters. Too many bins result in sharp discontinuities and too few bins are unable to capture the changes between bins of temperature, which will bias the results. The analysis in this paper used 10 bins, which provided the best fit across all temperature variables for revenue-per-acre and proportion of crops. Figure 1.7 provides robustness checks for different bin lengths from five to 30. The results remain stable across bins, except for degree days greater than 30°C for 20 and 30 bins. Nonlinear responses remain stable. The reason for the large positive benefits for degree days greater than 30°C at a first-order approximation is fewer observations on the upper bound. Sharp discontinuities in the data result in first-order approximations overstating the benefits.

Figure 1.8 shows the sharp discontinuities for interpolation using 30 bins for revenue-per-acre. While most of the crops share similar patterns for average temperature and degree day 10°C-30°C, hay and wheat have differences on the upper ranges of the binning method for degree days greater than 30°C. Due to a low number of observations in the upper range, marginal effects will overstate the results, thus providing positive benefits where negative benefits should exist and driving a large increase in revenue-per-acre at a first-order approximation. Accounting for average temperatures and degree days 10°C-30°C gives results that are robust across bins and temperatures and further support the results of this paper.

Evidence supporting the envelope theorem exists throughout the literature. Butler and Huybers

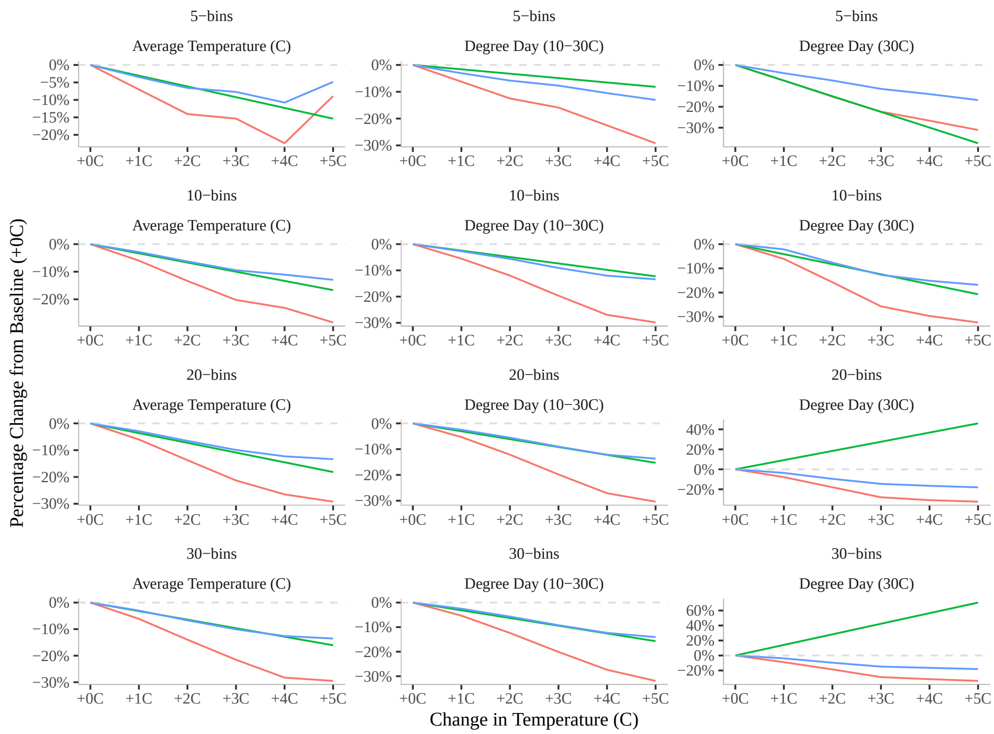

Figure 1.7: **Robustness Check on Nonlinear and First-order Approximation for All Measures of Climate**

Notes: Figure provides robustness checks of main results of the analysis. The green line represents first-order approximations (marginal), the blue line represents nonlinear effects with adaptation, and the red line represents nonlinear effects without adaptation. Bins from 5-30 are used to check the robustness of the results. Results are robust across a variety of bins, except for degree days greater than 30°C. The inconsistency for degree days greater than 30°C is due to sharp discontinuities from more bins and lower observations on the extreme end of the variable (see 1.8). As a result, marginal effects will overstate benefits.

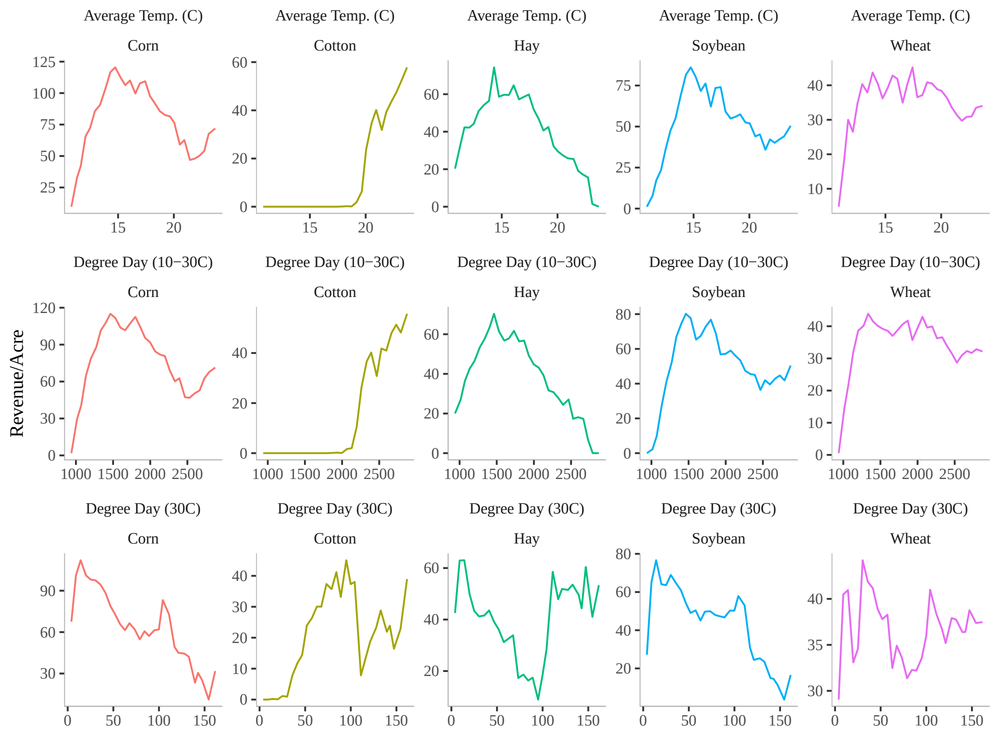

Figure 1.8: **Revenue-per-acre with 30-bins**

Notes: Figure provides revenue-per-acre with 30-bins. Degree days greater than 30°C are less robust than other measures when changing bins due to sharp discontinuities and lower observations on the extreme end. As a result, marginal effects will overstate benefits.

(2013) suggest that adaptation is transformation and higher crop sensitivity is costless. However, in a reply by Schlenker et al. 2013, the authors address this finding by showing that if the costs were zero, then farmers would already be adapting. They show that adaptation is not costless and the marginal benefit gained from adapting is met with equal marginal costs; thus, adaptation equals zero. The authors conclude that adaptation is a second-order effect.

In a paper by Lobell et al. 2014, the authors find that farmers are adapting by increasing sowing densities for a marginal benefit to production, but at an additional cost due to an increase in vapor pressure deficit (VPD) sensitivity. The results suggest the benefits from sowing densities are met with equal costs from VPD sensitivity to crop production.

There is also evidence in the literature that suggests that adaptation can be beneficial through crop choice (Kurukulasuriya et al. 2008; Seo and Mendelsohn 2008). These studies model crop choice as a discrete choice through a selection process (Heckman) where the outcomes are probabilities of selecting a variety of crops. However, probabilities are not observed events[3], so extrapolating results based on probabilities are not accurate interpretations. Further, these studies rely on cross-section associations, which are unable to control for unobserved heterogeneity. These studies are also not able to control for correlation between crop choices – an increase in corn acres implies other crop acres will decrease due to limited farmland acres – which may also produce biased estimates.

Future research can address these limitations by modeling crop choices as crop shares, or a proportion of total acres. This transformation ensures the interpretation of estimates and predictions provided changes in crop proportions. Another critical decision is to ensure that correlations between crop choices are incorporated into the model to reduce biased estimates. These limitations can be addressed using a system of equations (seemingly unrelated regression) where each equation is a specific crop proportion regressed on weather and climate variation (Wooldridge 2010). This setup also allows control of unobserved heterogeneity since the individual models are OLS with adjusted coefficients. A long panel of crop production outcomes using this setup would provide a more accurate depiction of how crop choices have changed due to weather and climate.

1.6 Conclusion

How farmers will adapt to climate change and what the economic impact of those decisions will be is an ongoing debate in agricultural economics. Much of the literature has focused on estimating economic impacts to determine how choices have adjusted historically while holding all else fixed. The literature does not provide consistent results across empirical strategies, and no evidence exists that suggests that adaptation will completely mitigate the damages from climate change.

The envelope theorem tells us that exogenous changes, such as weather or climate, to outcomes are unable to affect behavioral changes (indirect effect) because the behavior is already optimized.

[3]For example, if you flip a coin five times and heads turn up three times, the proportion of heads is 0.60, but the probability of heads is still 0.50

Therefore, to a first-order approximation, adaptation can be ignored. When considering adaptation, second-order effects such as extreme environmental changes matter, but first-order approximations without adaptation can provide reasonable estimates of nonlinear effects with adaptation.

In this paper, I first show the envelope theorem holds at discrete pivotal climate points where behavioral changes equal zero – where Guo and Costello (2013) suggested there were benefits from adaptation. I show this formally where the sum of pivotal climate points equals zero when considering a distribution of climates. The aggregate marginal effect integrated along the whole distribution of climates equals zero, thus returning to the envelope theorem.

I next provide an empirical analysis using a long history of agriculture outcomes from 1950 to 2010 to show that nonlinear effects without adaptation negatively affect revenue-per-acre. Accounting for nonlinearities with adaptation slightly improve impacts, but adaptation is unable to mitigate the effects of climate change completely. Finally, I show that the envelope theorem holds where a first-order approximation provides a reasonable estimate of nonlinearities plus adaptation. This result holds across different temperature variables and bins.

To conclude, the envelope theorem suggests that to a first-order approximation, adaptation can be ignored. As long as environmental changes are small and gradual behavioral changes equal zero, accounting for adaptation is not necessary for measuring damages. Second-order effects, or rapid changes to the environment, are more important situations in which to consider adaptation when estimating damages. The degree of climate change that requires consideration of adaptation and nonlinearity is an empirical issue.

CHAPTER 2
ADAPTATION TO CLIMATE CHANGE: DISENTANGLING REVENUE AND CROP CHOICE RESPONSES†

2.1 Introduction

To what extent can farmers adapt to climate by changing crops and production practices, and thereby mitigate the potentially negative impacts of climate change? A growing body of literature considers how agricultural outcomes respond to random, year-to-year weather variations in fixed locations. However, this approach cannot reasonably account for adjustments farmers would make in response to a permanent shift in weather (Deschênes and Greenstone 2007; Schlenker and Roberts 2009; Deschênes and Greenstone 2011; Dell et al. 2012).

Another body of literature, following seminal work by Mendelsohn et al. (1994), uses a hedonic approach to link farmland climate values, implicitly accounting for adaptation. However, this approach may suffer from omitted variable bias. For example, factors associated with climate, like the availability of irrigation water, can confound the causal interpretation of a cross-sectional association between local climate and outcomes (Deschênes and Greenstone 2007; Schlenker et al. 2005, 2006). A second problem with the hedonic approach is that land values do not impart the mechanism that underlies the association, even if the association is defensibly causal. Thus, most of the literature on potential climate-change impacts either ignores adaptation or uses less persuasive, cross-sectional identification strategies.

An exception is Burke and Emerick (2016), who account for adaptation by examining how corn yields respond to both short and long-run changes in weather. Over the long run, farmers can adjust practices, but they cannot over the short run. Adaptation accounts for this difference between short-run and long-run responses. This approach leverages the fact that climate trends can differ across regions, even those within relatively close proximity to each other and with similar baseline conditions. Even within states, climate trends over a 30-year horizon often vary by 1-2°C, mimicking the amount of change anticipated.

In this study, I use a long history of crop choice and productivity outcomes in the United States to consider how each responds to short-run weather fluctuations as well as longer-run changes in climate. To estimate climate, I use backward-looking, rolling means of weather measures, with the window length selected such that it best predicts weather in the current period. This identification strategy assumes that farmers rationally extrapolate from past weather to form expectations about current weather. Thus, farmers who experience warmer or cooler temperatures over many years may rationally adjust their expectations and hence adapt crop choices and production practices accordingly.

† This chapter is a result of collaboration with Michael J. Roberts.

Furthermore, climate trends differ across counties, with some cooling and others warming, especially concerning critical extremes. Thus, I identify adaptive effects by comparing changes in output and cropping patterns in places experiencing different changes in climate and weather. Unlike earlier models of crop choice that rely on cross-sectional identification strategies (Cohn et al. 2016; Miao et al. 2015; Seo and Mendelsohn 2008; Kurukulasuriya et al. 2008), all models in the present study include county fixed-effects to control for unobserved heterogeneity and isolate arguably random, within-county weather and climate fluctuations. This approach uses a similar identification strategy to Burke and Emerick, while broadening the scope beyond corn to simultaneously consider the five most abundant crops in the United States: corn, soybean, wheat, cotton, and hay. These crops comprise the majority of acres farmed and agricultural production (53% and 57% respectively, USDA 2017) from 1950-2010. These shares are even larger in the regions I focus on, specifically the Midwest and the South.

I first develop an econometric model of county-level revenue-per-acre, conditional on weather and climate (degree days and precipitation) from 1950-2010. Revenue-per-acre is calculated by multiplying the yield and long-run average price across the five crops. This approach resolves an endogeneity problem prevalent in earlier studies — prices and on-farm storage can respond to local weather shocks, attenuating the true weather effect (Fisher et al. 2012). This model uses measures of both weather (annual aggregations) and climate (historic rolling means of weather), as well as county fixed-effects and state-specific trends to account for unobserved heterogeneity and heterogeneous productivity gains as well as other time-varying factors.

To better understand the underlying mechanism of adaptation, I then develop two systems of equations: one that links the shares of each crop to climate measures and one that links the revenue-per-acre of each crop to weather measures. Unlike the aggregate revenue model, these crop-specific regressions cannot neatly account for production complementarities and heterogeneous land quality. However, they do provide clues about the nature and extent of likely adaptation.

From these sets of estimates, I: (1) use the crop-specific revenue-per-acre system to estimate how the revenue-per-acre of each crop responds to weather, holding land allocation and production practices fixed; (2) use the land-allocation system to estimate how crop choice changes with climate; and (3) use the aggregate revenue-per acre model to estimate the total effect of climate change, which may embody adjustments in production practices that are more subtle than crop-switching, such as adjustments in planting times, double-cropping and rotating complementarities of cropping systems. By considering all three models, I estimate overall climate impacts and disentangle some sources of adaptation.

I then make predictions under a series of uniform climate-change scenarios that increase temperatures from historical baselines by 1-5°C. The results suggest slight decreases in cotton, soybean, and hay acres, and small increases in corn and wheat acres. The total effect of climate change (3) closely matches the sum of the predicted impacts of crop-specific regression (1). The effects of adaptation appear modest and mostly harmful relative to predictions without climate change. These results suggest that farmers' short-run adaptation responses cannot mitigate the harmful effects of climate change, but long-run adaptation responses may provide slight improvements.

This paper offers a number of contributions. First, I develop climate measures that best predict realized weather and revenue-per-acre, which is consistent with the idea that adaptive decisions concern those who connect to anticipated as opposed to realized weather outcomes. Second, I develop novel, panel-based estimates of broad agricultural outcomes, inclusive of adaptation, that is identified using within-county temporal changes in climate. These changes are more plausibly exogenous than previous cross-sectional approaches. Third, I resolve endogeneity issues that arise when using revenue-based outcome measures as a result of storage and price responses to transitory weather shocks. Fourth, I disaggregate revenue-per-acre outcomes to identify crop choices and outcomes that are jointly affected by climate. I find that adaptation offers a slight reduction in damages relative to models that extrapolate only from weather responses.

The remainder of this paper is organized as follows. Section 2.2 outlines a theoretical framework for crop-switching similar to the modeling approach used in the present study. Section 2.3 discusses the empirical approach of this study. Section 2.4 describes the data in the present study. Section 2.5 presents the empirical results. Section 2.6 provides climate change impacts. Finally, Section 2.7 offers the conclusion.

2.2 Theoretical Framework

Adaptation is thought to be an important component in dealing with climate change to mitigate damages in the future. One approach is to switch to crops that are less sensitive to heat and away from those crops more sensitive. Each year, a farmer will decide how many acres to plant each crop from the current weather conditions on the farm and expectations about the climate based on previous weather patterns. As a result, the harvest at the end of the season will be impacted positively if the weather/climate was beneficial to crops or negatively if the weather/climate was harmful. Optimal decisions result in optimal harvested yield for the farmer. In this section, I outline a simple crop yield model that accounts for weather and climate. I then extend this approach to account for crop-switching as a decision of proportion of total acres. Finally, I outline how to compare short (weather-effect) and long-run (weather-climate-effect) adjustments that account for adaptation due to climate change, which follows my empirical approach below.

2.2.1 Crop-switching Model

Starting with a simple model of output by farm i in year t as $y_{it} = f(x_{it}, z_{it})$ where z_{it} are temperatures on the farm and x_{it} are decisions made by the farm that maximize profits. Farmers make decisions from expectations about weather based on historical weather. Farm-level decisions are a function of temperatures where the simple model transforms to $y_{it} = f(x_{it}(z_{it}), z_{it})$. I model a piece-wise production function with respect to low L, medium M, and high H temperature distribution bins,

$$y_{it} = \beta_0 + \beta_1 L_{it} + \beta_2 M_{it} + \beta_3 H_{it}$$

In this simple model, output y_{it} is measured as production per acre where q_{it} is production and a_{it} is harvested acres. Equation 1 defines aggregate output for all crops on the farm, or the sum of individual crop output,

$$y_{it} = \sum_{c=1}^{C} y_{cit} = \beta_0 + \beta_1 L_{it} + \beta_2 M_{it} + \beta_3 H_{it} \tag{2.1}$$

Farmers respond to year-to-year changes in temperature by making short-run adjustments to their decisions. However, changes observed over more extended periods would indicate permanent long-run shifts in their decision. I define these permanent shifts as an adaptation to climate change.

In the long-run, d years measure the expectations of farmers by comparing a long-run measure of climate using temperatures over specific intervals. Using a long-run measure of climate accounts for expectations about permanent changes in the weather and impacts on output. We, therefore, define aggregate output to include weather and climate as,

$$y_{itd} = \beta_0 + \beta_1 L_{it} + \beta_2 M_{it} + \beta_3 H_{it} + \beta_4 \overline{L_{id}} + \beta_5 \overline{M_{id}} + \beta_6 \overline{H_{id}} \tag{2.2}$$

Each year, farmers decide how many total acres to cultivate and how much of each crop to plant. A farmer's decision is not random and relies on knowledge from previous seasons. Because the output is a function of temperature, their expectations about temperature will drive their decisions. If, for example, temperatures are expected to increase, then the farmer will rely on their knowledge of each crops' sensitivity to temperature and plant accordingly. These changes in decisions are known as crop-switching.

To model how output changes based on crop-switching practices, I explicitly model the decision variable, x_{it}. In year t, farmer i faces the decision of how much to plant of each crop c, which is composed of five varieties: corn, cotton, hay, wheat, and soybean. Their decision is defined by the proportion of total acres a_{cit} they designate for each crop, which is an element of $x_{cit} \in [0, 1]$. For example, a farmer could decide to plant 20% of total acres with each crop, in which case $x_{cit} = 0.2$. Given this specification, the sum of x_{cit} is equal to one.

If farmers are responding to long-run expectations, then acreage decisions are $x_{cid} = f(\overline{z}_{id})$. As with production, I assume x_{cid} follows a piecewise process given by,

$$x_{cid} = \beta_0 + \beta_1 \overline{L_{id}} + \beta_2 \overline{M_{id}} + \beta_3 \overline{H_{id}}$$

It is important to consider that acres are finite for any farmer i; although, acres can change for each year t or interval d. As a result, changes in the proportion of individual crops acres are correlated across crops (i.e. an increase in corn acres decreases at least one other crop acre on the farm). Therefore, each crop is modeled as one element of a system of equations where the total cultivated acreage a_{id} is defined as,

$$x_{id} = \sum_{c=1}^{5} x_{cid} = \begin{cases} x_{1id} = \beta_{c0} + \beta_{c1}\overline{L_{id}} + \beta_{c2}\overline{M_{id}} + \beta_{c3}\overline{H_{id}} \\ \vdots \\ x_{5id} = \beta_{c0} + \beta_{c1}\overline{L_{id}} + \beta_{c2}\overline{M_{id}} + \beta_{c3}\overline{H_{id}} \end{cases} \tag{2.3}$$

This system defines each farmers' acreage decisions.

The correlation between crops carries over to annual outputs, therefore a similar system of equations is used to model total crop output y_{it}, given below,

$$y_{it} = \sum_{c=1}^{5} y_{cit} = \begin{cases} y_{1it} = \beta_{c0} + \beta_{c1}L_{it} + \beta_{c2}M_{it} + \beta_{c3}H_{it} \\ \vdots \\ y_{5it} = \beta_{c0} + \beta_{c1}L_{it} + \beta_{c2}M_{it} + \beta_{c3}H_{it} \end{cases} \tag{2.4}$$

Using equations 2.2, 2.3, 2.4 above, I can identify the effects of crop-switching with respect to short-run and long-run temperatures. First, I estimate aggregate output per acre in response to weather and climate changes in equation 2.2 which captures adjustments in production practices, such as planting times, double-cropping, and rotations. Next, I use crop-shares from equation 2.3 to estimate individual crop acre changes as the climate changes over the interval, d. Crop-specific output per acre in equation 2.4 estimates crop output responses to weather, holding land allocation and production practices fixed.

I then combine estimates from the crop-share and individual crop revenue-per-acre to measure the impact of crop-switching with a weather-effect (short-run) and a climate-weather-effect (long-run). Next, I lay out the same effects for the aggregate measure of output, which imply different forms of adaptation. The weather-effect assumes farmers are making decisions based on current weather conditions, while the weather-climate-effect assumes farmers are taking into account previous weather patterns (e.g., past 10-years). The weather-climate-effect suggests farmers can recognize changes in the climate and adapt their practices to mitigate damages. I can compare these two effects to determine whether adaptation can mitigate the adverse effects of climate change and by how much.

2.3 Empirical Approach

Previous studies have estimated short-run responses by accounting for the weather to draw inferences about short-run adaptation and make predictions under climate change. However, short-run responses may not translate to long-run responses, so estimates may over- or under-estimate potential damages. Other studies estimate long-run responses using cross-sectional identification and assume adaptation is implicit in revenue based on a discounted sum of all future net-benefits. However, cross-sectional approaches cannot easily defend the idea that climate is random and not associated with other factors correlated with climate. Moreover, cross-sectional studies linking land values or other aggregated outcomes can provide little insight into the likely mechanism that underlies adaptation.

I address these issues in my specifications by including exogenous measures of weather and climate, while controlling for time-invariant factors and state-specific trends. The specification compares short-run responses in weather (weather-effect) and short-run responses conditional on climate (weather-climate-effect), thus identifying short-run versus long-run adaptation.

To estimate the impacts of weather and climate, I rely on a nonlinear transformation of temperature, degree days, and precipitation. Degree days accounts for the rise and fall of temperatures in a day and has been used extensively to measure the effects of temperature on crop yields (Schlenker et al. 2006; Schlenker and Roberts 2009). Degree days are separated into bins to account for low (0-10°C), medium (10-30°C), and high (above 30°C) temperatures. I use variation in the short-run weather (year-to-year) and long-run climate (rolling mean window) for degree days and precipitation (see data section below).

The identification strategy relies on exploiting exogenous variation in weather and climate as a natural experiment. I assume differences in counties that warmed the most versus cooled the most in the short and long-run will see declines in revenue-per-acre and will also adjust their proportion of crop acres. As a result, counties will change their behaviors due to changes in weather and climate through adaptation practices, such as crop-switching. Figure 2.1 provides USDA Economic Research Services (ERS) agricultural regions and their demeaned county-level climate variation degree day 10-30°C and degree day 30°C. Figure 2.2 provides degree day 30°C rolling mean differences between the 1950's and 2000's at the county-level. I exploit this variation within county while also controlling for state-specific trends to estimate revenue-per-acre and proportional changes in individual crop acres.

Crop revenue-per-acre is used as the primary outcome variable in this analysis. The reason for this choice is because it is a simple way to value aggregate crops. By aggregating the value of crops in this way, it allows for productivity changes through cropping complementarities. Examples of productivity changes include crop rotations affecting pest management or improved time management due to different planting and harvesting times (Cai et al. 2013; Livingston et al. 2014).

To construct revenue-per-acre, I fix state-level prices of each crop at the average over all years.

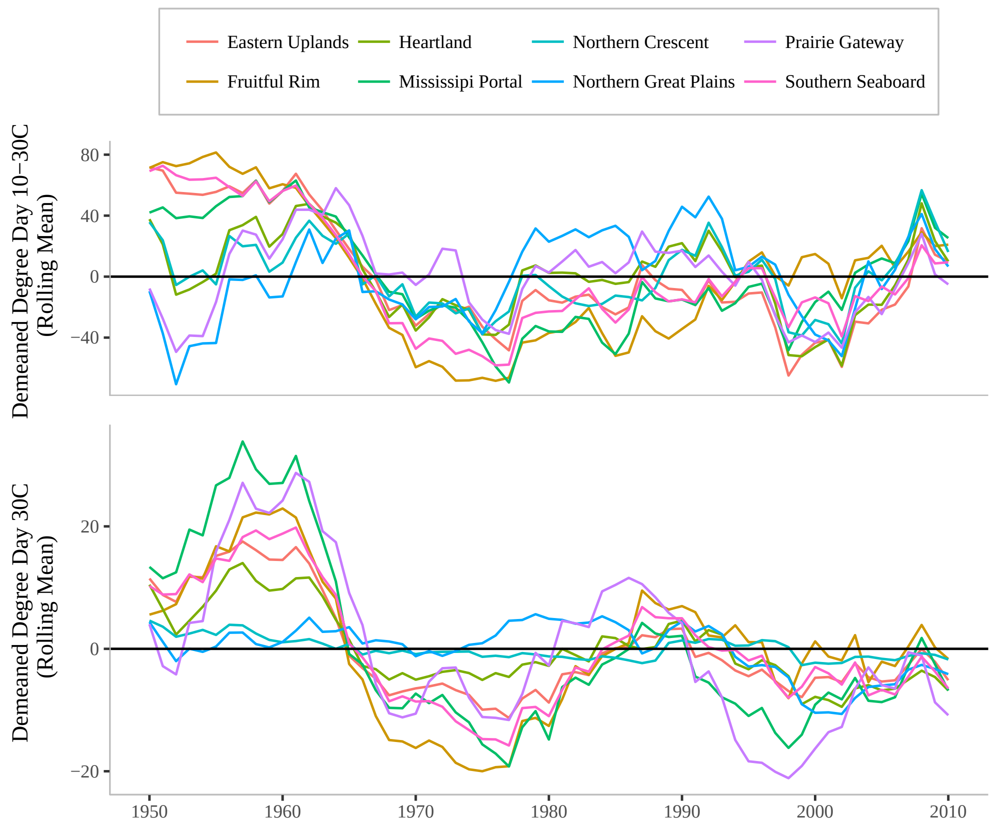

Figure 2.1: **Region specific climate variation from 1950-2010**

Notes: Figure provides county-level demeaned values for Degree Day 10-30°C and Degree Day 30°C with a rolling mean window from 1950-2010 for regions across the US east of the 100th degree meridian.

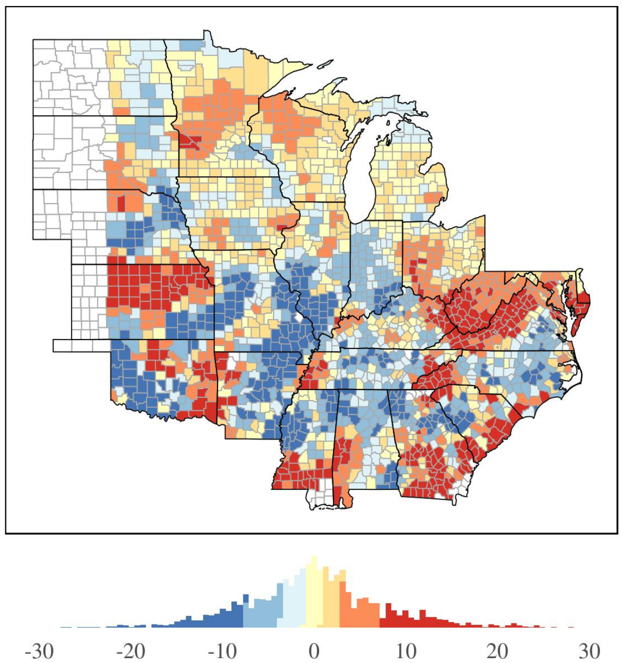

Figure 2.2: **Change 2000's to 1950's Degree Day 30°C (Rolling Mean)**

Notes: Map provides differences in Degree Day 30°C with a rolling mean window while controlling for weather, state specific trends, and county fixed-effects. Density below provides a histogram of colored regions.

Prices are adjusted for inflation using the GDP deflator with the base year of 2010. Fixing prices resolves problems with endogeneity, such as a storage and weather-induced price shocks. Prices are also volatile from year-to-year, so one year may not reflect the appropriate terms of trade between two crops. Averaging prices over the time series resolves the volatility in prices and follows that crop prices tend to move together over the long-run (Sumner 2009; Roberts and Schlenker 2013).

I consider three models to measure short and long-run adaptation, (1) aggregate revenue-per-acre, (2) crop-specific revenue-per-acre; and (3) crop-shares as a proportion of total acres. Aggregate revenue-per-acre includes measures of weather and climate where revenue-per-acre equals the sum of five main crop yields in the US (corn, cotton, hay, soybean, wheat) multiplied by a fixed state-level price. Crop-specific revenue-per-acre separates the five main crops and estimates them separately in a system of equations using only variations in weather. And the crop-share estimates the proportional change in individual crop acres also in a system using just variations in climate. From these three models, I can disentangle the source of adaptation and estimate effects with a warming climate.

2.3.1 Climate Measure

Measuring climate is difficult because the appropriate length of time a permanent shift in behavior has occurred is not well established. Previous research has estimated climate using a monthly weather average or nonlinear transformation of temperature over a period of 30-years (Mendelsohn et al. 1994; Schlenker et al. 2005). However, this isolates climate as an average, or mid-points, in the distribution of temperatures, and may mask changes I seek to identify in climate. Other studies have used rolling mean windows over 3-5 years (Henderson et al. 2017) or differences in decades (Burke and Emerick 2016). However, it is unclear whether these are the best measures of climate for identifying long-run changes.

The first empirical exercise focuses on developing a measure of climate that provides a reasonable expectation of the current weather. As has been shown previously, previous weather may influence current yields or revenue via price responses (Roberts and Schlenker 2013). Therefore, a climate measure that best predicts the current weather can account for expectations over the long-run.

Three transformations of degree days and precipitation are calculated as climate for three intervals: decade average, rolling mean window, and an instrument variable of lag year length. Decade average is simply the midpoint of the weather variable for 10-years. For the rolling mean and instrumental variable approach, window and lag lengths must be chosen that best predict the weather. A cross-validation exercise is discussed below for choosing the window and lag lengths. Once the climate measures are selected, another cross-validation exercise is done to determine which climate variable transformation best predicts revenue-per-acre, thus reducing RMSE from a baseline model. I utilize data on US counties from 1950-2010 and weather from 1900-2010 (see data described below).

The decade average approach averages county-level degree day bins over 10-years. I then estimate a simple linear model,

$$R_{it} = \frac{1}{10}\sum_{t=1}^{10} W_{it} + \gamma_s \tau_t + c_i + \varepsilon_{it}$$

where R_{id} is the revenue variable of interest for county i in state s and year t regressed on the climate variable, W_{it}. I included county-level fixed-effects, c_i and state trends, $\gamma_s \tau_t$.

Next, I determine the best rolling mean window that predicts weather to instrument an exogenous source of variation using climate. I utilize a "right" rolling mean window that measures an average window from the current year until n previous years. This measure allows for decisions to be based on previous weather and removes the ability to predict what will happen next year. To determine the best window, I regress the degree day interval or precipitation, W_{id}, on the degree day measure.

$$W_{it} = \frac{1}{n}\sum_{l=t-n-1}^{t-1} W_{il} + \gamma_s \tau_t + c_i + \varepsilon_{it}$$

where W_{it} is the degree day bin or precipitation regressed on the climate variable, W_{il}, with different rolling mean windows, n. Figure 2.3 provide results of this analysis. I extract the best window that predicts the weather for each degree day interval (Degree Day (0-10°C): 30-year window; Degree Day (10-30°C): 10-year window; Degree Day (30°C): 7-year window; Precipitation: 9-year window) and estimates a rolling mean using each window.

A linear regression is estimated similar to the decade average approach to estimate revenue,

$$R_{it} = \frac{1}{n}\sum_{l=t-n-1}^{t-1} W_{il} + \gamma_s \tau_t + c_i + \varepsilon_{it}$$

The final climate measure uses an autoregression to determine the best lag length as an instrument for predicting weather, similar to the rolling mean approach. However, instead of averaging over an interval I capture the previous k lags from the current year t. The regression setup is,

$$W_{it} = W_{it-1} + ... + W_{it-k} + \gamma_s \tau_t + c_i + \varepsilon_{it}$$

where W_{id} is the degree day bin or precipitation regressed on the climate variable W_{it-k} as different lag lengths. The results suggest the lag lengths that best predict weather are: Degree Day (0-10°C): 30-year lag; Degree Day (10-30°C): 10-year lag; Degree Day (30°C): 7-year lag; Precipitation: 9-year lag). Appropriate lag lengths are calculated for each degree day and precipitation, then used to predict revenue as the previous approaches,

$$R_{it} = W_{it-1} + ... + W_{it-k} + \gamma_s \tau_t + c_i + \varepsilon_{it}$$

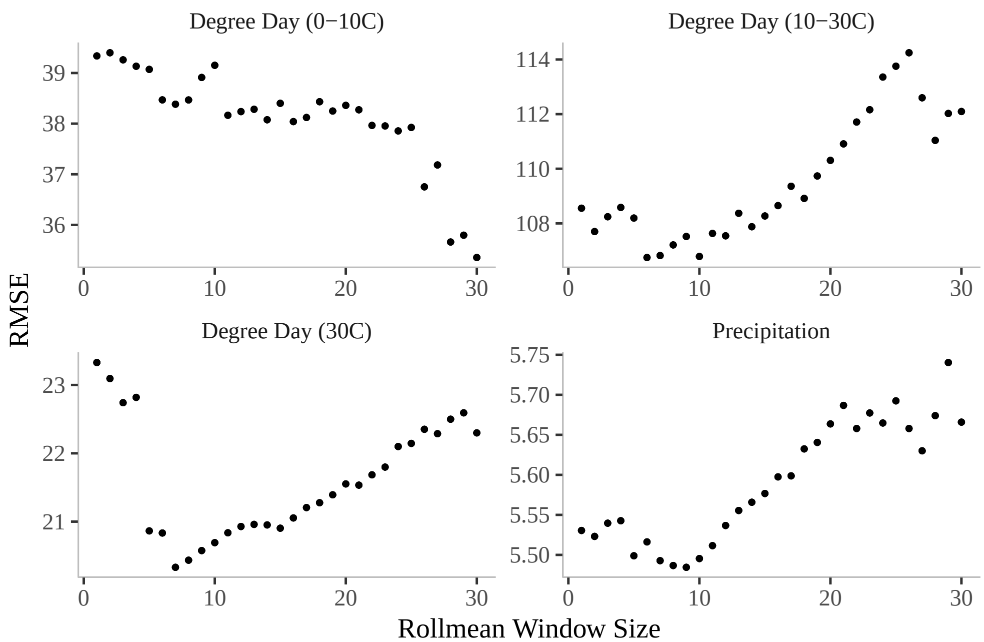

Figure 2.3: **Predicting Climate Intervals**

Notes: Figure provides results from our empirical exercise to determine the best predictor of climate on weather variables as an exogenous source of variation. Rolling mean window size represents the number of previous years used in the window size. We utilize a lag "right" rolling mean window to assume current decisions are made on the previous interval in years.

A cross-validation exercise is used to predict the current period revenue based on each climate measure, move forward in time, and repeat. For example, to predict revenue in 1930 using climate, I estimate each regression model using data from 1900-1929, predict revenue in 1930, and roll forward one year until 2012. I cross-validate in this way because I am seeking to predict the current revenue based on the past climate, so the order matters. With each iteration, I calculate the reduction in RMSE from a baseline model that includes county-fixed-effects and state trends. The baseline model can forecast revenue in a county for each year.

Figure 2.4 provides the results of this analysis. The top panel provides the reduction in RMSE from only the weather, while the bottom panel uses climate variation for each of the climate measures. Along the x-axis are specific revenue measures for the five main crops and an aggregate revenue that includes all crops. Utilizing only weather provides similar results to Schlenker and Roberts (2009), except I use the revenue for corn, cotton, and soybean. For each of the climate measures, the rolling mean measure of climate reduces the RMSE the most. Corn revenue performs the best (14%), which is similar to only using weather. For aggregate revenue, the results suggest climate reduces RMSE by around 4%. Wheat and hay degree day intervals are not well defined in the literature, which is driving such a small reduction; however, the evidence suggests it does better than decade average and an instrument variable using an autoregression. For the remainder of the paper, I estimate climate using the rolling mean measure for degree day bins and precipitation.

2.3.2 Aggregate Revenue

The next empirical exercise relies on estimating adaptation by accounting for an aggregate measure of revenue-per-acre using weather and climate variation within counties. Aggregate revenue-per-acre accounts for adjustments in production practices, such as planting times, double-cropping, and complements of rotating cropping systems. From the weather and climate measures, I can extract these forms of adaptation implicitly through revenue-per-acre.

To estimate aggregate revenue-per-acre, R_{itd}, degree days, $LDD_{it}, GDD_{it}, HDD_{it}$, and precipitation, P_{it}, P_{it}^2, for weather conditional on climate are included in the data generating process as,

$$\begin{aligned} R_{itd} \quad = \quad & \beta_1 LDD_{it} + \beta_2 GDD_{it} + \beta_3 HDD_{it} + \beta_4 P_{it} + \beta_5 P_{it}^2 + \\ & \beta_6 \overline{LDD_{id}} + \beta_7 \overline{GDD_{id}} + \beta_8 \overline{HDD_{id}} + \beta_9 \overline{P_{id}} + \beta_{10} \overline{P_{id}}^2 + c_i + \gamma_s \tau_t + \varepsilon_{itd} \end{aligned} \tag{2.5}$$

Equation 2.5 includes weather variation from year-to-year, t, and climate effects as a right rolling mean, d. County fixed-effects, c_i, are included to control for time-invariant factors to remove concerns of omitted variable bias. State-specific trends $\gamma_s \tau_t$ are used to control for technological differences within states. In this specification, I assume weather and climate are plausibly exogenous;

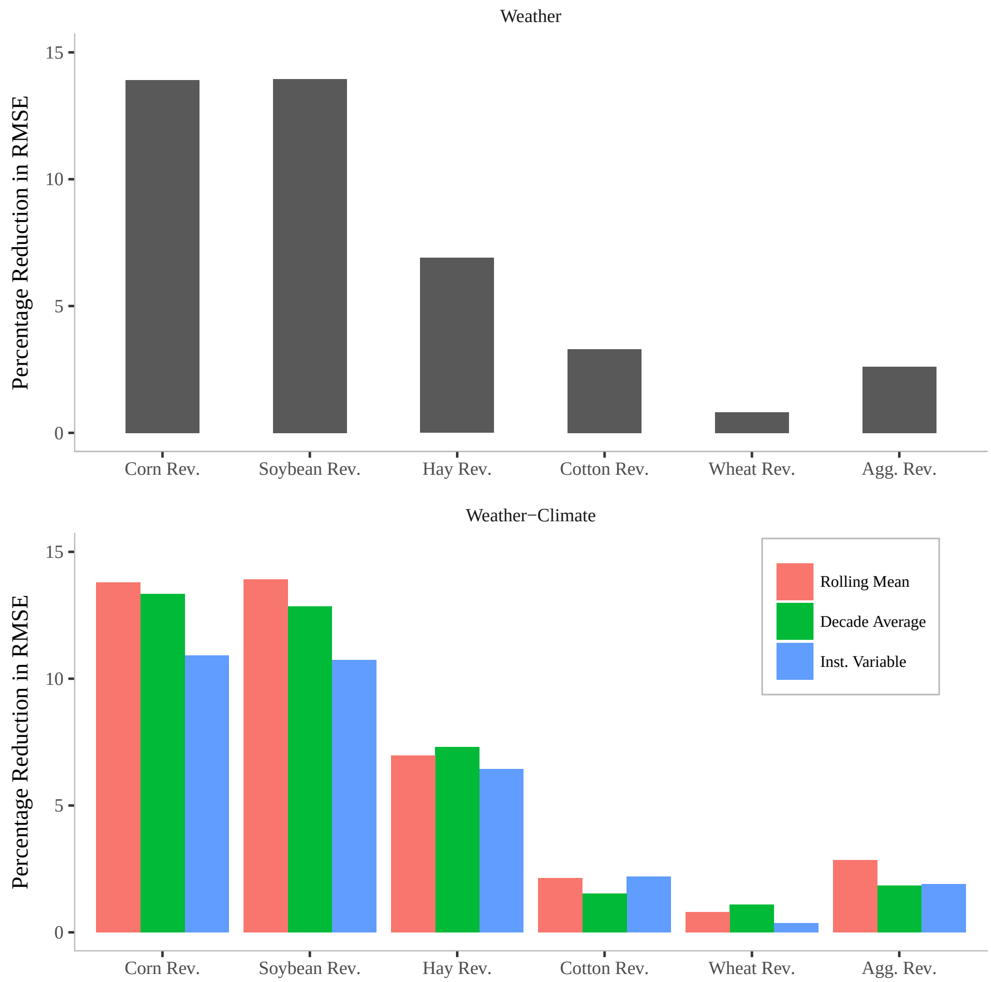

Figure 2.4: **Out-of-sample Climate Predictions**

Notes: Bar charts display percentage reduction in root-mean-square-error (RMSE) for out of sample predictions for various revenue measures (corn, soybean, hay, cotton, wheat, and aggregate revenue) conditional on weather and weather-climate. The top panel displays predictions using only weather. The bottom panel displays predictions using weather and various climate measure: Rolling Mean, Decade Average, and Instrument Variable. Each model is estimated by making predictions on next year's revenue from the previous year's climate measure. Performance is based on percentage reduction from a baseline model that includes state trends and county fixed-effects.

thus, I can identify a causal relationship to revenue-per-acre.

From equation 2.5 I look for long-run adaptation by comparing coefficients in weather and climate. To measure long-run adaptation to extreme heat, degree days above 30°C for climate should be less harmful (higher) than the coefficient on weather degree days. For degree days between 10-30°C, climate coefficient should be higher than the weather coefficient because farmers can take advantage of the increase in preferred growing degree days. I further seek to use these estimates to predict increases in temperatures from 0-5°C to better understand potential long-run adaptation to climate change.

2.3.3 Disaggregate Revenue

Aggregate revenue-per-acre accounts for changes in production practices implicitly. A widespread discussion among economists is that as the climate changes farmers will decide to plant crops that are less sensitive to extreme heat or plant a variety of crops to offset the impacts on yields (Adams et al. 1990; Kurukulasuriya et al. 2008). Disentangling this change from revenue-per-acre tests this hypothesis against an aggregate measure.

I first disaggregate revenue-per-acre by estimating individual crop revenue-per-acre, holding land allocation and production practices fixed. Next, I separate acres as individual crop-shares to evaluate crop-switching practices within counties. Individual crop revenue-per-acre for each crop is regressed on weather variables to account for short-run adjustments by county. If farmers are making decisions about the variety of crops they have available, the results should be similar to the aggregate weather coefficients. If, however, farmers favor one crop over another, then the weather coefficients will represent those preferences.

To measure the climate effect from crop-switching, crop-shares are regressed on climate variables to measure the changes in the proportion of crop acreage by county. Farmers will decide to plant crops based on how they observe the current season, so short-run changes may not translate to long-run changes. There may be slight variations from year-to-year but will likely not impact revenue in the long-run because of subtle differences in the short-run. Therefore, I focus on climate variation that captures changes in acreage over the long-run due to permanent changes in behavior. Individual crop revenue-per-acre and crop-shares are then combined to determine how changes in crop shares impact revenue-per-acre.

It is important to note that when disentangling revenue-per-acre, problems with correlations between dependent variables exists; that is, decisions to plant one crop will offset the production of another crop due to limitations in total available acres. For example, the decision to plant more corn will result in a decrease in soybean because corn and soybean are planting complementarities. This decision will also impact individual crop revenue-per-acre. For this reason, I model individual crop revenue-per-acre and crop-shares as a system of equations that can deal with correlations between dependent variables. I model each system as a seemingly unrelated regression (SUR).

The data generating process for revenue-per-acre is defined as,

$$R_{cit} = \beta_{c1} LDD_{it} + \beta_{c2} GDD_{it} + \beta_{c3} HDD_{it} + \beta_{c4} P_{it} + \beta_{c5} P_{it}^2 + c_i + \gamma_s \tau_t + \varepsilon_{cit} \tag{2.6}$$

$$R_{cit} = \begin{cases} R_{\text{corn},it} = \beta_{c1} LDD_{it} + \beta_{c2} GDD_{it} + \beta_{c3} HDD_{it} + \beta_{c4} P_{it} + \beta_{c5} P_{it}^2 + c_i + \gamma_s \tau_t + \varepsilon_{cit} \\ R_{\text{cotton},it} = \qquad \vdots \\ R_{\text{hay},it} = \qquad \vdots \\ R_{\text{soybean},it} = \qquad \vdots \\ R_{\text{wheat},it} = \beta_{c1} LDD_{it} + \beta_{c2} GDD_{it} + \beta_{c3} HDD_{it} + \beta_{c4} P_{it} + \beta_{c5} P_{it}^2 + c_i + \gamma_s \tau_t + \varepsilon_{cit} \end{cases}$$

where individual log crop revenue-per-acre is R_{cit} regressed on weather variables $LDD_{it}, GDD_{it}, HDD_{it}$ and precipitation, P_{it}, P_{it}^2. County fixed-effects, c_i, are included to control for time invariant factors. State specific trends, $\gamma_s \tau_t$ are also included to control for state differences in productivity. The system adjusts the coefficients to account for the correlation in errors.

For crop-shares, the system of equations is defined as,

$$P(A_{cid}) = \beta_{c1} \overline{LDD_{id}} + \beta_{c2} \overline{GDD_{id}} + \beta_{c3} \overline{HDD_{id}} + \beta_{c4} \overline{P_{id}} + \beta_{c5} \overline{P_{id}}^2 + c_i + \gamma_s \tau_t + \varepsilon_{cid} \tag{2.7}$$

$$P(A_{cid}) = \begin{cases} P(A_{\text{corn},id}) = \beta_{c1} \overline{LDD_{id}} + \beta_{c2} \overline{GDD_{id}} + \beta_{c3} \overline{HDD_{id}} + \beta_{c4} \overline{P_{id}} + \beta_{c5} \overline{P_{id}}^2 + c_i + \gamma_s \tau_t + \varepsilon_{cid} \\ P(A_{\text{cotton},id}) = \qquad \vdots \\ P(A_{\text{hay},id}) = \qquad \vdots \\ P(A_{\text{soybean},id}) = \qquad \vdots \\ P(A_{\text{wheat},id}) = \beta_{c1} \overline{LDD_{id}} + \beta_{c2} \overline{GDD_{id}} + \beta_{c3} \overline{HDD_{id}} + \beta_{c4} \overline{P_{id}} + \beta_{c5} \overline{P_{id}}^2 + c_i + \gamma_s \tau_t + \varepsilon_{cid} \end{cases}$$

where individual crop-shares are defined as a proportion of total acres $P(A_{cid})$ regressed on climate variables $\overline{LDD_{it}}$, $\overline{GDD_{it}}$, $\overline{HDD_{it}}$ and precipitation, $\overline{P_{id}}, \overline{P_{id}}^2$. County fixed-effects, c_i, are also included in this specification along with state specific trends, $\gamma_s \tau_t$. The system further adjusts coefficients in the same manner.

In the crop-share model, the dependent variable measures the proportion of crop acres as a share of total acreage in a county. As a result, the sum of all five crops equals one. The goal of this system is to estimate the proportional changes while taking into account the crop-switching effect (changes in one crop acres results in changes in other crops without exceeding a sum of one). A problem arises here because the dependent variable is not linear (proportions are from zero to one). Since the end goal is to make predictions on climate change, I need to ensure the model estimates an accurate measure of proportional changes when making predictions so that the dependent variables sum to one across all crops.

Therefore, I either need to ensure the coefficients sum to zero (sum of dependent variables equal one) by constricting the coefficients or adjust the dependent variable into a linear transformation and transform back to crop-shares. Constricting the coefficients may introduce problems I prefer to not deal with, such as inefficient estimates; therefore, I prefer to transform the dependent variable into z-scores, estimate the system, and transform back into crop shares. To accomplish this, I transform each crop-share by adding 0.001 and then divide by 1.00101 to restrict the bounds to be greater than zero and less than one. A quantile distribution (inverse CDF) is produced for each crop share and regressed on climate variables in the system while controlling for county-level differences and state-trends. I then transform back using a CDF, so predictions are in crop shares.

From these two SUR models, I estimate the short-run weather responses in revenue-per-acre and the long-run climate responses in crop-share changes. These two estimates are used to study how changes to the weather and climate affect changes in individual crop revenue-per-acre. These two estimates provide a weather-climate-effect I use to compare against the short-run effects for individual crop and an aggregate measure of revenue-per-acre.

2.4 Data

Agriculture data comes from the National Agricultural Statistics Service (NASS) released by the United States Department of Agriculture (USDA). The data reports county-level estimates for production and acreage for 1,923 counties from 1950-2010. The primary crops are corn, cotton, hay, wheat, and soybean which make up a majority of the crops in the U.S. (57% of total U.S. agricultural production). From here, I keep counties that have at least one observation for crop acres in 1950. To deal with irrigation, counties west of the 100th-degree meridian are removed. Figure 2.5 shows a map of counties and crop acres in 1950 included in the sample.

To calculate crop revenue-per-acre, crop yields and average state-level prices are multiplied together – adjusted using the GDP deflator with the base year of 2010. Prices are unavailable at the county-level, so I utilize state-level prices. Crop shares are estimated using county-level crop acreage divided by total crop acreage in that county. If a county does not report crop acres, a zero is reported.

The mechanism between crop yields and weather is well defined in the literature using an agronomic variable known as degree days (Schlenker and Roberts 2009). Degree days are defined as the amount of time during the day the environment is exposed to a threshold of temperature. For example, suppose for half of the day it is 30°C. A simple calculation of degree days above 25°C would involve 5 degrees for half of a day; therefore, for that particular day, the degree days above

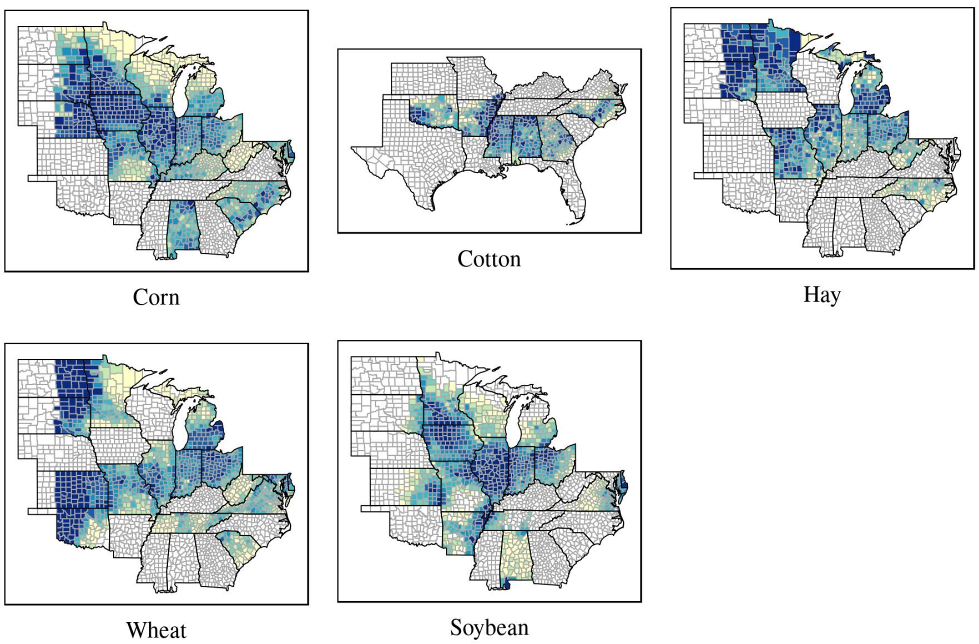

Figure 2.5: **Sample of Counties and Crop Acreage**

Notes: Map of county-level crop acreage used in this paper. Darker colors in blue report more acres relative to % of other county acreages. Light colors in yellow are fewer acres relative to other acres.

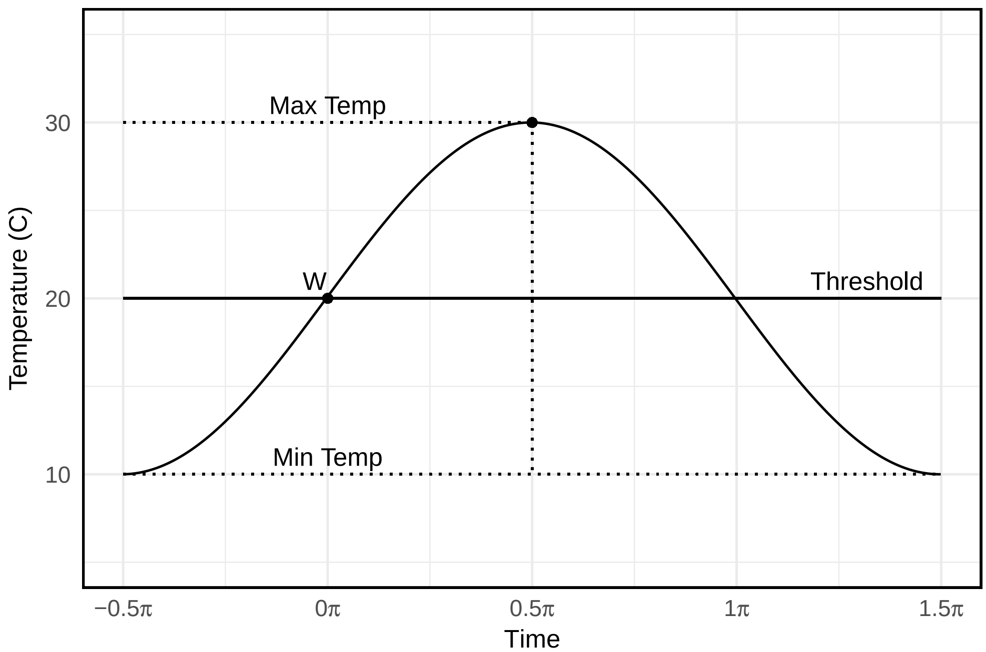

Figure 2.6: **Integrated Sine Approach for Calculating Degree Days**

Notes: Figure provides an example of how degree days are calculated using the integrated sine approach. In each day, a minimum temperature of 10°C and max temperature of 30°C is fit to a sine curve where W is the average temperature and a threshold is defined. Time begins at the radian of the sine curve from -0.5π to 1.5π to account for lower temperature (min) increasing until the hottest temperature (max) and falling to the lowest temperature. In this example, to find degree days above the average temperature (threshold), the sine curve is inverted and integrated between the temperature threshold (20°C - 30°C).

25°C would be 2.5. Exposure to heat during the day becomes longer as it becomes warmer and shorter if it becomes cooler. However, this simple example is not easily calculated using the described method because the exact time for each degree throughout a full day is unknown – weather data provides a minimum and maximum temperature but masks the nonlinear temperature changes occurring during the day (i.e., rise and fall of temperatures). I use an integrated sine approach to capture the nonlinear temperature changes through the day (Arnold et al. 1960; Baskerville and Emin 1969). The method accounts for nonlinear changes by integrating over the minimum and maximum temperature of a day and within a defined threshold (see Figure 2.6).

Temperature data before 1950 is needed to estimate the rolling mean degree day measures. However, temperature is not available at the county-level pre-1950 (see Schlenker and Roberts 2009 for post-1950 data). To overcome the lack of historical data, daily temperatures are developed using a relative anomaly interpolation technique from 1895-2010 using monthly observations from PRISM (2018). I focus on the growing season from March to October and aggregate to county-level for each year[1].

To maintain consistency across crops, I identify three thresholds: low degree days (LDD), preferred growing degree days (GDD), and high heat degree days (HDD). Degree days between 0-10°C are identified as temperature effects below the growing degree day threshold. Equation 2.8 outlines the measure for degree days by estimating the density of time (T) at each degree in each day j in year t and county i. These measures are summed across all days from March to October.

$$LDD_{it} = \sum_{j=\text{March 1st}}^{\text{October 31st}} \int_{T=0}^{\infty} \min(T-0, 10) h_{itj}(T) dT \tag{2.8}$$

Equation 2.9 estimates the growing degree days from 10-30°C. Typically, this measure is defined as growing degree days because it represents an optimal level of temperature for crop growth. Warmer temperatures within this threshold represent an increase in the marginal revenue, while a decrease in these temperature thresholds represents a slight decline in marginal revenue.

$$GDD_{it} = \sum_{j=\text{March 1st}}^{\text{October 31st}} \int_{T=10}^{\infty} \min(T-10, 30) h_{itj}(T) dT \tag{2.9}$$

And degree days above 30°C are estimated in equation 2.10. In general, this measure attributes to harmful temperature effects since an increase in HDD represent an increase in extreme temperatures above 30°C. However, crops less sensitive to warmer temperatures may see less of an effect, or even positive effects depending on the region.

$$HDD_{it} = \sum_{j=\text{March 1st}}^{\text{October 31st}} \int_{T=30}^{\infty} (T-30) h_{itj}(T) dT \tag{2.10}$$

Each crop will have a different response to each threshold. Corn, for example, will see slight negative revenue effects from 0-10°C, positive effects from 10-30°C, and negative effects for degree days above 30°C. Soybean and cotton will have a similar impact. Hay and wheat responses are not well established in the literature, but I seek to estimate the effects here.

2.5 Empirical Results

The analysis focuses on the effects of short and long-run changes in weather and climate from an aggregate and disaggregate measure of revenue-per-acres. I focus on county-level activity from 1950 - 2010 for the five main crops corn, cotton, hay, soybean, and wheat to study the effects of implicit adaptation in the aggregate measure. I also isolate a popular adaptation technique that allows for

[1]Documentation for developing fine-scale weather data and calculating degree days is available at https://github.com/johnwoodill/Fine-Scale-Weather-Interpolation

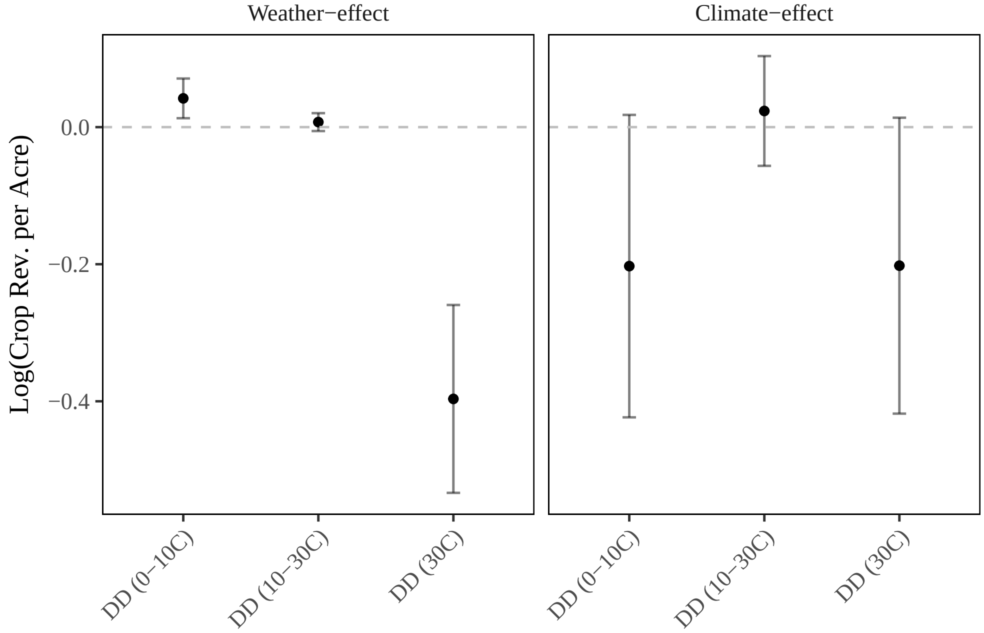

Figure 2.7: **Aggregate Revenue per Acre Temperature Coefficients**
Notes: Figure reports aggregate revenue-per-acre for weather and climate effects for degree days 0-10°C, 10-30°C, and greater than 30°C. Weather-effect captures short-run year-to-year changes and climate-effect capture long-run rolling mean climate intervals. Standard errors are clustered by state and bands represent a 95% confidence interval.

crop-switching between the five main crops of disaggregating revenue-per-acre. These two measures are then compared to identify short and long-run changes as sources of adaptation.

2.5.1 Aggregate Revenue Results

Table 2.1 provides results from the aggregate revenue-per-acre specification. A plot of coefficients with 95% confidence intervals is also provided in Figure 2.7. The results reported in Table 2.1 have been multiplied by 100, so the effects can be interpreted as follows: for every 100-degree day increase, the percentage change in aggregate revenue-per-acre changes by the reported coefficient. Column 5 estimates in Table 2.1 are as expected: the weather effect is positive and significant for degree days 10-30°C (0.01) and negative for extreme heat degree days (-0.40). Precipitation also reports a positive and significant effect (3.20) and negative effects for excess precipitation (-0.03). For climate effects, the extreme heat effects are less than the weather effects (-0.19), which suggests some adaptation in the long-run. Degree days between 0-10°C are similar to the weather effects but insignificant.

The implications of these results suggest there is some implicit level of adaptation in aggregate revenue-per-acre; however, I do not know if a single adaptation response or a combination provides

Table 2.1: Regression Model explaining Crop Revenue per Acre

	Log(Crop Revenue per Acre)				
	(1)	(2)	(3)	(4)	(5)
Weather-effect					
Degree Days (0-10°C)	**0.16**	**0.16**	**0.04**	**0.01**	**0.04**
	(0.00)	(0.00)	(0.00)	(0.00)	(0.01)
Degree Days (10-30°C)	**0.00**	**0.01**	**0.01**	**0.02**	0.01
	(0.00)	(0.00)	(0.00)	(0.00)	(0.01)
Degree Days (30°C)	**−0.41**	**−0.56**	**−0.41**	**−0.39**	**−0.40**
	(0.01)	(0.01)	(0.01)	(0.00)	(0.07)
Precipitation	**3.85**	**3.59**	**3.54**	**3.80**	**3.20**
	(0.13)	(0.11)	(0.09)	(0.08)	(0.52)
Precipitation Squared	**−0.07**	**−0.07**	**−0.07**	**−0.07**	**−0.07**
	(0.00)	(0.00)	(0.00)	(0.00)	(0.01)
Climate-effect					
Degree Days (0-10°C)	**0.04**	**0.48**	**−0.12**	**0.03**	**−0.20**
	(0.01)	(0.01)	(0.01)	(0.01)	(0.11)
Degree Days (10-30°C)	**−0.02**	**−0.06**	**0.04**	**0.04**	0.02
	(0.00)	(0.00)	(0.00)	(0.00)	(0.04)
Degree Days (30°C)	**−0.45**	**−1.02**	**−0.26**	**−0.09**	−0.19
	(0.01)	(0.01)	(0.01)	(0.01)	(0.11)
Precipitation	**15.39**	**9.57**	**0.96**	**6.90**	−0.54
	(0.26)	(0.48)	(0.37)	(0.34)	(3.73)
Precipitation Squared	**−0.34**	**−0.20**	**−0.04**	**−0.14**	−0.01
	(0.01)	(0.01)	(0.01)	(0.01)	(0.07)
Constant	**24.22**				
	(6.23)				
Fixed-effect	–	County	County	County	County
National Trend	–	–	Yes	–	–
National Quad. Trend	–	–	–	Yes	–
State Trend	–	–	–	–	Yes
Clustered SE	–	–	–	State	State
Observations	117,303	117,303	117,303	117,303	117,303
R^2	0.27	0.48	0.69	0.74	0.73
Adjusted R^2	0.27	0.47	0.68	0.74	0.73

Notes: Table reports regression coefficients for log crop revenue-per-acre using weather (year-to-year) and climate (rolling mean) degree day and precipitation variables from 1950-2010. Crop revenue-per-acre is calculated by summing production (lbs) per acre times average crop price for corn, cotton, hay, soybean, and wheat. Climate variables use a 'right' rolling mean window based on best predictions of weather for each variable. Regression estimates are weighted by total county-level total acres (smoothed using a 3-year rolling mean). Estimates in **bold** are statistically significant at 95%. Coefficients have been multiplied by 100.

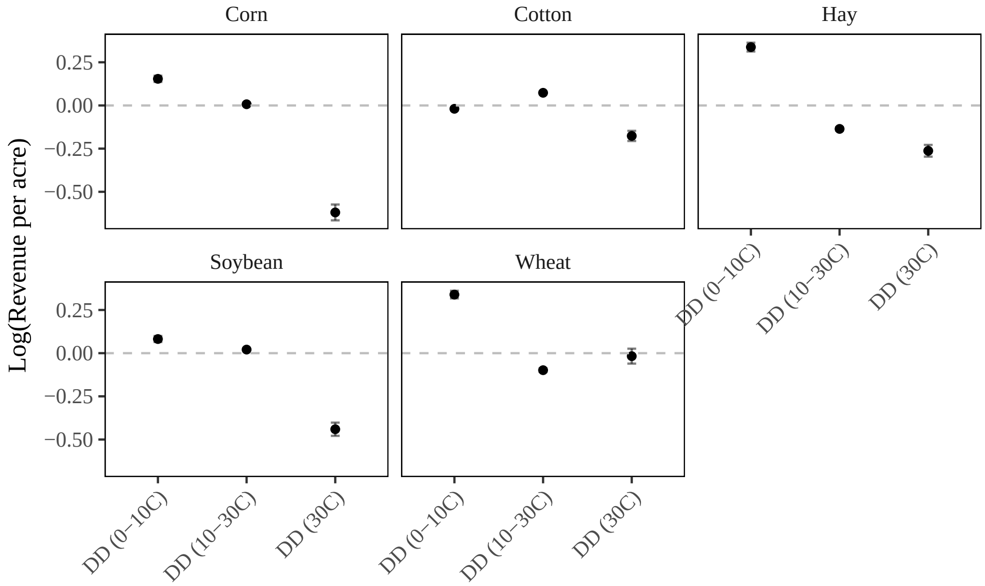

Figure 2.8: **Individual Crop-revenue Regression Estimates**

Notes: Figure reports individual crop revenue-per-acre for weather variables degree days 0-10°C, 10-30°C, and greater than 30°C. Weather variables captures short-run year-to-year changes. Standard errors are bootstrapped by strata state and bands represent a 95% confidence interval.

improvements to revenue-per-acre. In the long-run, farmers are adapting their farming practices which allows them to mitigate some of the effects of climate change. Relying on statistically significant estimates suggests farmers are not able to adjust to reduce the effects climate change entirely.

2.5.2 Disaggregate Revenue Results

Estimating individual crop revenue-per-acre is the first step in disaggregating revenue-per-acre. By separating each crop from the aggregate, a SUR model is used to capture how farmers' decisions respond to changes in weather while controlling for correlation between crop revenues. I merge individual crop revenue-per-acre with the crop-share to estimate changes due to crop-switching.

Table 2.2 provides results for individual crop revenue-per-acre for corn, cotton, hay, soybean, and wheat. Figure 2.8 also provides a plot of coefficients with 95% confidence intervals. Effects of extreme heat are most pronounced for corn (-0.62), followed by soybean (-0.44), hay (-0.26), cotton (-0.18), and then wheat (-0.02). Additionally, the results for degree days 10-30°C are positive across all crops and stable with ranges from 0.01-0.03. For every 100 degree days, I expect individual revenue-per-acre to increase/decrease by the coefficient reported as a percentage change.

From the crop specific results, I extract two implications. First, as previous research has reported, in the short-run, I expect negative effects to crop yield/revenue from extreme heat. I also

Table 2.2: Seemingly Unrelated Regression (SUR) Model explaining Crop Revenue/Acre

	Log(Revenue/Acre)						
	Corn	Corn	Corn	Cotton	Hay	Soybean	Wheat
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Weather-effect							
Degree Days (0-10°C)	**−0.04**	**0.15**	**0.15**	**−0.02**	**0.34**	**0.08**	**0.34**
	(0.01)	(0.01)	(0.01)	(0.00)	(0.01)	(0.01)	(0.01)
Degree Days (10-30°C)	**0.03**	**−0.01**	**0.01**	**0.07**	**−0.14**	**0.02**	**−0.10**
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.01)
Degree Days (30°C)	**−1.28**	**−0.99**	**−0.62**	**−0.18**	**−0.26**	**−0.44**	**−0.02**
	(0.02)	(0.02)	(0.02)	(0.01)	(0.02)	(0.02)	(0.01)
Precipitation	**2.47**	**3.32**	**3.97**	**1.78**	**2.71**	**3.81**	**4.27**
	(0.30)	(0.25)	(0.35)	(0.23)	(0.37)	(0.34)	(0.02)
Precipitation Squared	**−0.06**	**−0.06**	**−0.07**	**−0.02**	**−0.09**	**−0.06**	−0.12
	(0.00)	(0.00)	(0.01)	(0.00)	(0.01)	(0.01)	(0.41)
Fixed-effect	–	County	County	County	County	County	County
State Trend	–	–	Yes	Yes	Yes	Yes	Yes
Bootstrap SE	–	–	Yes	Yes	Yes	Yes	Yes
Observations	117,303	117,303	117,303	117,303	117,303	117,303	117,303

Notes: Table reports regression coefficients for log crop revenue-per-acre using weather (year-to-year) degree days and precipitation variables from 1950-2010. Crop revenue-per-acre equals crop yield per acre times average state-level crop price for corn, cotton, hay, soybean, and wheat. Standard errors are bootstrapped by strata state. Estimates in **bold** are statistically significant at 95%. Coefficients have been multiplied by 100.

provide new results for hay and wheat that show similar effects. A second implication is that these results provide insight into crop-switching practices: with less harmful effects to revenue, I expect farmers to switch to those crops when temperatures increase to offset the negative effects from extreme heat.

Results from the crop-share model are provided in Table 2.3. Figures 2.9 provide plots of degree day coefficients with 95% confidence intervals. The coefficients report climate effects regarding z-scores or the number of standard deviations from the mean. A coefficient of one means acreage increases by one standard deviation for each degree day interval. Interpreting degree days is difficult because of the nonlinear transformation of temperature. Further, by transforming the dependent variable into z-scores for the crop-share model to deal with the correlation of multiple dependent variables also makes interpreting these results difficult. Therefore, I rely on transforming the results of the model to crop-share acres and make predictions with uniform temperature increases to better understand the results.

Table 2.3: Seemingly Unrelated Regression (SUR) Model explaining Crop Share

	Z-score						
	Corn	Corn	Corn	Cotton	Hay	Soybean	Wheat
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Climate-effect							
Degree Days (0-10°C)	**−0.08**	**0.86**	**−0.04**	**0.05**	**0.10**	**−0.06**	**−0.01**
	(0.01)	(0.02)	(0.00)	(0.00)	(0.00)	(0.00)	(0.01)
Degree Days (10-30°C)	**0.03**	**−0.17**	**0.05**	**−0.03**	**−0.07**	**0.06**	**−0.03**
	(0.00)	(0.01)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Degree Days (30°C)	**−1.16**	**−1.51**	**−0.17**	**−0.31**	**0.07**	**−0.47**	**0.03**
	(0.01)	(0.04)	(0.02)	(0.01)	(0.02)	(0.02)	(0.00)
Precipitation	**2.52**	**−7.78**	**−3.77**	**−1.83**	**5.62**	**3.11**	**0.16**
	(0.24)	(1.04)	(0.34)	(0.24)	(0.33)	(0.29)	(0.02)
Precipitation Squared	**−0.07**	**0.05**	**0.02**	0.00	**−0.05**	**−0.01**	**1.11**
	(0.00)	(0.02)	(0.01)	(0.01)	(0.01)	(0.01)	(0.26)
Fixed-effect	–	County	County	County	County	County	County
State Trend	–	–	Yes	Yes	Yes	Yes	Yes
Bootstrap SE	–	–	Yes	Yes	Yes	Yes	Yes
Observations	117,303	117,303	117,303	117,303	117,303	117,303	117,303

Notes: Table reports regression coefficients for a seemingly unrelated regression (SUR) using transformed crop shares to z-scores for climate (rolling mean) degree days and precipitation variables from 1950-2010. Z-scores are calculated using individual crop shares as a proportion of total acres. Climate effects use a right rolling mean window. Standard errors are bootstrapped by strata state. Estimates in **bold** are statistically significant at 95%. Coefficients have been multiplied by 100.

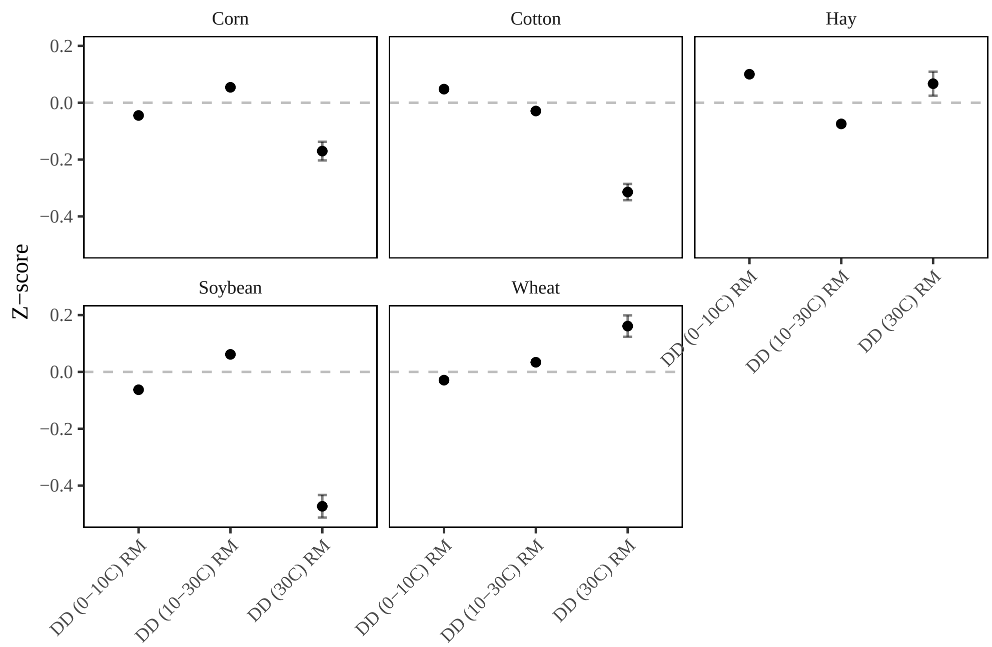

Figure 2.9: **Share Regression Coefficients**

Notes: Figure reports crop shares as z-scores for rolling mean climate variables degree days 0-10°C, 10-30°C, and greater than 30°C. Robust standard errors are provided and bands represent a 95% confidence interval.

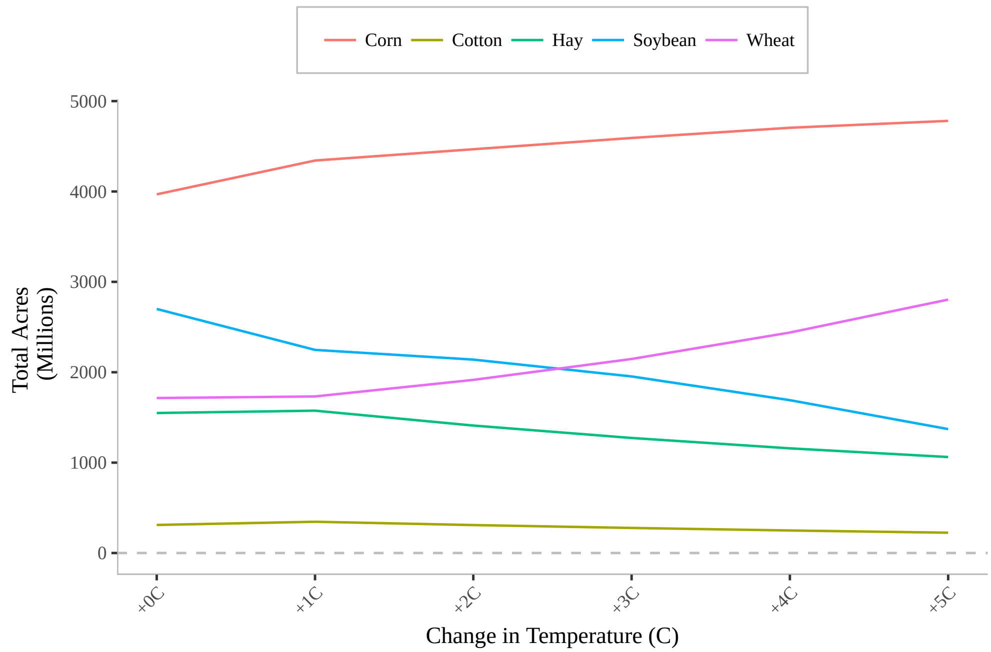

Figure 2.10: **Predicting Acreage Changes**

Notes: Figure reports predicted change in estimates of crop acres from the SUR share model. Crop acres at the county-level vary based on the variation in degree days and precipitation for climate intervals. Temperature increases along the x-axis from 0 - +5°C.

2.6 Climate Change Impacts

The empirical analysis so far has relied on estimating the effects of weather and climate from 1950-2010 for revenue-per-acre and crop-shares. I end the empirical analysis by estimating these effects with a uniform increase in temperatures from 0-5°C. This allows the estimates to be interpreted as the climate changes and also to trace out climate impacts for weather-effects and weather-climate-effects.

To accomplish this, I first increase minimum and maximum temperatures for each county in each year by 1-5°C in the sample. Degree day bins are then re-calculated for each county. The sample for predictions remains the same as the initial analysis, but the degree days are now adjusted to account for the uniform increases in temperature.

Results of crop-share predictions are provided in Figure 2.10. The predictions show slight increases in corn and wheat, and slight decreases in cotton, hay, and soybean. The increases in corn is possibly a result of longer growing seasons and farmers able to plant earlier. Thus, a possible trade-off exists between soybean and corn as complements, which sees a decline in crops shares as temperatures increase. Geographical shifts can also explain these results in production to more northern counties.

For each model, I predict using the reported coefficients and the new data for each increase in C. From these predictions, I estimate aggregate revenue-per-acre weather-effects by allowing the weather to vary while holding climate constant. I also estimate weather-climate-effects by allowing both weather and climate to vary across counties. Weather-effects captures short-run effects while weather-climate-effects captures long-run effects, which include adaptation practices that may mitigate the harmful effects of temperatures. If long-run adaptation is occurring, weather-climate-effect predictions will be less than the weather effect.

To estimate predictions for the disaggregate revenue-per-acre, I disentangle revenue using individual crop revenue-per-acre in the SUR model. I first predict revenue-per-acre for each crop and multiply the prediction by crop acres to calculate county-level revenue. Next, the estimates from the SUR crop-share are used to make predictions from the effects of crop-switching and estimate crop acres as temperatures changes. The predicted crop acre changes are used as the divisor to revenue at the county-level to estimate the weather-climate-effect for the disaggregated revenue-per-acre. This effect only captures crop-switching practices and does not take into account other adaptation practices that adjust decisions around acres.

Figure 2.11 provides the main results of this paper. The left panel reports the disaggregate revenue-per-acre for the weather effect and weather-climate-effect. The right panel reports the predictions from the aggregate revenue-per-acre. The weather-effect is similar for both the disaggregate and aggregate revenue, which suggests farmers are making decisions across crop choices and are not focused on specific crops. The weather-climate-effect for aggregate and disaggregate revenue-per-acre shows modest improvements from adaptation but does not entirely mitigate the effects of climate change.

There are important implications from these findings. First, crop-switching does appear to mitigate some of the adverse effects of climate change. If farmers continued practices in the short-run, they would not be able to reduce these effects, but by crop-switching between corn, cotton, hay, soybean, and wheat they offset the effects to modest increases; however, not accounting for the extensive margin may bias these results. I do not explicitly control for differences in crop spacing or quality of land, so the results assume farmers plant any of the five crops in equal conditions. Therefore, I run into an endogeneity problem that is difficult to control. Additional research in this area would provide valuable insight into whether crop-switching is beneficial when differences in cropland quality are controlled.

Another import insight is that in the aggregate, there appears to be a slight long-run adaptation, such as planting times, double-cropping, and rotating complementarities of crop systems. In the long-run, farmers are accounting for previous climate conditions and adjusting to mitigate the effects of harmful temperatures. However, these practices are not able to ultimately reduce the impact of climate change and offer only slight improvements.

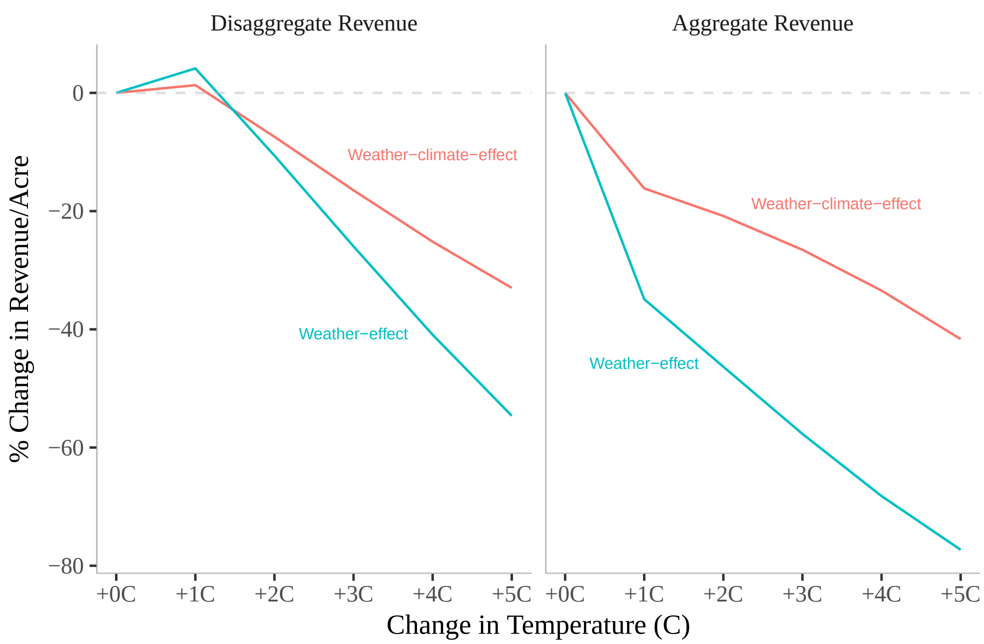

Figure 2.11: **Weather and Climate-effects**

Notes: Figure provides weather and weather-climate effects for a percentage change in revenue-per-acre for our aggregate and disaggregate measure as temperatures increase from 0-+5°C for aggregate and disaggregate measures. Weather-effects allow the weather to vary across increases in temperature, while holding climate constant. The weather-climate effect allows both weather and climate to vary across temperature increases. Temperature increases along the x-axis from 0 - +5°C.

2.7 Conclusion

Adaptation is thought of as an important component in mitigating the impacts of climate change on agriculture production. One of the more popular discussions is how farmers will switch crops to offset damages from increases in temperature. By switching crops, farmers account for short-run weather events, such as planting soybean when winter goes long, or long-run climate events, such as consistent early springs means planting corn early to avoid extreme heat in the summer. However, previous studies have relied on estimating short-run adaptation, which may not translate to long-run adaptation or might be measuring something else entirely. Studies that have estimated long-run adaptation rely on cross-section identification, which introduces omitted variable biases and assumptions about implicit values in the discounted sum of all future net-benefits. These techniques are not able to capture long-run adjustments explicitly.

In this study, I exploit a long history of crop choice and productivity outcomes in the US to understand how short-run weather and long-run climate variation impact each measure. I estimate short-run weather using year-to-year variation and climate using a right rolling mean window. Next, I estimate an aggregate short and long-run effect which captures changes in production practices that are implicit in revenue. I then disaggregate revenue-per-acre by estimating individual crop revenue-per-acre and changes in crop-shares. These estimates are then used to identify changes to revenue-per-acre from crop-switching. Using a uniform increase in temperatures across counties, I predict weather and weather-climate-effects, thus identifying short and long-run sources of adaptation.

This study only considers understanding how farmers have adjusted to weather and climate historically while holding all other factors constant. The analysis only considers predicting changes in uniform temperatures from +1°C to +5°C and the impact on revenue-per-acre and crop choice decisions. However, it is important to note that climate change is expected to impact other factors in the environment, such as precipitation, humidity, vapor pressure deficit, and carbon dioxide concentrations. Agriculture policy, such as crop subsidies and insurance can also change, which will further complicate understanding adaptation mechanisms and the impacts from climate change. While this study only accounts for increases in temperature, hold all else constant, it is important to consider additional factors when accounting for the full impacts of climate change on agriculture.

I assume farmers make decisions based on the intensive margins (revenue-per-acre). I do not account for the extensive margin in the analysis, which introduces endogeneity in crop-shares – cropping and land requirements are not equal. While I allow for complementarities, such as corn and soybean, other crops, such as cotton, are not able to be planted as complements. As a result, there exists an unequal trade-off between crop choice decisions which I do not account for. I allow farmers to plant a variety of the five main crops without limitation or costs. Farmers will, therefore, plant crops with higher revenue-per-acre estimates, which does not account for changes in farmland. I do not address this issue here but suggest it is a critical component when considering changes in crop-shares as climate changes.

Results show the total effect closely matches the sum of predicted impacts from individual crop revenue-per-acre. These results suggest that farmers are making planting decisions based on a combination of crops and not favoring one over the other. I next show that the effects of adaptation appear modest and mostly harmful relative to predictions without climate change. These results suggest short-run adaptation responses by farmers cannot mitigate the harmful effects of climate change, but those long-run adaptation responses may provide slight improvements. The results also suggest adaptation from crop-switching may mitigate the harmful effects of climate change in the long-run. However, these results may be overstating the benefits of crop-switching by not accounting for changes in the extensive margin.

CHAPTER 3
OPTIMAL SPRAYING AND HARVESTING STRATEGY TO COMBAT CBB: A DYNAMIC APPROACH†

3.1 Introduction

The Coffee Berry Borer (CBB), *Hypothenemus hampei*, is one of the most destructive pests to coffee throughout the coffee growing regions of the world, second only to coffee rust. In 2010, the discovery of CBB in Kona, Hawaii resulted in farmers reporting up to 80 - 90% infestation levels of their coffee crop. More recently, bearing acreage is down 16% from the 2013-2014 season to 2015-2016, processors rejected 2.6 million pounds of cherry (ripe fruit containing the coffee bean) in 2015 - 2016 season and, in the same year, value of utilized production was down \$49 million – \$62 million from the 2014-2015 season (USDA 2016). Coffee farmers in Hawaii operate on small margins where costs and uncertainty in production can force farmers out of business (Woodill et al. 2014). With the discovery of CBB, farms are expected to shut down as costs of control increase and production value decreases.

Recommendations for dealing with CBB are outlined in an Integrated Pest Management (IPM) program and include methods for monitoring, spraying, and harvesting, as well as ways to dispose of infested cherry (Kawabata et al. 2017). One of the preferred ways to control CBB is to spray a biological insecticide, *Beauveria bassiana*. At this time, the effectiveness is not well known for different concentrations and spray intervals, especially under different environmental conditions. Spraying is costly regarding application and labor, and farm-level factors such as elevation, terrain, and farm size characteristics add to the uncertainty of controlling for CBB. Preliminary growth rates of CBB in Hawaii have been modeled using a constant growth rate of 35% during a growing season (Woodill et al. 2017). However, many environmental factors influence infestation rates during a growing season, so different infestation rates are needed to account for those factors.

In this paper, we first consider changes in the infestation levels during a coffee growing season using field-level data to track the percentage change in infestation from month-to-month. To model infestation levels, we calibrate a time-inhomogeneous Markov-chain for each month that allows us to determine the effectiveness of spraying and track changes temporally. We model one decision: to spray or not to spray a pesticide, based on the economic impact of that decision. The shift in infestation will be more significant if the farmer decides not to spray, and smaller if he/she chooses to spray; however, the decision to spray also incurs costs. If the expected damage from not spraying is higher than the cost to spray, then it is beneficial to spray.

We then model the expected net-benefit over all months to optimize the decision to spray or not spray throughout a coffee growing season. The calibrated Markov-chain estimates infestation

†This chapter is a result of collaboration with Stuart T. Nakamoto, Andrea Kawabata, and PingSun Leung.

growth rates of CBB for the decision in each month. Each decision leads to separate paths in previous months. A dynamic programming (DP) model is used to estimate the net-benefit (the difference between revenue and total costs) for each path. We then trace out the optimal path of monthly spray/no spray decisions that maximizes net-benefit for the entire season.

A convenient feature within our approach to this paper is that growing coffee in Kona is finite for a season. There is variation from year-to-year, but in general, flowering begins in March or April, the trees start to fruit in May or June, and harvesting runs from August through December. In January, if following best practices, farmers start cleaning their farms, stripping/pruning their trees, and maintaining the farm until flowering begins and cherries form again in March or April[1]. This finite feature allows estimating infestation levels to be isolated to a single growing season since farmers start fresh each year regarding cherry on tree and infestation levels.

Additionally, farm-level decisions made in each month are reliant on decisions made in previous months. The problem of optimizing decisions based on previous months presents an ideal example of dynamic programming (DP). We characterize the problem by maximizing the net-benefit (revenue from cherry harvest minus costs) in each previous month and sum across the growing season. We utilize the variant nature of the Markov-chains to estimate economic damage in each month and model optimal decisions. From this setup, we then estimate economic damages from CBB, the optimal decision in each month, and net-benefits across different infestation levels.

Pest management decisions have been modeled using Markov-chains for a variety of applications. Marcos et al. 2013 discuss the basic concepts for applying Markov-chains to pest behavior in agriculture (see Feller 1968; Heyman and Sobel 1984; Lee et al. 1965 for a formal mathematical outline). Lee et al. 2007 use Markov-chains for zebra mussel invasion in Florida to estimate the net economic impact. In another paper, Lockwood et al. 1988 apply Markov-chains to determine the growth of the rangeland grasshopper populations with pesticide use to improve management in Montana and Wyoming. Other examples include modeling pest control through the introduction of a predator (Kyriakidis 1993) and mitigation and adaptation (Perrings and Walker 2004).

Finally, Haight and Polasky 2010 discuss the problem of imperfect information and the level of infestation. The authors show Markov-chains can be used to determine the impact of different decisions and the power of improved information from farm-level knowledge. For an older survey of applications see White 1993.

Estimating the temporal impacts of pests on agriculture using Markov-chains is a standard approach in the literature. Using this information to determine optimal management practices in a dynamic setting has been discussed in Lee et al. 2007. Other papers use stochastic dynamic programming or optimal control to estimate the impact of invasive species and different policy decisions (Onstad and Rabbinge 1985; Eiswerth and Johnson 2002; Eiswerth and Van Kooten 2002;

[1]If farmers do not follow best practices at the end/beginning of the season, there will be spillover effects from the previous season where CBB is still prevalent in the field in cherries on the ground or trees. As a result, initial infestation levels will be higher and will significantly impact decision making and net-benefit.

Lee et al. 2015).

This paper contributes to the literature by introducing the first estimates for CBB growth rates in Hawaii using farm-level data in Kona and a Markov-chain. The benefit to these estimates is that they can account for farm-level characteristics to identify the effectiveness of spraying or not spraying from month-to-month. These estimates are then used to model decisions to spray.

We then estimate a dynamic programming model to optimize total net-benefit over a growing season. The estimated Markov-chain to model CBB infestation levels is incorporated directly into the DP model to estimate expected damage from infestation and the impact on monthly net-benefits. This approach allows us to trace out an optimal spraying strategy to combat CBB in Hawaii. Next, we compare our economic model against IPM recommendations, the decision to always spray, or to never spray. Our dynamic programming model performs best to optimize final net-benefit.

The paper unfolds as follows: first, we provide background for CBB in Hawaii in Section 3.2; we then outline our Modeling and Empirical Approach in Section 3.3. In Section 3.4, we discuss our main results; Section 3.5 compares and discusses findings with different decisions, and Section 3.6 concludes.

3.2 Background

The Coffee Berry Borer (CBB) is well known in every coffee growing region in the world as one of the most destructive pests. CBB is native to South Africa, but they are have now found in almost every coffee growing region in the world. The main reason CBB is hard to combat is that once they have entered the coffee bean, they are impervious to available pesticides and free to start the next generation. Once CBB has burrowed through the skin and reached the bean, it is almost impossible to eradicate them. Therefore, it is essential for farmers to apply pest management strategies early to mitigate any damage to the bean and further infestation of the crop.

The life-cycle of CBB begin inside the coffee cherry where CBB use the coffee bean as a food source. Initially, an adult female infests a single coffee cherry by burrowing through the outer skin and into the coffee bean. If the bean has matured and hardened enough, the female will begin laying 1-3 eggs per day for the next 15-20 days, and then egg production begins to diminish. The larva will then burrow more deeply into the coffee bean to feed until they mature. In Hawaii, the evolution from egg to adult takes about 20-30 days. As many as 25-30 individual CBB can remain in a single coffee bean, continually reproducing and feeding on the coffee bean. If the bean is overwhelmed with offspring, adult females will leave in search of another suitable coffee bean and repeat this cycle over again (Barrera 2008).

Hawaii was free of CBB infestation until August 2010, when researchers found them in South Kona, Hawaii (HDOA 2010). After the beetle was confirmed, a survey of farms showed the beetle was present in many more regions and had been present for some time. Since then, the beetle has spread to multiple regions on Hawaii Island and neighboring islands Maui and Oahu. It is not yet

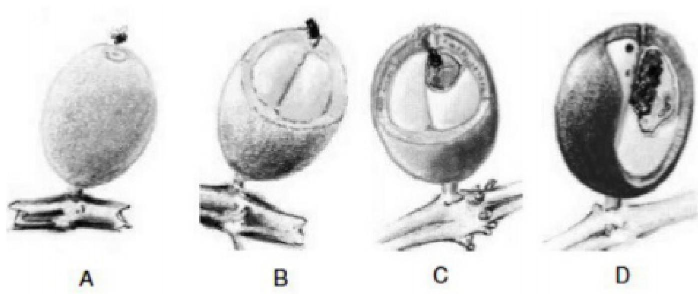

Figure 3.1: CBB Positions
Note: Figure displays four positions CBB can take: A, B, C, D.

clear how CBB made it to Hawaii.

After confirmation of CBB in Hawaii, integrated pest management strategies (IPM) were adapted and updated from other coffee growing areas to mitigate the damage to coffee (Kawabata et al. 2017). These strategies include improved pruning practices, field sanitation, and monitoring, pest spraying using a biological insecticide, and harvesting strategies. While following these guidelines will not completely remove CBB from the crop, they can improve infestation levels at harvest.

The location of CBB in the coffee cherry can take four positions: (1) A, the beetle has landed on the coffee cherry and is beginning to eat away at the skin; (2) B, the beetle has eaten through the outer skin, but has not reached the coffee bean; (3) C, the beetle has now begun eating away at the coffee bean; and (4) D, the beetle has done significant damage to the coffee bean and has already started, or is in the process of, laying eggs (see Figure 3.1 for CBB positions). The movement between each position depends on many variables in the field, such as the maturity of the cherry and environmental factors, such as temperature. In A or B, the CBB is exposed and vulnerable to pesticides. When a coffee berry is immature, CBB will remain in AB for extended months. Once berries mature, CBB can move from A or B to C or D position in a matter of hours, ensuring the demise of the cherry.

A farmer can estimate the level of infestation and CBB positions on their farm by randomly selection 30 trees, sample cherries, cut open the cherries, and inspecting the position of CBB (Kawabata et al. 2017). IPM strategies recommend monitoring the farm 30 days after the first flowering and then to sample at least every two weeks in the early part of the season and monthly after that. This strategy is done to ensure consistent monitoring of the farm and allows the farmer to make decisions based on the levels of infestation.

It is difficult for a coffee farmer to make the critically important decisions about whether to spray or to not spray. Ideally, a farmer will choose to spray when the majority of CBB on the farm are in the AB position because this is when CBB are most vulnerable. When CBB move into the CD position, they are no longer susceptible to spraying efforts and will begin damaging the coffee bean crop. As a result, damaged beans are not marketable and the price received for coffee declines.

3.3 Modeling and Empirical Approach

Decisions about farm-level practices are related to the environment since coffee cherry production is reliant on the current seasonal conditions, such as the timing and abundance of the spring rains and when cherry begins to ripen. Farmers must consider a vast array of different aspects to improve coffee production, thus improving their revenue and reducing costs.

The natural growth rate of CBB is also dependent on environmental factors, such as temperature and humidity. While these vary from year-to-year, in general, warmer conditions will result in increases in the CBB population. The situation on the farm also plays a role in the CBB population, such as cherry on the ground from the previous season, which CBB can use as a food source when the cherry is not on the tree.

Without utilizing CBB management practices on a farm, the CBB population can grow to the point where processors will reject an entire crop (e.g., 25%). Even if not rejected, the price received for damaged coffee (which the processor must then sort) may not pay for even the farmer's harvesting costs. It is, therefore, critical that they have access to reliable data to make informed decisions.

One of the challenging aspects of CBB is understanding the environmental factors that allow them to reproduce and how those factors impact coffee production. Because coffee production occurs on the side of a mountain, we need to factor in elevation effects as well as microclimates due to differences in topography. Hawaii's unique weather patterns also vary from year-to-year and across farms — even within a farm — so conclusively determining how the effects of environmental factors play a role is difficult. There is, however, a direct relationship between coffee cherry production on the farm and CBB infestation: CBB gains new sources of food as coffee trees produce fresh cherry in the field, which allows them to reproduce faster. This relationship also makes a farmer's decision difficult because they need to consider all conditions to make an optimal decision.

Ideally, data on how environmental conditions and farm-level practices affect coffee production, CBB, and infestation levels would provide us with the necessary insights to estimate the economic damages. This information would allow us to establish an ideal model to optimize a profit-maximizing farm. Unfortunately, data is not available that will enable us to determine the environmental link between cherry production, CBB population, and infestation levels. Further, data on farm-level practices and the impact on production and CBB infestation are not available either. To overcome these limitations, we first discuss the ideal model, discuss challenges, and then simplify to an operational model below.

3.3.1 Ideal Model

We start with a model that captures coffee cherry production, $Coffee_t$, in each month, t as a function of weather, w_t and farm-level practices, z_t.

$$Coffee_t = f(w_t, z_t)$$

Weather, w_t, takes into account many of the various environmental factors, such as rain, duration of daylight, temperature spans, and humidity, which directly affects cherry production. It is also essential to account for farm-level practices, z_t, such as trimming/pruning, and fertilizer choice and frequency of application because these decisions also affect the quality and quantity of cherry production. Taking account of these factors provides the direct relationship needed to model cherry growth.

Another critical component is modeling CBB population dynamics, which also carry similar relationships to cherry growth. CBB prefer warmer days and ample moisture, so accounting for these in a similar setup provides a realistic idea about how CBB grow during the season.

CBB population on a farm can be modeled as,

$$CBB_t = g(w_t, Coffee_t, z_t)$$

where CBB in month t is modeled as a function of weather, which accounts for the same dynamics as cherry production, $Coffee_t$. CBB depend on cherry that is available to infest, so increases in cherry imply potential increases in CBB. We also consider farm-level practices, z_t, such as spraying and stripping cherry from trees.

Next, we account for the infestation of CBB in the coffee cherry. There is a direct relationship to cherry production and the CBB population, so we model infestation levels as a function of cherry production and the CBB population,

$$InfestedAmount_t = h(Coffee_t, CBB_t)$$

The infestation amount, $InfestedAmount_t$, is tied directly to the amount of coffee cherry and CBB on the farm, so we can track how much CBB infests cherry on the farm. To calculate infestation levels, divide the total infestation cherry by total amount of cherry on the farm, $Infestation_t = InfestedAmount_t/Coffee_t$. An important point to note here is that the infestation level in each month is related to the amount of available cherry on the farm (i.e., a 1% infestation level on 1,000 lbs of cherry versus 1% infestation level on 10,000 pounds of cherry).

We next account for the different positions CBB take in the coffee cherry. As the IPM outlines, there are four positions CBB can take: A, B, C, or D. These positions also have a time component and vary from month-to-month with CBB attacking newly formed cherry. We want to model each of these positions, $Position_{it}$, as a proportion of the total infestation levels, $\xi_{it}(Infestation_t)$, on the farm. Summing across all positions gives us the total infestation on the farm, $Infestation_t$.

$$Infestation_t = \sum_{i=1}^{4} Position_{it} = \xi_{it}(Infestation_t)$$

At this point, we know the percentage of the damaged crop from CBB in the CD position that is not marketable. We also know the percentage of CBB in the AB position that is likely to move into the CD position. We can now select from $Coffee_t$ and CBB_t to determine the level of infestation on the farm. We interpret CBB infestation positions as the probability of CBB being in A, B, C, D was taken from a sample of the coffee cherry. Selecting from $Coffee_t$ and CBB_t would be analogous, or would substitute, for the grower taking a physical sample from the farm to survey their coffee farm and make appropriate decisions.

Generally, the grower decides whether to spray or not spray after sampling the farm. The IPM provides a lookup table with suggested spray decisions based on a percentage of AB that are alive and the total infested cherry on the farm (Kawabata 2017, Table 1 and 3.5). The IPM focuses on AB alive because this population of CBB in the AB position represents potential movement into CD and because sprays are not effective against CBB in the CD position.

We proposed in a previous paper to model the decision as an economic trade-off that accounts for expected damages from CBB in the CD position from not spraying versus the costs of spraying and the trade-off being spraying versus not spraying (Woodill et al. 2017). If the expected damage from not spraying is greater than the cost to spray, then it is beneficial to spray. From our model, CBB in the CD position comes from the difference in the expected infestation minus the current CD infestation, defined as,

$$E(Damage_{t+1}) = E(Position_{CD,t+1}) - E(Position_{CD,t})$$

The main difference between these two strategies is that the IPM recommendations rely on current levels while our economic decision relies on subsequent month levels. Therefore, our economic model requires data for an entire growing season while the IPM recommended strategy requires data only for the current month.

We next seek to optimize a net-benefit function,

$$\text{Total NB} = \sum_{t=\text{January}}^{\text{December}} NB_t = \sum_{t=\text{January}}^{\text{December}} \{\underbrace{P_t H_t(1 - D_t)}_{\text{Total Revenue}}\} - \{\underbrace{c_s S_t + c_h H_t}_{\text{Total Cost}}\} \tag{3.1}$$

for a coffee growing season from January to December where net-benefit is defined as revenue $P_t H_t(1 - D_t)$, minus total costs, $c_s S_t + c_h H_t$, where revenue equals price, P_t, times the current harvest, H_t, times the remaining percentage of coffee that is marketable $(1 - D_t)$ from the harvest. Total cost is equal to costs to spray, c_s, times decision to spray or not to spray, S_t, plus labor costs to harvest, c_h, times the current harvest, H_t. The price of cherry, P_t, is based on the current infestation of CBB in the CD position and are damaged, D_t (See Table 3.1).

Table 3.1: Harvest Pricing per lbs of Coffee Cherry

Infestation CD	Price per pound.
0-5%	$1.80
6-10%	$1.70
11-15%	$1.60
16-20%	$1.45
21-30%	$1.20
31-40%	$0.60
41-58%	$0.59-$0.35

The decision to spray in the current month, S_t, is true ($S_t = 1$) if the costs to spray, c_s, times remaining months, k, are less than expected revenue loss in cherry without spraying plus subsequent month. The loss is equal to the price in the subsequent month, P_{t+1}, times the expected non-marketable cherry percentage, D_{t+1}, times harvest level, H_{t+1} plus subsequent months, $t+k$. In contrast, if the total costs to spray, $k \cdot c_s$, are greater than expected revenue loss, then the decision is not to spray ($S_t = 0$).

$$S_t = \begin{cases} k \cdot c_s >= P_{t+1}H_{t+1}D_{t+1} + ... + P_{t+k}H_{t+k}D_{t+k} & \text{No spray} \\ k \cdot c_s < P_{t+1}H_{t+1}D_{t+1} + ... + P_{t+k}H_{t+k}D_{t+k} & \text{Spray} \end{cases}$$

We have outlined an ideal model to improve decision making on a farm; however, due to limitations of data availability, we are not able to utilize our ideal model until such information becomes available. To overcome these limitations, we rely on a single season of infestation data and expert knowledge to present a scenario that resembles a typical farm in Kona. We estimate a Markov-chain that models changes in infestation positions, while implicitly accounting for weather and farm-level practices. We then optimize net-benefit using a dynamic programming (DP) approach. While nonlinear programming methods are available to solve our problem, we rely on DP to simplify the approach and to take advantage of the finite nature of coffee production in Kona, Hawaii. We discuss these challenges and our simplified approach in the next section.

3.3.2 Challenges and solutions

Due to data limitations, we seek to simplify our modeling strategy to deliver the most reliable estimates for the decision-making process. Data on weather and farm-level practices that would allow us to estimate population dynamics of CBB and coffee cherry are not available. However, we do have field-level data where farmers sample their farm and report infestation levels. As a result, we can assume the weather and farm-level practices are external factors that are implicit in the data.

To simplify the functional form $Coffee_t$, we assume cherry production follows a logistic growth function (see Figure 3.2). The logistic growth function is used in two ways: (1) to account for the amount of coffee cherry CBB can attack each month; and (2) the amount harvested in each month. At the beginning of the season, there will be a minimal cherry on the farm. As the season continues, trees will flower, fruit, and cherry production begins to increase rapidly starting in the summer. This simplification allows us to account for the total cherry on the farm at any given month t. The amount harvested is also removed from the cherry available on the farm.

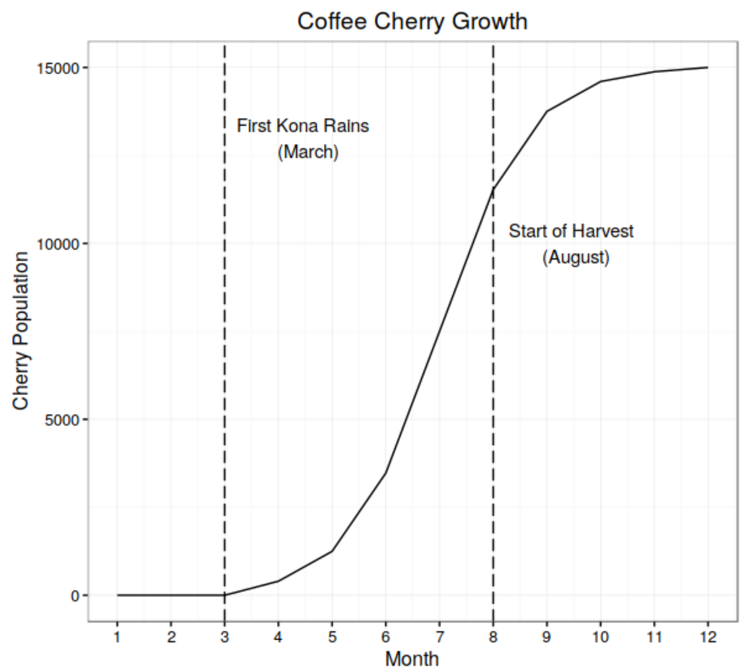

Figure 3.2: **Coffee Cherry Growth**
Note: Figure displays logistical growth function to estimate coffee cherry population dynamics from January to December.

CBB population dynamics will follow cherry production on the farm since cherry is the food source. As more cherry becomes available, CBB will leave their host cherry and infest newly formed cherry. To model CBB_t, we assume CBB population dynamics are similar to our functional form for cherry production (see below for calibrating CBB population dynamics).

One of the biggest challenges is accounting for cherry production and CBB population dynamics to make inferences about infestation levels on a farm. Without knowing current infestation levels, farmers have limited information that allows them to make optimal decisions. From our simplification of coffee production and CBB population, we can simplify infestation levels by utilizing a Markov-chain setup that tracks infestation levels for a farm using available farm-level data. This setup allows us to estimate the proportional change in each position during the season, thus allowing us to model $Infestation_t$ directly. From here, we assume a sample from the farm will provide us with the percentage of infested cherry in each position. We use this farm-level knowledge as a basis for decision making.

We next utilize the power of a dynamic programming (DP) setup that allows for a stage-wise structure that can deal with month-by-month solutions based on previous months and also allows easy integration of our Markov-chain simplification. We optimize a monthly net-benefit objective function that relies on the previous month's results. We also account for infestation levels in our Markov-chain setup to estimate expected damages from spraying versus not spraying. A forward-recursive DP model then estimates each previous month's optimized net-benefit, then summing across the entire optimized net-benefit for the growing season. Our final result will provide an optimal spray strategy and final net-benefit.

3.3.3 Infestation and Markov-chains

We use a Markov-chain to model temporal changes and to account for the external factors implicitly through field-level data. The infestation level for the current month utilizes information on the infestation level of previous months, which makes a Markov-chain a natural fit. We explicitly model the changes in infestation levels for each position (state) by using a time-inhomogeneous Markov-chain to identify the movement of CBB.

The Markov-chain estimates the probability of CBB moving into each position and adjusts current infestation levels to reflect each month change. For example, a sample of 100 cherries from the field is collected and dissected. After dissecting the cherries, 20 are in the AB position and ten are in the CD position. Next month, another 100 dissected cherries show 15 are in the AB position and 15 are in the CD position. From the first month to the next month five cherries have moved from the AB position into the CD position. Therefore, 5% of the cherries will move from AB to the CD position. This behavior can be modeled using a Markov-chain to identify these changes for each position (state) in each month (month) as well as the percentages in each position.

The Markov-chain can also model the changes in infestation levels based on a farmer's decision to spray or not spray. When a farmer decides to spray, there is a lower probability of CBB moving into the CD position than if the farmer decides not to spray. We identify two separate Markov-chains based on the decision to spray or not spray. An important reason we use time-inhomogeneous Markov-chains is that it allows changes between months to be variant (e.g., change from AB to CD will be different between March-April and April-May). Each month will have a different level of change and the infestation levels adjusted based on the decision to spray or not spray.

Formally, a time-inhomogeneous Markov-chain is defined as,

$$\mathbb{P}(X_t = x_t | X_{t-1} = x_{t-1}, X_{t-2} = x_{t-2}, ..., X_0 = x_0) = \mathbb{P}(X_t = x_t | X_{t-1} = x_{t-1})$$

where the probability, $\mathbb{P}$, of a stochastic process, X, in month t is equal to x_t conditional on the previous month's stochastic process, X_{t-1}, which is equal to x_{t-1} in the previous month until $t = 0$. The intuition here is that the current state is based on the previous month's infestation level until the state returns to the beginning, t_0. Therefore, at X_0 initial values are estimated to start the stochastic process through t. In terms of infestation levels, the stochastic processes, X_t, are the different positions of CBB and the path throughout a growing season's $t = 0, \ldots, T$. Initial values are estimated as,

$$V_0 = \begin{bmatrix} v_{1,0} & v_{2,0} & v_{3,0} & v_{4,0} \end{bmatrix}$$

where vector V_t contains four elements (probability in each state that sum to 100%): (1) $v_{1t} = NI$: % not infested or those cherries with no holes, (2) $v_{2t} = ABL$: those cherries with a hole and have live CBB in the AB position, (3) $v_{3t} = ABD$: those cherries with a hole and have dead or missing CBB, and (4) $v_{4t} = CD$: cherries with a hole and have CBB in the CD position.

$$V_0 = \begin{bmatrix} NI_0 & ABL_0 & ABD_0 & CD_0 \end{bmatrix}$$

Next, we define two sets of transition matrices for spraying (SP_t) and not spraying (NSP_t) where each matrix defines a probability for each month, t, with event probabilities, a_{ijt} and b_{ijt}, in the probability space. Formally, the transition matrices are defined as,

$$SP_t = \begin{bmatrix} a_{11t} & a_{12t} & a_{13t} & a_{14t} \\ a_{21t} & a_{22t} & a_{23t} & a_{24t} \\ a_{31t} & a_{32t} & a_{33t} & a_{34t} \\ a_{41t} & a_{42t} & a_{43t} & a_{44t} \end{bmatrix}$$

$$NSP_t = \begin{bmatrix} b_{11t} & b_{12t} & b_{13t} & b_{14t} \\ b_{21t} & b_{22t} & b_{23t} & b_{24t} \\ b_{31t} & b_{32t} & b_{33t} & b_{34t} \\ b_{41t} & b_{42t} & b_{43t} & b_{44t} \end{bmatrix}$$

For SP_t, the movement from state i to state j in a_{ijt} is the probability that the current state, or infestation level, will transition to a different state in a_{ijt}. For example, NI to ABL is defined as a_{11t} to a_{12t}, from ABL to ABD is a_{22t} to a_{23t}, and ABL to CD is a_{22t} to a_{24t}. However, movements in ABD or CD are not possible since CBB is already dead or missing in the ABD position and once a cherry is damaged in CD it cannot be undone; therefore, a_{33t} and a_{44t} are defined as one in the matrices. Zeros in the matrices prevent impossible movements, such as CBB that are in AB position and dead to CD. For not spraying, the matrix elements determine the same position movements at SP_t, although, probabilities between movements may be higher due to not spraying. New transition matrices with these assumptions are defined as,

$$SP_t = \begin{bmatrix} a_{11t} & a_{12t} & a_{13t} & a_{14t} \\ 0 & a_{22t} & a_{23t} & a_{24t} \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

$$NSP_t = \begin{bmatrix} b_{11t} & b_{12t} & b_{13t} & b_{14t} \\ 0 & b_{22t} & b_{23t} & b_{24t} \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

To track the current levels of infestation, vector V_t is defined as,

$$V_t = \begin{bmatrix} NI_t & ABL_t & ABD_t & CD_t \end{bmatrix}$$

where each position of CBB is defined as an element in each month t. The choice to spray is a binary decision (zero or one), S_t, is equal to one if the farmer decides to spray, and zero if the farmer decides not to spray. Therefore, to estimate the current infestation, given a choice decision, the vector, V_t, is equal to,

$$V_t = V_{t-1} \cdot [S_t \cdot SP_t] + V_{t-1} \cdot [(1 - S_t) \cdot NSP_t]$$

To find the non-marketable cherry, CD, in the current month, we identify the damage, D_t, as the fourth element of the vector V_t,

$$D_t = CD_t$$

To demonstrate how this Markov-chain works, a simple matrix algebra example for month one is estimated as,

$$V_1 = V_0 \cdot [S_1 \cdot SP_1] + V_0 \cdot [(1 - S_1) \cdot NSP_1]$$

$$D_1 = CD_1$$

where a vector, V_0, is equal to the initial infestation levels of each position of CBB, times the choice to spray ($S_1 = 1$) or not spray ($S_1 = 0$) using the transition matrices defined above. The current non-marketable cherry is equal to the element in the vector where $D_1 = CD_1$. For month two, the current infestation levels are equal to,

$$V_2 = V_1 \cdot [S_2 \cdot SP_2] + V_1 \cdot [(1 - S_2) \cdot NSP_2]$$

$$D_2 = CD_2$$

where a vector of infestation levels, V_2, in month two uses the previous infestation levels, V_1, to estimate current levels based on the decision to spray or not spray utilizing the transition matrices in month two. The vector V_t can then be multiplied by harvest amount to get the quantities of cherry that are NI, ABL, ABD (all acceptable in the market) and CD (damaged and not marketable, in our model).

3.3.4 Calibration

Ideally, to calibrate these Markov matrices, we would want two identical farms in the same location where one sprayed all year and the other did not. This data provides us with changes in infestation when choosing to spray and when deciding not to spray. A comparison of the two farms would give us the growth rate difference resulting from the primary decision to spray. Unfortunately, data are not available – no two coffee farms are the same and the landscape is different between farms.

To overcome this limitation, we rely on field-level data collected during May-December 2016 from farmers in Kona, Hawaii, to calibrate the Markov-chains. In each month, the sampling procedure records the field-level infestation and the decision to spray or not. We also rely on experts who have been studying the growth patterns of CBB in Hawaii through sampling. By combining data and expert knowledge, we can estimate the change in the growth rate of CBB from spraying and not spraying.

Data that we do not explicitly have is observations for the early part of the year (January - May) and infestation levels when a farm does not spray every month. First, we focus on filling in the spray pattern for the early part of the year from January to May. We interpolate to an expected initial infestation based on what field-level experts would suggest; that is, AB live should be higher in the early part of the year because of fewer cherries on the tree and slowly declining as cherries fill in. We then link these estimates to the field-level data we have for May. These estimates provide infestation levels from January to December for a farm that always sprays.

Next, we calibrate infestation levels for a farm that never sprays. Again, we rely on expert field level knowledge to obtain these infestation levels by calibrating a spray pattern as described above, change the initial and end infestation levels to the expectation when not spraying, and interpolate between the values. While not ideal, these are the best estimates we can obtain for a farm that does not spray during the season.

A maximum likelihood estimator and the data described above are used to calibrate the Markov-chains based on the likely probability of moving between each state given the data[2]. Formally, a

[2]Maximum likelihood description is taken from the vignette documentation and calibration of the Markov-chain matrices from the R package `markovchain` (Spedicato 2017).

maximum likelihood estimator calculates the probability of a change in the field-level infestation, $\hat{p}_{ijt}$, where n_{ijt} represents a number in the sequence of the current state, X_i, that contains a set of Markov-chain state spaces, s_t. Therefore, the current state depends on the next state as, $X_i = s_t$ and $X_j = s_{t+1}$ in between time, t and $t+1$. The estimate calculates the proportion of the sequence of infestation level divided by the sum of total sequence for every infestation level.

$$\hat{p}_{ijt}^{\text{MLE}} = \frac{n_{ijt}}{\sum_{u=1}^{k} n_{iut}}$$

The calibration technique produces two sets of Markov-chain transition matrices (decision to spray and to not spray) for the nine months in a growing season that includes the four positions, NI, ABL, ABD, and CD (see Figure 3.3). We can now estimate infestation levels in each month based on whether a farmer decides to spray and track these levels through time. We use these estimates in our dynamic model to incorporate optimal farm-level decision making.

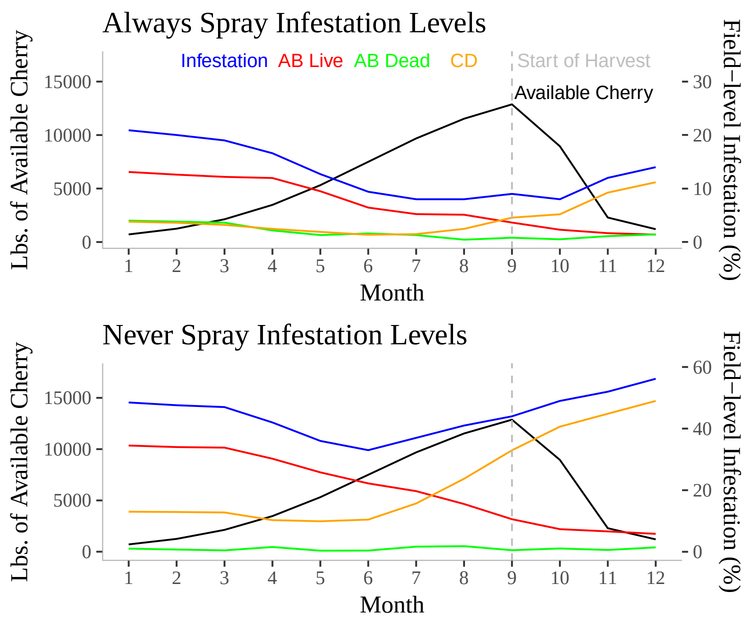

Figure 3.3: **Markov-chain Calibration Data**
Note: Figure displays data used to calibration Markov-chains for tracking CBB positions and infestation levels. Data was generated from field-level observations for 2015 coffee growing season and expert knowledge. Available cherry represents cherry that is available on the farm and is ready for harvest.

3.3.5 Dynamic Programming

In this section, we develop a dynamic programming economic model where the optimal strategy for a farmer is the set of monthly spray decisions that maximize the net-benefit over the entire growing season. Within this dynamic framework, we also account for available cherry through a logistic growth function and then optimize harvesting of cherry during a month of harvest. One advantage of using dynamic programming to address the optimal strategy for a coffee farmer is that it fits nicely into decisions farmers make on a monthly basis. The real strength of the dynamic programming model is that it can determine a set of monthly decision strategies based on the impact of previous decisions while also accounting for expected subsequent infestation levels using our Markov-chains. As a result, we are able to model an entire season of decisions and extract optimal spraying and harvesting strategies.

We start with determining available ripe cherry to harvest. To model this, we need to account for the transition from flowering, to green cherry, to ripe cherry, to overripe cherry (raisins), which introduces complexity beyond the scope of this paper due to limitations in data. For simplification, we assume ripe cherry growth is a logistic growth function, G_t that provides available cherry to harvest,

$$G(K, r, t) = \frac{K}{1 + e^{-rt}}$$

where K is the total expected cherry on the farm at the end of the season, r, is the steepness of the curve, and t is the time component.

We then harvest from available cherry on the farm and assume there is a proportion of total cherry, $\rho(c)$ harvested in September, October, November, and December (32%, 48%, 12%, 8%, and 0% for all other months). In H_t we define percentages that are infested in each position and price based on CD infestation level[3]. The amount harvested in each month equals,

$$H_t = \rho(c) \cdot K$$

On a farm, as the amount of harvested cherry increases, available ripe coffee cherry declines. To account for this, the amount harvested, H_t is subtracted from current available cherry G_t.

To estimate damage to harvested cherry from CBB, we use the results from the Markov-chain above. The proportion of CBB in the CD position is defined as $D_t = CD_t$. To calculate the total amount of harvested cherry in the CD position, the current CD infestation level is multiplied by the current amount harvested, H_t.

Our objective function, net-benefit, is defined as the revenue generated from cherry crop minus any costs (see equation 3.1). The revenue includes reduction due to economic damages from CBB and costs include spraying and harvesting. The economic damage from CBB in the CD position, $D_t = CD_t$, is obtained from the vector that tracks infestation levels, V_t. We utilize this result to account for damages to revenue. Spraying costs are included if the decision to spray is made. Harvesting costs include a labor rate applied to the amount harvested.

[3] At the mill, a sample of harvested cherry is collected, dissected, and CBB position infestation levels are calculated to price the harvest. We model this behavior in our DP setup to provide a realistic scenario for coffee farms.

We now derive a dynamic programming model utilizing all components discussed above. Formally, a value function, $f()$, in month one is equal to the net-benefit in period one, NB_1, given harvest, H_1, and current levels of infestation, V_1,

$$f_1^* = NB_1(H_1, V_1)$$

moving forward to month two carries with it the optimal results from previous month one, f_1^*, which includes total harvested, H_1, and infestation levels based on the decision to spray or not spray, V_1. We now define month two as,

$$f_2^* = \max_{NB}\{NB_2; \beta\{NB_2(H_2, V_2) + f_1^*\}$$

where the net-benefit is maximized for month two, NB_2, with a discount factor, β, plus the optimal net-benefit from month one, f_1^*. The dynamic nature of the model includes previous optimal value function and optimization in the current month; thus, the Bellman equation can be written as,

$$f_t^* = \max_{NB}\{NB_t; \beta\{NB_t(H_t, V_t) + f_{t-1}^*\}$$

where the optimal function in month t is defined as f_t^*, which maximizes the current month, t, plus previous optimal value function, f_{t-1}^* is given a discount factor β.

An essential feature of this model is that we consider the variant nature of infestation levels between months and base the decision to spray on whether the damage to cherry from not spraying is higher than the cost to spray. By optimizing the net-benefit given the level of infestation, V_t, and expected changes in CD_{t+1}, we compare the results in the next month to determine this decision. As a result, we derive an optimal decision path for a coffee growing season.

3.4 Results

Our main results utilize parameters in Table 3.2 which describe a typical farm in Kona with two acres of coffee farmland and a projected yield of 7,500 pounds of cherry per acre. Farm labor per hour equals \$15 and harvest labor is \$0.50 per pound of cherry. If the farm decides to spray, we assume one spray per month with a total cost of \$214[4].

Initial infestation levels are needed to start our Markov-chain in the first month. We utilize the first-month infestation levels reported for a farm that always sprayed. We assume the farm followed best practices to ensure low initial infestation levels at the beginning of the season follows: 5.5% AB live, 2.5% AB dead, 2% CD, and 90% NI.

Our economic results are described in Table 3.3 and Figure 3.4. The optimal spray schedule is

[4]Total costs include pesticide costs (\$140) plus labor costs (\$30) plus water (\$20) plus surfactant for Beauveria bassiana (\$24) equals \$214. We do not account for sampling/monitoring costs in our economic model.

Table 3.2: Model Parameters for a Typical Farm in Kona, Hawaii

Parameter	Unit	Estimate
Acres	Acres	2
Projected Cherry	Lbs. per acre	7,500
Farm Labor	Dollars per hour	15
Spray Labor	Hours per acre	1
Harvest Labor	Dollars per lbs.	0.5
Pesticide	Quart per acre	1
Pesticide Costs	Dollars per quart	70.35
Water	Gallons per acre	100
Water Cost	Dollars per 1k gallons	1
Surfactant	Ounces per acre	45
Surfactant Costs	Dollars per quart	8

not to spray January - May, spray from June - November, and not spray in December. The damage to the crop as a result of the spray schedule totals 562 pounds of cherry. Due to changes in CD infestation levels, price varies from \$1.80 to \$1.70 per pound of cherry (see Table 3.1). The projected loss in revenue from coffee cherry damage is \$1,012. The total net-benefit for a typical farm in Kona is \$17,916. Table 3.4 provides Infestation levels for each month. The final CD field-level infestation level in December is 9.23%; although, as described previously this level is based on the available cherry on the farm and does not represent the total CD infestation over the season – loss in revenue accounts for varying damage in each month and is a better indicator of damages to the cherry on a farm.

Table 3.3: Economic Model Results

Month	Spray Decision	Harvested Cherry	Harvested Damage	Harvested Cost	Spray Cost	Price	Net-benefit	Net-benefit (Cum. Sum)
1	No Spray	0	0	0	0	1.8	0	0
2	No Spray	0	0	0	0	1.8	0	0
3	No Spray	0	0	0	0	1.8	0	0
4	No Spray	0	0	0	0	1.8	0	0
5	No Spray	0	0	0	0	1.8	0	0
6	Spray	0	0	0	1	1.8	-214	-214
7	Spray	0	0	0	1	1.8	-214	-428
8	Spray	0	0	0	1	1.8	-214	-642
9	Spray	4,800	114	2,400	1	1.8	6,026	5,384
10	Spray	7,200	208	3,600	1	1.8	9,146	14,530
11	Spray	1,800	128	900	1	1.7	1,946	16,476
12	No Spray	1,200	110	600	0	1.7	1,440	17,916

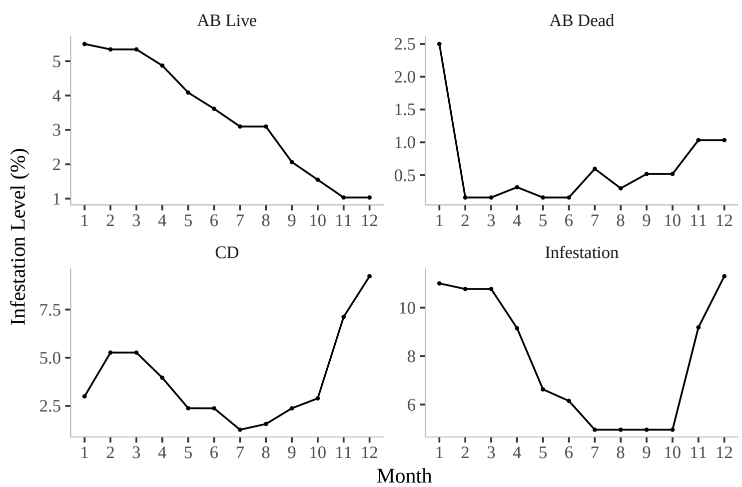

Figure 3.4: **Economic Model Infestation Levels**
Note: Figure displays CBB infestation levels from economic model results from January through December for AB Live, AB Dead, CD, and Total Infestation.

Our results suggest that if best practices are followed to ensure a low initial infestation level then spraying in the early part of the season is not necessary. This result is from there being fewer cherries on the coffee trees and ground that CBB can use as a food source. Because of expected damages, it is beneficial to begin spraying when a farm expects cherry growth to start increasing in the next month (around May or June), thus providing a food source for CBB to begin populating cherry. Our economic model captures this by not spraying January-May where field-level CD infestation

Table 3.4: Field-level Infestation Results from DP model

Month	Spray Decision	AB Live (Field)	AB Dead (Field)	CD (Field)	Infested (Field)
1	No Spray	5.5	2.5	3	11
2	No Spray	5.34	0.16	5.27	11
3	No Spray	5.34	0.16	5.27	11
4	No Spray	4.87	0.31	3.96	9
5	No Spray	4.09	0.16	2.38	7
6	Spray	3.61	0.16	2.38	6
7	Spray	3.1	0.59	1.27	5
8	Spray	3.1	0.3	1.57	5
9	Spray	2.07	0.52	2.38	5
10	Spray	1.55	0.52	2.9	5
11	Spray	1.03	1.03	7.12	9
12	No Spray	1.03	1.03	9.23	11

levels are decreasing (due to increases in cherry growth), and a low final CD infestation level of 9.23%.

3.5 Discussion

Our economic model described in this paper seeks to optimize a net-benefit function in previous months given CBB infestation levels based on expected damage from not spraying versus spraying. If the expected damage from not spraying is higher than the cost to spray, then it is beneficial to spray. However, this strategy is not practiced in the field because farmers do not have sufficient information about the following months. In this section we identify three possible scenarios: IPM choice, always spray or never spray. Under IPM choice, growers will sample their farm to determine the level of AB live versus total infestation to decide whether to spray or not spray. We estimate this strategy using our DP model setup and the IPM recommendations outlined in Figure 3.5. We also run our DP model set up where a farmer decides to spray every month, and never spray. We then compare our economic model against alternative strategies.

% A/B Alive

% infestation	1%	2%	3%	4%	5%	10%	15%	20%	25%	30%	35%	40%	45%	50%	55%	60%	65%	70%	75%
1%	0.01	0.02	0.03	0.04	0.05	0.1	0.15	0.2	0.25	0.3	0.35	0.4	0.45	0.5	0.55	0.6	0.65	0.7	0.75
2%	0.02	0.04	0.06	0.08	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1	1.1	1.2	1.3	1.4	1.5
3%	0.03	0.06	0.09	0.12	0.15	0.3	0.45	0.6	0.75	0.9	1.05	1.2	1.35	1.5	1.65	1.8	1.95	2.1	2.25
4%	0.04	0.08	0.12	0.16	0.2	0.4	0.6	0.8	1	1.2	1.4	1.6	1.8	2	2.2	2.4	2.6	2.8	3
5%	0.05	0.1	0.15	0.2	0.25	0.5	0.75	1	1.25	1.5	1.75	2	2.25	2.5	2.75	3	3.25	3.5	3.75
10%	0.1	0.2	0.3	0.4	0.5	1	1.5	2	2.5	3	3.5	4	4.5	5	5.5	6	6.5	7	7.5
15%	0.15	0.3	0.45	0.6	0.75	1.5	2.25	3	3.75	4.5	5.25	6	6.75	7.5	8.25	9	9.75	10.5	11.25
20%	0.2	0.4	0.6	0.8	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
25%	0.25	0.5	0.75	1	1.25	2.5	3.75	5	6.25	7.5	8.75	10	11.25	12.5	13.75	15	16.25	17.5	18.75
30%	0.3	0.6	0.9	1.2	1.5	3	4.5	6	7.5	9	10.5	12	13.5	15	16.5	18	19.5	21	22.5
35%	0.35	0.7	1.05	1.4	1.75	3.5	5.25	7	8.75	10.5	12.25	14	15.75	17.5	19.25	21	22.75	24.5	26.25
40%	0.4	0.8	1.2	1.6	2	4	6	8	10	12	14	16	18	20	22	24	26	28	30
45%	0.45	0.9	1.35	1.8	2.25	4.5	6.75	9	11.25	13.5	15.75	18	20.25	22.5	24.75	27	29.25	31.5	33.75
50%	0.5	1	1.5	2	2.5	5	7.5	10	12.5	15	17.5	20	22.5	25	27.5	30	32.5	35	37.5

=0-0.99 – Spraying not recommended; will cost more than the expected value of coffee saved from CBB
=1-1.99-Consider spraying, especially early in the season
=2-4.99 – Especially early in the season, this is a critical level to start spraying to avoid economic loss.
=5-9.99 – You are starting to lose money due to CBB damage. Losses will be greater if you don't spray.
=10-19.99 – You are losing money due to CBB damage, but you may still want to spray.
=>20 – Processors may reject your harvest. The value of your harvest may not cover picking cost, so consider focusing on your next crop (i.e. strip pick, stump prune)

Figure 3.5: **IPM Recommendations based on AB**
Note: Figure displays IPM recommendations based on dissected AB infestation and total infestation levels.

Each strategy relies on different assumptions and costs. Our economic model relies on good (perfect) information and expectations about future infestation levels. With sufficient data, the economic model will provide the optimal spraying strategy and monitoring/sampling are not needed. However, the IPM choice strategy requires constant monitoring/sampling each month to make spray decisions, so we include the associated costs of 2-labor hours per acre ($30). When a farm decides to spray regardless of information or sampling/monitoring results, they incur only costs to spray. This strategy of spraying on a schedule can be considered a mechanism to cope with inadequate information. We also compare against a farm that decides never to spray, such as an abandoned farm where the farmer allows the farm to produce coffee unmanaged. We examine each strategy against these assumptions to see which performs best regarding net-benefit.

Figure 3.6 and 3.7 provides the results from each model run. Our economic model performs best when optimizing net-benefit, with IPM choice performing next best, then always spray, and never spray. When deciding to never spray, final net-benefit is low due to the high levels of infestation in CD position and low price for cherry. Final CD infestation levels when deciding to always spray is lowest (8.76%), followed by the economic model (9.23%), IPM choice (14.97%), and never spray (40.99%). Marketable cherry describes cherry that is free from CD damages at the mill. Always spraying provides the highest level of marketable cherry (14,512 pounds), followed by the economic model (14,438 pounds), IPM choice (14,437 pounds.), and then never spray (10,984 pounds). CD damage is the difference between full harvest and marketable cherry, which follow the same order.

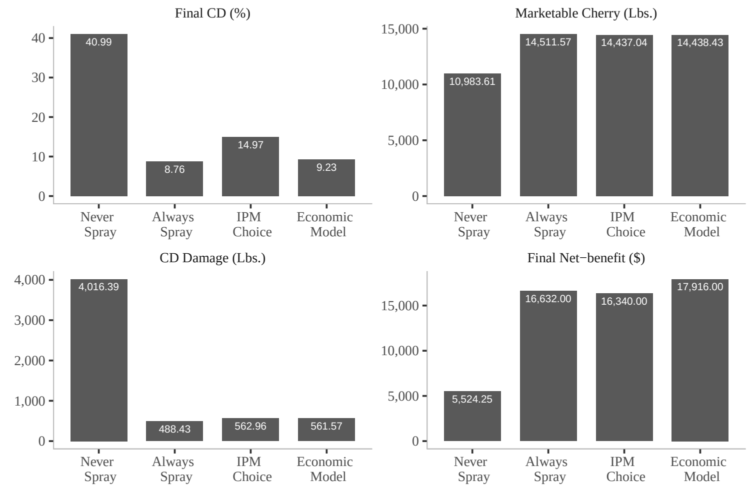

Figure 3.6: **Comparing Model Results**
Note: Figure displays results from our economic model, IPM choice, always spray, and never spray for final CD infestation levels, marketable cherry, damaged cherry, and total net-benefit.

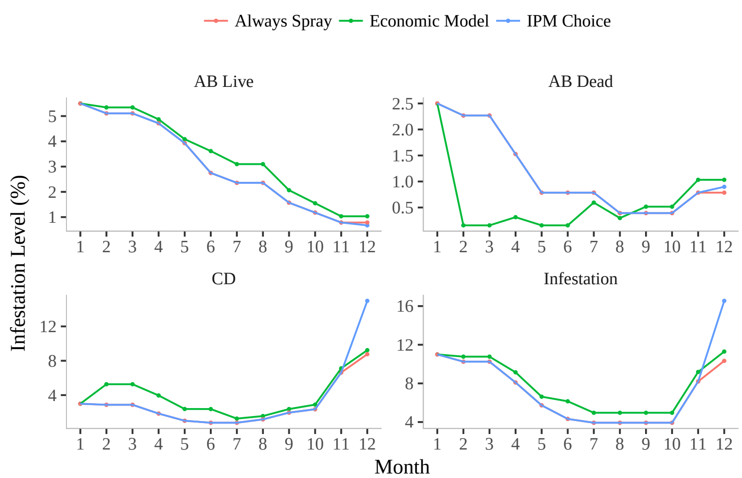

Figure 3.7: **Comparing Infestation Levels by Decision Model**
Note: Figure displays CBB infestation levels for decisions Economic Model, IPM Choice, and Always Spray. CBB Infestation levels are reported from January through December.

While our economic model performs best, the difference from always spraying is $1,284. This result suggests that the costs from monitoring/sampling in prior seasons to gather useful information about infestation levels for use in the economic model may not be worth it. If the costs to collect the data for use in the model are higher than $1,284, then it would be best to always spray the farm each month. However, if we are relying on a single season of data, it may still be beneficial to monitor/sample to gather useful information for the farm. If better knowledge about current environmental conditions portrays a light season of CBB – cooler conditions – then it may not be beneficial to spray in each month because the natural environment reduces CBB population without intervention.

Additionally, the difference between our economic model and IPM choice is $1,576. A large portion of the difference is in the monitoring/sampling costs, which equals $540 (2-acres x 2-hours x $15-labor rate x 9-months of spraying). This difference further suggests our economic model is the preferred decision strategy.

3.6 Conclusion

Understanding CBB infestation levels in Hawaii is critical when determining economic damages to a farm and potential revenue losses. In this paper, we have provided first estimates for calibrating and modeling CBB infestation levels to calculate damage to coffee production. Using a time-inhomogeneous Markov-chain, we track changes in infestation levels during each month of the coffee growing season, based on the decision to spray or not spray. We then incorporate these estimates into a forward-recursive dynamic programming model that captures a farmer's expected net-benefit in each month. We also extract an optimal decision path for spraying or not spraying based on the trade-off between the cost to spray and expected damage from not spraying. The results suggest it is best not to spray January-May, then spray from June-November. The final CD infestation level is 9.4% and total net-benefit is $17,916. We then compare our economic model against IPM choice, deciding to always spray or to never spray. These results confirm our economic model provides the highest net-benefit.

Our previous work utilized a decision tree analysis to optimize spraying decisions through a growing season using a constant growth rate of CBB infestation (Woodill et al. 2017). We have improved upon this previous work by modeling a time-inhomogeneous Markov-chain to account for the temporal changes in infestation levels from deciding to spray or to not spray. This approach accounts for implicit factors in the infestation level, such as environmental conditions and farm-level practices. Additionally, this paper utilizes a dynamic programming model to account for the previous month's decision while also accounting for subsequent month's infestation levels. While our decision tree approach accounted for each period's optimal net-benefit, our dynamic approach allows us to recursively estimate decisions through a net-benefit function and optimize over the whole growing season.

This paper focuses on individual farm-level decisions and neglects potential spillover effects from spray decisions by neighbors. There are concerns about the negative impacts on infestation levels by neighbors, especially for feral and abandoned or poorly managed farms. If CBB is widespread on farm borders they could begin infesting and spreading through the well-managed farm. Accounting for beneficial or negative spillover effects from neighbors spraying or not would provide a more complete picture. However, it is not yet known if and by how much spillover effects play in affecting infestation levels.

Our analysis relies on a single-season of data and knowledge from Hawaii-coffee experts to model infestation levels. The data reports infestation levels and the position of CBB in each month. Unfortunately, the data does not provide environmental conditions or CBB population dynamics to model interactions. Ideally, we prefer multiple season data that include environmental conditions, CBB population dynamics, and sampling/monitoring results from farms with different characteristics, such as elevation changes, farm management practices, and larger sample observations. To overcome these limitations, we utilize the data available and expert knowledge to model infestation levels in a Markov-chain setup that accounts for these dynamics implicitly. The transition to each CBB position in each month estimates the probability of movement based on the decision to spray or not spray. Our simplification provides the first estimates of CBB population dynamics and how decisions affect changes in each month. Ongoing research and data collection in Hawaii will only improve our modeling results, thus tightening the relationships we seek to identify in future research.

To conclude, we provide the results of our economic dynamic programming model and the optimal decision path to optimize net-benefit in each month. We provide first estimates of CBB population dynamics and infestation levels using Markov-chain setup. We then compare against IPM choice, always spray or never spray. Modeling these results in a dynamic setting allows for a better understanding of the processes underlying decisions at the farm-level and how parameters affect the final net-benefit for farmers.

BIBLIOGRAPHY

Adams, R. M., C. Rosenzweig, R. M. Peart, J. T. Ritchie, B. A. McCarl, J. D. Glyer, R. B. Curry, J. W. Jones, K. J. Boote, and L. H. Allen Jr (1990). Global climate change and US agriculture. *Nature 345*(6272), 219.

Arnold, C. Y. et al. (1960). Maximum-minimum temperatures as a basis for computing heat units. In *Proceedings. American Society for Horticultural Science*, Volume 76, pp. 682–92.

Barrera, J. F. (2008). Coffee pests and their management. In *Encyclopedia of Entomology*, pp. 961–998. Springer.

Baskerville, G. and P. Emin (1969). Rapid estimation of heat accumulation from maximum and minimum temperatures. *Ecology 50*(3), 514–517.

Burke, M. and K. Emerick (2016). Adaptation to climate change: Evidence from US agriculture. *American Economic Journal: Economic Policy 8*(3), 106–40.

Cai, R., J. D. Mullen, M. E. Wetzstein, and J. C. Bergstrom (2013). The impacts of crop yield and price volatility on producers' cropping patterns: A dynamic optimal crop rotation model. *Agricultural systems 116*, 52–59.

Cohn, A. S., L. K. VanWey, S. A. Spera, and J. F. Mustard (2016). Cropping frequency and area response to climate variability can exceed yield response. *Nature Climate Change 6*(6), 601.

Dell, M., B. F. Jones, and B. A. Olken (2012). Temperature shocks and economic growth: Evidence from the last half century. *American Economic Journal: Macroeconomics 4*(3), 66–95.

Deschênes, O. and M. Greenstone (2007). The economic impacts of climate change: Evidence from agricultural output and random fluctuations in weather. *American Economic Review 97*(1), 354–385.

Deschênes, O. and M. Greenstone (2011). Climate change, mortality, and adaptation: Evidence from annual fluctuations in weather in the US. *American Economic Journal: Applied Economics 3*(4), 152–185.

Eiswerth, M. E. and W. S. Johnson (2002). Managing nonindigenous invasive species: Insights from dynamic analysis. *Environmental and Resource Economics 23*(3), 319–342.

Eiswerth, M. E. and G. C. Van Kooten (2002). Uncertainty, economics, and the spread of an invasive plant species. *American Journal of Agricultural Economics 84*(5), 1317–1322.

Feller, W. (1968). *An Introduction to Probability Theory and its Applications: volume I*, Volume 3. John Wiley & Sons New York.

Fisher, A. C., W. M. Hanemann, M. J. Roberts, and W. Schlenker (2012). The economic impacts of climate change: Evidence from agricultural output and random fluctuations in weather: Comment. *American Economic Review 102*(7), 3749–60.

Guo, C. and C. Costello (2013). The value of adaption: Climate change and timberland management. *Journal of Environmental Economics and Management 65*(3), 452–468.

Haight, R. G. and S. Polasky (2010). Optimal control of an invasive species with imperfect information about the level of infestation. *Resource and Energy Economics 32*(4), 519–533.

HDOA (2010). Serious coffee pest detected in Kona.

Henderson, J. V., A. Storeygard, and U. Deichmann (2017). Has climate change driven urbanization in Africa? *Journal of Development Economics 124*, 60–82.

Heyman, D. P. and M. J. Sobel (1984). *Stochastic Optimization*. McGraw-Hill.

Kawabata, A., S. T. Nakamoto, and R. Curtiss (2017). Recommendations for coffee berry borer integrated pest management in Hawaii 2017.

Kurukulasuriya, P., R. Mendelsohn, et al. (2008). Crop switching as a strategy for adapting to climate change. *African Journal of Agricultural and Resource Economics 2*(1), 105–126.

Kyriakidis, E. G. (1993). A markov decision algorithm for optimal pest control through uniform catastrophes. *European Journal of Operational Research 64*(1), 38–44.

Lee, D. J., D. C. Adams, and F. Rossi (2007). Optimal management of a potential invader: the case of zebra mussels in Florida. *Journal of Agricultural and Applied Economics 39*(s1), 69–81.

Lee, D. J., M. Motoki, C. Vanderwoude, S. T. Nakamoto, and P. Leung (2015). Taking the sting out of little fire ant in Hawaii. *Ecological Economics 111*, 100–110.

Lee, T., G. Judge, and T. Takayama (1965). On estimating the transition probabilities of a markov process. *Journal of Farm Economics 47*(3), 742–762.

Livingston, M., M. J. Roberts, and Y. Zhang (2014). Optimal sequential plantings of corn and soybeans under price uncertainty. *American Journal of Agricultural Economics 97*(3), 855–878.

Lobell, D. B., M. J. Roberts, W. Schlenker, N. Braun, B. B. Little, R. M. Rejesus, and G. L. Hammer (2014). Greater sensitivity to drought accompanies maize yield increase in the US midwest. *Science 344*(6183), 516–519.

Lockwood, J. A., W. P. Kemp, and J. A. Onsager (1988). Long-term, large-scale effects of insecticidal control on rangeland grasshopper populations (Orthoptera: Acrididae). *Journal of Economic Entomology 81*(5), 1258–1264.

Marcos, W., M. Ferreira, A. Cortes, and J. Tavares (2013). Markovian decision problems applied in pest control of agriculture production systems. *IFAC Proceedings Volumes 46*(24), 159–164.

Mendelsohn, R., W. D. Nordhaus, and D. Shaw (1994). The impact of global warming on agriculture: A ricardian analysis. *The American Economic Review*, 753–771.

Miao, R., M. Khanna, and H. Huang (2015). Responsiveness of crop yield and acreage to prices and climate. *American Journal of Agricultural Economics 98*(1), 191–211.

Nordhaus, W. D. (2010). The economics of hurricanes and implications of global warming. *Climate Change Economics 1*(01), 1–20.

Onstad, D. W. and R. Rabbinge (1985). Dynamic drogramming and the computation of economic injury levels for crop disease control. *Agricultural Systems 18*(4), 207–226.

Perrings, C. and B. Walker (2004). Conservation in the optimal use of rangelands. *Ecological Economics 49*(2), 119–128.

PRISM Climate Group (2018). Historical past (1895-1980) and recent years (1981-2017). `http://prism.oregonstate.edu`. Accessed: 01-01-2018.

Roberts, M. J. and W. Schlenker (2013). Identifying supply and demand elasticities of agricultural commodities: Implications for the US ethanol mandate. *American Economic Review 103*(6), 2265–95.

Schlenker, W., W. M. Hanemann, and A. C. Fisher (2005). Will US agriculture really benefit from global warming? accounting for irrigation in the hedonic approach. *The American Economic Review 95*(1), 395–406.

Schlenker, W., W. M. Hanemann, and A. C. Fisher (2006). The impact of global warming on US agriculture: An econometric analysis of optimal growing conditions. *Review of Economics and Statistics 88*(1), 113–125.

Schlenker, W. and M. J. Roberts (2009). Nonlinear temperature effects indicate severe damages to us crop yields under climate change. *Proceedings of the National Academy of Sciences 106*(37), 15594–15598.

Seo, S. N. and R. Mendelsohn (2008). An analysis of crop choice: Adapting to climate change in South American farms. *Ecological Economics 67*(1), 109–116.

Silberberg, E. (1971). *The Le Chatelier Principle as a Corollary to a Generalized Envelope Theorem.* Institute for Economic Research.

Spedicato, G. A. (2017). Discrete time markov chains with R. *The R Journal*. R package version 0.6.9.7.

Sumner, D. A. (2009). Recent commodity price movements in historical perspective. *American Journal of Agricultural Economics 91*(5), 1250–1256.

USDA (2016). Hawaii coffee marketings: Preliminary season estimates.

White, D. J. (1993). A survey of applications of markov decision processes. *Journal of the Operational Research Society 44*(11), 1073–1096.

Woodill, A. J., D. Hemachandra, S. T. Nakamoto, and P. Leung (2014). The economics of coffee production in Hawaii. *Economic Issues*.

Woodill, A. J., S. T. Nakamoto, A. M. Kawabata, and P. Leung (2017). To spray or not to spray: A decision analysis of coffee berry borer in Hawaii. *Insects 8*(4), 116.

Wooldridge, J. M. (2010). *Econometric Analysis of Cross Section and Panel Data.* MIT press.