Three essays on applied economics

Minkyong Ko

Dissertation submitted to the Faculty of the
Virginia Polytechnic Institute and State University
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy
in
Economics, Agriculture and Life Science

Chi Ta, Co-chair
Ford Ramsey, Co-chair
Klaus Moeltner
Jesse Tack

December 1, 2023
Blacksburg, Virginia

Keywords: climate change, production, outfarm migration, energy conservation, behavioral economics

Copyright 2023, Minkyong Ko

Three essays on applied economics

Minkyong Ko

(ABSTRACT)

This dissertation is centered on applying 'nonlinearity' across various fields, a decision informed by the understanding that linear models frequently fail to capture the full extent of real-world complexities. This approach is driven by the distinct insights that nonlinearity offers, insights crucial for a more profound and precise understanding of diverse phenomena. In this regard, I have explored a variety of empirical methodologies and theoretical frameworks, each chosen for its effectiveness in unraveling and accurately depicting the complexity inherent in different subjects. The first paper, "Warming Temperatures and Potential Adaptation through Breeding: Evidence from U.S. Soft Winter Wheat," examines the impact of climate change on wheat production. Given wheat's role as a major staple for much of the global population, its susceptibility to rising temperatures presents significant challenges to food security. Despite its importance, comprehensive production data for wheat considering extensive U.S. regions is relatively scarce. To address this, I compiled a unique dataset on wheat production comprising 35,000 observations over 50 years from USDA-ARS hard-copy publications. Findings obtained through the mixed-effects model reveal significant variations in the influence of nonlinearly specified temperature on wheat production. This analysis identifies a decline not only in wheat yield across various U.S. farming sites but also in wheat quality, an aspect often neglected in similar studies. The issue is further compounded by simulations I conducted, which predict worrying decreases in both yield and quality due to rising temperatures. Despite these challenges, my analysis of varietal improvements indicates modest yet significant progress in countering the effects of warming, offering

viable strategies for agricultural adaptation. In the second chapter, "Semiparametric Analysis of Out-farm Migration in China," I explore the nonlinear relationship between sectoral migration and the income gap within China. This work builds on my co-authored publication, "Intersectoral Labor Migration and Agriculture in the United States and Japan," published in Agricultural Economics. While the earlier study employed discrete thresholds and kink approaches to explore migration patterns in developed countries such as Japan and the United States, it did not reveal significant nonlinear relationships. This led me to investigate whether the results were influenced by the economic development status of the countries in question. Focusing on China, a developing country with distinct labor dynamics, I employ semiparametric methodologies to assess migration patterns, diverging from the linear assumptions common in existing literature. By using nationally representative data, it suggests a potential nonlinear relationship between farmers' sectoral migration and the income gap, providing new insights into labor migration in developing contexts. The last chapter, "Enhanced Salience of Nonlinear Pricing and Energy Conservation," explores the energy consumption of residents of Hanoi in Vietnam, using a large-scale randomized control trial. I study whether enhancing salience of information with respect to the nonlinear pricing can help energy conservation. The novelty of the project lies in its experimental design and the utilization of digital tools such as smart meters and mobile apps, adopting technologies with the potential to alter consumer behavior. Currently, we are in the post-intervention data collection phase. Supported by the International Growth Centre (IGC), the project aims to bridge the research gap in energy consumption behavior in developing countries, thereby contributing to policymaking in energy management and development in these regions. Through these diverse yet interconnected chapters, I attempt to use the varied applications of nonlinearity in studying economic and environmental issues. The main objective is to contribute to both academic knowledge and practical policymaking in these fields, addressing complexities that are often oversimplified. This approach aims to provide

a more comprehensive understanding of the intricate dynamics in these areas.

Three essays on applied economics

Minkyong Ko

(GENERAL AUDIENCE ABSTRACT)

This dissertation focuses on the application of 'nonlinearity' across various fields, acknowledging that linear models often fail to fully capture the complexities of real-world scenarios. This approach yields essential insights for a more precise understanding of diverse phenomena, incorporating a range of empirical methods and theoretical models, each selected for their efficacy in accurately depicting complexities in various subjects. The first part of my dissertation examines the agricultural sector, specifically the impact of weather on crop yields. Instead of traditional methods that use basic temperature data like minimum, maximum, or average values, I use a nonlinear approach that aligns with the specific growth stages of crops. In the area of labor economics, the dissertation explores migration patterns in China. It questions the traditional linear relationship between migration and income disparities, suggesting a more complex model that better represents the dynamics of a developing economy. The final chapter addresses the field of energy policy, examining consumer responses to nonlinear electricity pricing models in practice. This section explores the challenges faced by individuals in understanding such policies and assesses the impact of providing real-time information about nonlinear pricing on promoting energy conservation.

Dedication

To my mother and my husband, Dongwoo: I dedicate this dissertation to you both for your enduring support and the joy you have shown, which surpasses even my own.

Acknowledgments

I would like to thank my esteemed advisors, Austin Ford Ramsey and Chi Ta. Your support and encouragement were crucial in my pursuit and completion of this journey, and in starting a new one. Your guidance provided me with the confidence to persevere in my PhD, especially during moments of doubt. I am truly grateful for your patience and generosity throughout this course. You are my mentors who will continue to impact me both professionally and personally. I also want to express my gratitude to my committee members, Klaus Moeltner and Jesse Tack, for their assistance, sincere guidance, and detailed feedback. Having you as part of my academic journey has been a privilege. I wish to thank my husband for his support throughout this journey. These past years have been some of our most challenging, yet his understanding and presence by my side have helped me to walk this path. Lastly, I am thankful to God for blessing me with such precious times and people.

Contents

List of Figures xi

List of Tables xiii

1 Warming Temperatures and Potential Adaptation through Breeding: Evidence from U.S. Soft Winter Wheat 1

1.1 Abstract . . . 1

1.2 Introduction . . . 2

1.3 Literature Review . . . 5

1.4 Background of US wheat and its breeding system . . . 8

1.5 Data . . . 9

1.5.1 Wheat Data . . . 9

1.5.2 Weather Data . . . 11

1.6 Empirical Approaches . . . 12

1.6.1 Basic Model . . . 14

1.6.2 Predicting Warming Scenario . . . 16

1.6.3 Evaluating Breeding Efforts: An Extended Model Approach . . . 17

1.6.4 Robustness Checks . . . 17

1.7 Results . . . 18

1.8 Conclusion . . . 25

1.9 Appendix . . . 27

2 Semiparametric Analysis of Out-farm Migration in China **34**

2.1 Abstract . . . 34

2.2 Introduction . . . 35

2.3 Literature Reviews . . . 36

2.4 Background of Chinese Inter-sectoral Migration . . . 40

2.4.1 Overview of Chinese Out-farm Migration . . . 40

2.4.2 Chinese Unique Land System, Hukou . . . 44

2.5 Theoretical Framework . . . 45

2.6 Model . . . 48

2.6.1 Data . . . 48

2.6.2 Parametric Estimation and Results . . . 51

2.6.3 Semiparametric Partially Linear Regression Model . . . 54

2.7 Results and Discussion . . . 57

2.7.1 Semiparametric Results . . . 57

2.7.2 Specification Test . . . 60

2.7.3 Robustness Check: Controlling for Heteroskedasticity . . . 62

2.7.4 Remarks . . . 64
2.8 Conclusion . . . 65

3 Enhanced Salience of Nonlinear Pricing and Energy Conservation 68

3.1 Abstract . . . 68
3.2 Introduction . . . 69
3.3 Conceptual Model . . . 73
3.4 Research Setting . . . 75
3.4.1 Background of Vietnam's Energy Use . . . 75
3.4.2 Sampling and Randomization Process . . . 77
3.5 Experimental Design . . . 80
3.6 Data . . . 83
3.6.1 Variables . . . 83
3.6.2 Balance Check . . . 85
3.7 Empirical Analysis . . . 91
3.7.1 Treatment Effects . . . 91
3.7.2 Heterogeneous Treatment Effects . . . 99
3.8 Conclusion . . . 105

Bibliography 107

List of Figures

1.1 Wheat Trial Sites 13

1.2 Warming Scenarios 22

1.3 Resilience and Introduction Year 24

1.4 Resilience and Mean Effect 24

1.5 Quantity and Quality by Year 31

1.6 Marginal Effects across Three Seasons 32

1.7 Warming Scenario across Three Seasons 33

2.1 Labor Share, Productivity and Migration Rates 42

2.2 Semiparametric Results 58

2.3 Semiparametric Results after controlling Heteroskedasticity 63

3.1 Distribution of Monthly Electricity Use among App Users 78

3.2 Mean Comparisons between the rest of the sample (370K) and 30K 80

3.3 App Interface Variations: Control Group vs. Treatment Groups 83

3.4 Mean Comparisons across control and treatment groups 88

3.5 Time Effect by Week (study period) 98

3.6 Heterogeneous Effect: Billing Cycle (study period only) 103

3.7 Heterogeneous Effect: Billing Cycle (all periods) 104

List of Tables

1.1 Descriptive Statistics 12

1.2 Weather Effect on Wheat Yields 20

1.3 Weather Effect on Wheat Quality 21

1.4 Seasonal Weather Effect on Wheat Yields 27

1.5 Seasonal Weather Effect on Wheat Quality 28

1.6 Robustness Check for Yield 29

1.7 Robustness Check for Quality 30

2.1 Summary Statistics and Variable Description 50

2.2 OLS Estimation 53

2.3 Estimates from Quadratic and Cubic Models 54

2.4 Estimates from the Semiparametric Model 58

2.5 Controlling Heteroskedasticity 63

3.1 Nonlinear Pricing Structure Implemented by EVN Hanoi, Vietnam 77

3.2 Kolmogorov-Smirnov Tests between the rest of the sample (370K) and 30K . 81

3.3 Intervention Strategies by Groups 84

3.4 Balance analysis: mean comparisons across control and treatment groups . . 87

3.5 Kolmogorov-Smirnov Tests across control and treatment groups 89

3.6 Distribution of Consumption in Tiers 90

3.7 Treatment effects: Cross-Sectional variation 93

3.8 Treatment effects: panel 94

3.9 Heterogenous Effects: Billing Start Timing 102

Chapter 1

Warming Temperatures and Potential Adaptation through Breeding: Evidence from U.S. Soft Winter Wheat

1.1 Abstract

Using annual wheat trial data, this study examines the impact of climate change on crop quality and quantity, as well as the possibility of mitigating these effects through breeding efforts. While most contemporary research focuses on the relationship between grain yields and warming, we consider quality as a critical component of crop revenue that must be stressed in climate studies since it can be associated with reduced protein content or size. Due to the importance of crop quality for nutrients, farmers' livelihoods, and consumers' health, studies of the effects of temperature on agriculture disregarding quality risk underestimating or overestimating the overall impact of warming (Ramsey et al., 2021). As a result, in this paper, we quantify quality by measuring wheat test weights as well as yield. Data from U.S. Uniform Southern Soft Winter Wheat Nursery Reports (1962-2017) and PRISM weather data are used, involving over 17,000 observations from 270 cultivars and 80 locations. Our

findings indicate that elevated temperature conditions adversely affect both the yield and quality of wheat and these losses increase as temperatures rise. Additionally, recognizing the significance of U.S. wheat as a major staple crop essential for ensuring future food security, our study assesses the advancements in breeding programs over time, with a specific focus on heat resilience. Our research reveals that recently developed wheat cultivars exhibit improvements compared to their predecessors. This finding holds important implications for comprehending ongoing wheat breeding efforts in response to rising temperatures.

1.2 Introduction

As a principal staple food for 35% of the world's population, the demand for wheat continues to rise as the global population is projected to surge from its current level of 7.20 billion to 9.55 billion by 2050 (Curtis and Halford, 2014). Such a scenario could make food security difficult if demand levels cannot be met in the next few decades. The Intergovernmental Panel on Climate Change (IPCC, 2021) has reported that the past four decades have been warmer than any decade since 1850. Three major cereal crops—wheat, rice, and maize—will decline in most farming areas due to the rise in temperatures which is projected to increase by $2.4 - 6.4°C$ within this century alone (IPCC, 2014). Numerous studies support these findings by demonstrating the negative relationship between extreme temperatures and crop productivity (Lobell and Asner, 2003, Lobell and Field, 2007, Schlenker and Roberts, 2009, Tack et al., 2017). If climate change becomes a significant challenge in sustaining human life and the environment, it will be imperative to find effective adaptation and mitigation strategies.

In the context of agriculture, adaptation encompasses modifications to farming practices, management approaches, and strategies to reduce vulnerabilities and enhance the resilience

of crop production to climate-related challenges. For instance, producers might adjust planting or harvesting schedules or experiment with different crop locations. Governments could employ policy interventions, and international trade can also serve as an adaptation strategy to warming (Burke and Lobell, 2010, Ortiz-Bobea and Just, 2013, Vital et al., 2022). When these agents proactively adapt their methods to address weather changes, it can lead to outcomes that significantly differ from those where no adaptations are made. Hence, incorporating these adaptations is crucial in accurately estimating the effects of weather on agricultural practices.

Among this wide range of actions, one strategy is technological change through plant breeding and improved management practices, which has enhanced farm productivity and increased agricultural output. There is a possibility that gains from improved management and new crop varieties could cancel out lower yields from warming. Previous research has identified the important relationship between yield risk, shifts in production technologies, and economic outcomes (Burke and Emerick, 2016, Shi et al., 2013a). This study primarily addresses how breeding efforts on wheat varieties can serve as an adaptation strategy against rising temperatures.

While many studies primarily focus on crop yield, it is important to recognize that the quantity side is not the sole economic factor of significance. Crop quality is also a crucial factor that should be stressed in breeding research as well as climate studies. Buyers purchase crops based on quality, and the price is determined accordingly. This price is closely related to farmers' livelihoods. Quality is also important to processors to maintain the desired end-product standards since the majority of wheat is not bought in its raw and unprocessed form but rather as a raw material for further processing. Consequently, studies that overlook the dimension of quality may risk underestimating or overestimating the overall impact of warming. Lyman et al. (2013) and Ramsey et al. (2021) reveal that when studies fail to

account for changes in quality, it leads to an underestimation of the impacts of warming on outcomes, and the negative effect on market value would be even worse. On the contrary, Kawasaki and Uchida (2016) show that quantity and quality effects may operate in different directions. In any case, taking account of quality is critical for avoiding misleading economic results.

In this article, we quantify quality by measuring wheat test weight. Test weight is a density measurement that is used as an indicator of grain quality as it provides insights into kernel density and potential milling yield; The unit is pound per bushel (lbs/bushel), whereas grain yield is expressed in bushels per acre (NDSU Agriculture Communication, 2012). Hence, a higher test weight generally suggests that the kernels are well-filled. Note that test weight and yield in wheat are two distinct characteristics, and they do not necessarily correlate with each other. While yield refers to the quantity of wheat produced per unit of land area, test weight denotes the weight of a bushel of wheat grains. In other words, it is possible to have a high-yielding wheat crop with good quantity but lower test weight, indicating that the grains are lighter or less dense. Conversely, a lower-yielding grain might have a heavier test weight. The relationship between these two factors is not straightforward and can be influenced by various agronomic practices, environmental conditions, and wheat varieties. Farmers and grain producers consider both yield and test weight to determine the overall suitability and market value of their wheat crops along with other quality factors like protein content and moisture levels. According to the USDA official standard for wheat, there are five grades depending on test weights.[1] In other words, the primary grade of wheat is determined by the test weight as it indicates how much flour can be extracted. Although we acknowledge the influence of temperature on overall wheat grade can be heterogeneous and

[1]In case of soft wheat case, 1st grade is 60 lbs, 2nd grade is 58 lbs, 3rd grade is 56 lbs, 4th grade is 54 lbs and 5th grade is under 51 lbs. Other than test weights, the official grades include foreign materials and damaged or shrunken kernels.

the test weight alone may not capture wheat quality, it is a widely accepted measure as it serves as a fundamental measure of wheat quality and influences pricing and marketability in the agricultural and grain trading industries.[2]

In short, this paper has three objectives: a) to investigate the nonlinear relationship between temperature and yield and quality, respectively; b) to simulate the impact of a 1°C temperature increase on wheat yield and test weight; and, c) to examine whether newer varieties of wheat are more or less susceptible to extreme temperatures. Results have important implications for understanding how breeding efforts on wheat have mitigated the effects of climate change and to what degree new varieties are resilient to long-term changes in temperature. We use data from U.S. Uniform Southern Soft Winter Wheat Nursery Reports between 1962 and 2017. These reports contain the results of variety trials by public institutions throughout the southern United States. The data contain over 17,000 year-location-variety observations across over 271 cultivars and 82 locations. Data from the trials are combined with weather data available from PRISM over a 2.5 by 2.5-mile grid of the contiguous United States. We measure variability in yield and test weight using a fixed effects model.

1.3 Literature Review

Many studies have examined the relationship between weather and crop productivity, as climate change has become a serious challenge in sustaining human life and the environment (Lobell and Asner, 2003, Lobell and Field, 2007, Schlenker and Roberts, 2009, Tack et al., 2015, 2017). Aligned with a report from IPCC (2014), for instance, Schlenker and Roberts (2009) evaluate the relationship between weather and yields for three major crops in the

[2]Additional quality tests include moisture content, protein content, falling number, impurities, etc., but they are not part of the US grain standards. Also, farmers often receive payment based on test weight, with other aforementioned factors typically not directly influencing their compensation. However, the actual payment structure can differ based on local agricultural practices and market demands.

US: corn, soybeans, and cotton. They predict yield increases gradually up to a certain temperature level but eventually decreases sharply for all three crops. Tack et al. (2017) show that sorghum yields start to decline above 33°C. Lobell and Asner (2003) use corn and soybean data and found that yields decreased by 17% with an increase of 1°C. Lobell and Field (2007) found that prolonged maximum temperatures reduced rice yields. As a result, there is a possibility that the lower yields due to warming temperatures cancel out a substantial portion of the gains from management skills and technological improvements (Lobell et al., 2011). Even though the methodologies, study regions, periods, and crops vary to some extent, the common pattern of the literature is to simulate the impacts of weather in the future, which shows poor prospects in most cases.

However, climate change is not only expected to affect crop quantity; it is also expected to influence quality. Relatively fewer, but a growing number of studies have analyzed the impact of warming on crop quality (Kawasaki and Uchida, 2016, Lyman et al., 2013, Ramsey et al., 2021, Wassmann et al., 2009). Some research on Japanese rice cases (Morita et al., 2016) shows that an elevated climate during the growing period dramatically degrades rice quality. As a primary staple of Japan, data on rice in that region are systemically organized and open to the public. Hence, rice grading data are relatively accessible. However, in many other cases where grading data are unavailable, the most challenging part of studying crop quality is how to measure it and find alternative methods since crop quality is often multidimensional and there are few relevant data available (Ramsey et al., 2021). As a means of measuring the crop quality other than direct grade data, Lyman et al. (2013) calculate the portion of chalkiness and the broken kernel to determine the rice milling quality in Arkansas, while Ramsey et al. (2021) use kernel size data to represent peanut quality in Virginia, North Carolina, and South Carolina.

It should be noted that the two studies above indicate that a failure to account for quality

changes in an analysis may result in an underestimate of the impact of warming on outcomes and the resulting negative impact on market value may be even more severe than previously thought. On the other hand, Kawasaki and Uchida (2016) have demonstrated that, while the quantity effect increases as the temperature rises, the quality effect reverses in the case of rice. It implies that the directions and magnitudes of quantity and quality may not always work in the same way, so the results should be explicitly compared to avoid misleading inference. In this regard, this research assesses the intertwined dynamics of quantity and quality in response to temperature variations within wheat crops to offer a more accurate understanding of how these factors collectively influence wheat production.

Another line of research highlights effective adaptation and mitigation strategies to address warming effects on crop production (Peng et al., 2004, Schlenker and Roberts, 2009). Adaptation involves adjusting and modifying strategies to suit a new environment. While immediate weather shocks may be unpredictable and challenging to manage, agents or firms can anticipate and respond to long-term climate shifts. Agricultural adaptation to climate change is a broad term that covers various actions and solutions aimed at addressing the challenges posed by changing weather conditions. The research encompasses agronomic practices like crop rotation, inter-cropping, cover cropping, and adjustments to planting or harvesting dates, as well as policy, institutional analysis, and international trade studies (Kawasaki and Uchida, 2016, Ortiz-Bobea and Just, 2013, Vital et al., 2022). Within this extensive field of research, some literature focuses on examining the breeding efforts and varietal improvements, whether the gains from improving management and switching to new crop varieties may be able to compensate for the lower yields or quality caused by the warming of the climate. The studies on varietal improvement are well-established and focus on the attribution of changes in yields to different sources including plant breeding advancements (Nalley et al., 2008, Nolan and Santos, 2012, Wang et al., 2022). In particular, Voon and Edwards (1992)

measure significant research benefits from quality improvements in wheat and Nogueira et al. (2015) find significant welfare impacts of yield improvement within different classes of wheat.

Given previous literature, this paper assesses the response of wheat to warming temperatures and also measures whether breeding efforts may mitigate temperature effects. To the best of our knowledge, there is little literature approaching this topic using US wheat data covering a broader geographical region and a longer period, despite the pivotal role of US wheat in the agricultural landscape.

1.4 Background of US wheat and its breeding system

The significance of US wheat extends both domestically and globally. It ranks among the top three field crops in the United States, with substantial acreage, production, and farm receipts, following corn and soybeans. Internationally, the US has contributed approximately 6-7 percent of the world's total wheat production, firmly establishing itself as one of the world's leading wheat exporters. This dual role highlights the importance of maintaining a stable supply of high-quality wheat within the US, not only to ensure national food security but also to address global food challenges exacerbated by the effects of climate change.

Wheat is adaptable to various climates across most US states and demands less irrigation compared to some crops. However, its extended growth cycle makes it susceptible to weather variations and pests. Particularly, winter wheat, the focus of our study, is planted in the fall, lies dormant during winter, and resumes growth for summer harvest. As a result, measures are necessary to strengthen its resilience against these challenges.

Kusunose et al. (2023) document several challenges faced by the US wheat breeding system. Economically, the market price for wheat in the US is often lower than that of crops such

as corn and soybeans. When combined with production costs, this results in narrow profit margins for wheat. Consequently, the majority of wheat research is undertaken by academic institutions, predominantly land grant universities. In contrast, the private sector's engagement in wheat breeding is noticeably limited.

From a market perspective, US wheat is segmented into various classifications based on factors like its growth season and end-use. Examples of these classifications include Soft Winter Wheat, Hard Spring Wheat, and Durham. Each category has its unique set of breeding objectives and marketing tactics. The complexity of this classification system, together with the variable market sizes and demands, makes it challenging for research investment decisions.

While genetic modification(GM) technology has advanced for many crops, its adoption for wheat has been slow. Wheat's role as a primary food source for human beings makes its genetic modification a topic of caution. Due to these concerns, the industry still largely depends on traditional breeding techniques, especially when trying to improve aspects like how well wheat can handle changing climate conditions and how much yield it can produce.

Considering the challenges mentioned, it is important to analyze how the breeding efforts in the US wheat have been adapted to address the effects of global warming. Our study aims to examine this adaptation in detail through varietal improvements over time.

1.5 Data

1.5.1 Wheat Data

The wheat data used in this study came from the U.S. Uniform Southern Soft Winter Wheat Nursery Reports between 1962 and 2017. These reports provide the results of variety trials

that have been conducted mainly by public institutions throughout the southern United States. The data contain over 17,000 year-location-variety observations across over 280 cultivars and 82 locations. Our measurements include both quantity (yield) and quality (test weight). All the wheat data is specific to dryland soft red winter wheat, which is grown without irrigation and relies solely on rainfall. The data form an unbalanced panel as locations and varieties enter and exit the nursery reports over time.

The unbalanced nature of the panel is particularly apparent in the variety dimension. The nursery reports were developed through cooperative research conducted by the USDA Agricultural Research Service and State Agricultural Experiment Stations on new wheat varieties. Each cultivar is entered into the trial and then grown at all of the cooperating sites. As most cultivars in these trials are not commercially available and are being grown for research purposes, they are generally not kept in the trials for long periods. Rather, new breeding lines and varieties are entered each year to develop new varieties and conduct experiments.

The primary objective of conducting wheat trials on university-granted lands or experimental sites is to evaluate the performance of novel wheat varieties or breeding lines in terms of yield, disease resistance, and quality across diverse environmental conditions. The top-performing varieties are released for commercial production, with seeds distributed to farmers either directly or through intermediaries such as government extension services or private seed companies. However, the exact point at which a wheat variety is commercialized is complex to determine because it is subject to various regional regulations and involves multiple stakeholders, such as government agencies, private seed companies, and farmers.[3]

Check varieties, which serve as standard or reference strains, are incorporated annually into trials to compare and evaluate other new varieties entering the trials. These benchmark va-

[3]Factors such as intellectual property rights, licensing agreements, and other legal and regulatory considerations can also affect the commercialization process.

rieties typically remain in trials for multiple years, providing a consistent basis for evaluation across different time periods. For instance, AGS2000, a variety developed at the University of Georgia appeared as a check variety from 1999 to 2016. AGS2000 is the most common cultivar in the trials with a total of 434 location-year observations in our data set.

1.5.2 Weather Data

The wheat trial data is combined with weather data obtained from PRISM over a 2.5 by 2.5-mile grid throughout the contiguous United States. Since the trial data only contains approximate locations of the growing sites, the center of the town where the trials took place has been geocoded and matched to the corresponding PRISM grid cell.[4] In terms of weather data construction, we employed growing degree days as our primary metric. The use of growing degree days in comparison to average temperature or minimum and maximum temperature appears to be more advantageous in terms of having a smoothing effect on measured weather variations (Tack et al., 2015).[5] Hence, the daily minimum and maximum temperatures in the PRISM data are interpolated using a sinusoidal curve and then used to calculate the number of growing degree days across the entire wheat growing season, spanning from October to next year June, as in Schlenker and Roberts (2009).[6] Total precipitation is calculated by adding up the daily precipitation reported in the PRISM data throughout the growing season.

Table 1.1 presents the descriptive statistics of our dataset, while Figure 1.1 provides a graphical representation of the geographical locations of our trial sites. Our dataset encompasses

[4] Miller et al. (2021) point out that potential measurement errors due to imprecise location information can lead to underestimations or overestimations in assessing climate impacts.

[5] When relying on daily minimum and maximum temperatures or average temperatures, short-term fluctuations can significantly affect the data.

[6] Growing degree days are a heat-based metric to quantify the accumulated warmth necessary for the growth and development of crops. While there are several approaches to calculating degree days, our study employs a method of interpolating daily minimum and maximum temperatures using a sinusoidal curve.

the majority of the cultivation regions for soft red winter wheat.

Table 1.1: Descriptive Statistics

Variable	Mean	S.D.	Min	Max
yield (bushel/acre)	57.54	23.59	0.10	168.50
test weight (lbs/bushel)	56.21	3.90	3.60	68.60
Number of sites by year	21.56	4.95	11.00	34.00
Number of years by site	13.15	11.39	2.00	49.00
Number of varieties by year	15.24	5.39	5.00	27.00
Number of years by variety	2.82	2.08	2.00	18.00
low (degree days)	2435.19	305.08	1252.75	2916.79
med (degree days)	862.10	284.09	140.55	1491.14
high (degree days)	275.06	171.45	2.87	853.74
precipitation (cumulative)	939.56	276.04	103.24	1832.81

Note: We have 82 trial sites and 270 distinct varieties over 50 years (1962-2017). *low* is defined as degree days above 0 °C minus the degree days above the lower threshold. Similarly, *med* represents the degree days above the lower threshold subtracted from the degree days above the upper threshold. Lastly, *high* is defined as the degree days exceeding the upper threshold. The specific values for these thresholds vary according to the three distinct growing seasons. The lower threshold is 10 °C (Fall), 5 °C (Witner), and 18 °C (Spring), and the upper threshold is 18 °C (Fall), 10 °C (Witner), and 30 °C (Spring), respectively.

1.6 Empirical Approaches

Before moving on to the model part, there are three methodological issues to be addressed in this research. A primary concern is the selection of samples and the rate of attrition. As previously mentioned, cultivars change frequently from year to year. To a lesser degree, the locations of the trials varied across time, making it challenging to obtain a balanced panel. If sample selection or attrition occurs, three methods may be used: *i*) disregarding these

Figure 1.1: Wheat Trial Sites

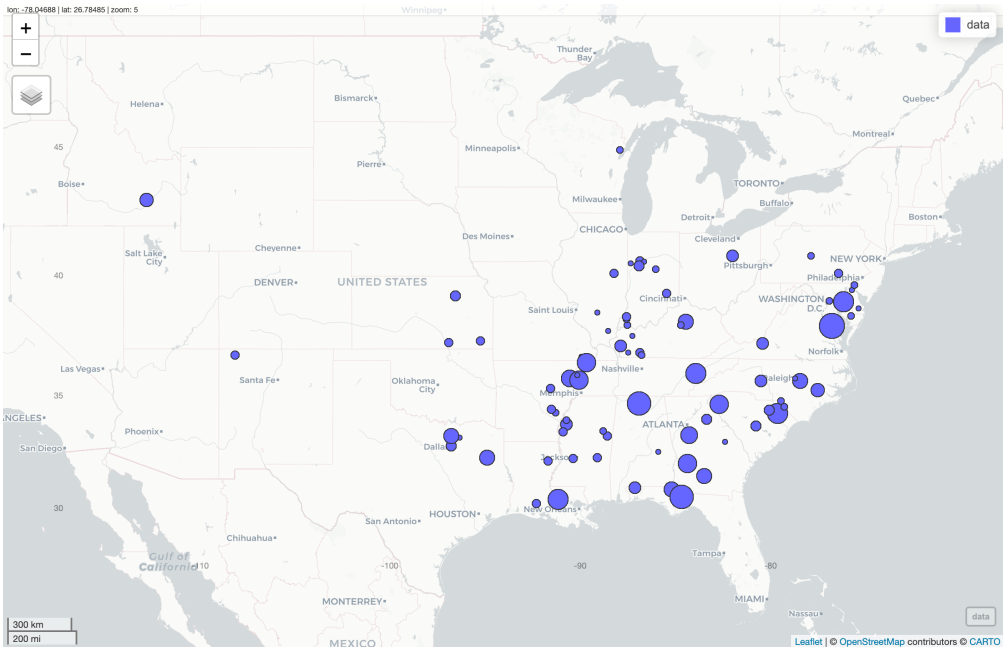

issues and using an unbalanced panel, *ii*) removing the observations that have the issues and making a balanced subset, and *iii*) constructing an econometric model that takes these issues into account (Wooldridge, 2015). This paper adopts the first option, but we must add an assumption that weather variables are given randomly and exogenously. Using this assumption, a consistent estimator can be obtained and unobserved heterogeneity can be controlled by fixed effects terms. Additionally, we include only those locations and cultivars that have been observed at least twice as in Tack et al. (2017) and Shew et al. (2020).

It is also important to consider other factors than climate change that are responsible for positive trends in wheat production over time, including technological advances (Shi et al., 2013b). Neglecting these factors would result in bias in determining the exact contribution of weather variables to yield and quality. Therefore, our model includes a time trends variable with year and year squared to capture technical improvements and other time-dependent

trends instead of time-fixed effects following Schlenker and Roberts (2009).[7]

Lastly, Lobell et al. (2008) emphasize the importance of considering human factors across locations, including management skills and adaptations, which may have a significant impact on crop growth. For example, adjusting crop planting and harvesting dates to align with local climate conditions, as well as customizing fertilizer based on soil types and conditions, can optimize crop productivity. Such location-specific human activities are crucial considerations in agricultural research. As stated above, however, as the wheat varieties in our data are being grown for experimental purposes, they are likely to be protected from pests and deficiencies in water and nutrients. For this reason, we assumed that all trial sites would be managed with best management practices.

1.6.1 Basic Model

We regress log wheat yield on temperature, precipitation, location, and cultivar fixed effects. Specifically, we consider the following multiple regression model with fixed effects to estimate the weather effect on wheat quantity:

$$lny_{ijt} = \gamma_i + \tau_j + \alpha_1 t + \alpha_2 t^2 + f(X_{jt}; \beta) + \epsilon_{ijt}, \tag{1.1}$$

where lny_{ijt} denotes either the log of wheat yield or test weight for variety i at trial location j in year t; γ_i is a vector of cultivar fixed effects; τ_j denotes a vector of location fixed effects; $\alpha_1 t + \alpha_2 t^2$ captures the time trends, and $f(\cdot)$ represents the weather function, while X_{ijt} is a vector of weather-related variables and the β are the slope parameters for X_{jt}; ϵ_{ijt} is the

[7] Although our chosen form assumes constant coefficients, it efficiently captures long-term trends. Fixed effects consume degrees of freedom, potentially reducing parameter estimate precision and obscuring the impacts of technological improvements we seek to isolate. In this regard, more standard time-fixed effects are considered in an alternative specification.

error term.

We add intercepts γ_i to control for genetic differences in yield and τ_j to control for spatially invariant unobserved effects such as soil quality. To avoid omitted variable bias, these fixed effects are included to capture additive time invariant influences. The time trend is included in this model as a quadratic form to capture the genetic improvements at a diminishing rate over time.[8] Since breeding programs often focus on optimizing elite wheat lines, this property needs to be reflected. Next, we cluster standard errors by location, which takes account of correlation in the error terms within the same location over time.

This paper presumes a nonlinear relationship between temperature and quantity or quality rather than linearity. In other words, we assume that when the climate is below or above a certain threshold, two dependent variables - quantity and quality - may decline dramatically. From this point of view, we use the function for capturing weather effects as follows:

$$f(X_{jt};\beta) = \sum_{s=1}^{3} \left(\beta_1 \cdot low_{jt} + \beta_2 \cdot med_{jt} + \beta_3 \cdot high_{jt} + \beta_4 \cdot P_{jt} + \beta_5 \cdot P_{jt}^2\right), \tag{1.2}$$

where low_{jt} represents the degree days from 0 to a season-specific lower threshold, while med_{jt} covers the degree days between these lower and upper thresholds. On the other hand, $high_{jt}$ measures the degree days that exceed the upper threshold. Furthermore, the terms P_{jt} and P_{jt}^2 stand for the quadratic effects of precipitation during the wheat's growing season, affecting both the yield and test weight.

We define the wheat growing season as the period from September 15 to June 15. [9]This

[8]Note that the time trend can also indirectly reflect the influence of other factors affecting wheat production over time. These include technological changes, shifts in consumer preferences, market prices and conditions, trade agreements, and federal and state policies supporting wheat farming.

[9]The defined wheat growing season is a general timeframe that captures the key phases of wheat growth, from planting to harvest. However, we acknowledge that the specific start and end dates of the wheat growing season can vary across different regions and years.

season is broken down into three distinct stages to capture the effects of seasonal weather patterns on wheat: Fall (from September 15 to November 30), Winter (from December 1 to February 28), and Spring (from March 1 to June 15). Each stage reflects the varying impact of temperature and precipitation during its specific time frame. The suggested thresholds for degree days are referred to by Tack et al. (2015). Specifically, during Fall, the temperature ranges from a low of 10°C to a high of 18°C. The Winter season experiences a chillier bracket, from 5°C as its low to 10°C as its high. Meanwhile, Spring spans from 10°C to 30°C.

1.6.2 Predicting Warming Scenario

To predict the potential effects of climate change on wheat yield and quality, we also employ a simulation approach by shifting average weather patterns. This is to contrast the effects of projected weather averages, W_1, against known historical averages, denoted as W_0.

Precisely, the point of reference is the baseline climate average, W_0, which represents the average weather conditions observed during the study period. We then introduce a set of hypothetical newer weather averages, W_1. To model potential outcomes under various warming scenarios, we manipulate observed daily maximum and minimum temperatures of W_0 by systemically adding between 1°C and 5°C. Following these temperature modifications to create W_1, associated degree days are recalculated for each scenario.

Finally, the impact on yield or quality due to this shift is calculated via the following formula:

$$\text{Expected impact} = 100 \times \left[e^{(W_1 - W_0)\cdot\beta} - 1\right] \tag{1.3}$$

This equation represents the percentage change in yield or quality relative to our baseline climate.

1.6.3 Evaluating Breeding Efforts: An Extended Model Approach

To evaluate the climate-mitigating effect of the breeding efforts, especially the heat resilience of various wheat varieties, a hierarchical varying-slope model is employed. We modify the base model to understand how different wheat varieties respond to extremely high temperatures during their growing seasons.

The modified model is given as follows:

$$lny_{ijt} = \alpha + \alpha_1 t + \alpha_2 t^2 + f(X_{jt}; \beta) + u_{0i} + u_{1i} high + u_{2j} + \epsilon_{ijt}, \tag{1.4}$$

where the random intercepts for varieties and locations are captured by u_{0i} and u_{2j}, respectively. They allow for the possibility that some varieties or locations inherently have higher or lower yields or test weights, regardless of the temperature or other factors. $u_{1i} high$ is the term of primary interest. The effect of high temperature is allowed to vary across each variety i through u_{1i}. Hence, each variety may have its unique response to high temperatures. Some varieties might be more resilient, exhibiting minor yield or quality decreases when exposed to heat, while others could be more susceptible, showing significant drops under similar conditions. All other terms are the same with the base model except for α, which represents the overall intercept now.

1.6.4 Robustness Checks

To check for robustness, we apply the following three alternative specifications by *i*)using the dependent variables in level form, *ii*)using a daily average temperature, and, *iii*)maximum and minimum temperatures.

One of the primary reasons for using a logarithmic transformation in regression models is

to ensure the outcome is greater than zero. It also can help address issues related to non-normality, and heteroskedasticity, or capture proportional changes in the dependent variable. However, in the context of our study, especially in the case of test weight, this requires a closer inspection.

Our preliminary analysis shows that test weight is distributed quite evenly, with minimal fluctuations over time. In such a case, a log form, which is typically used to capture proportional changes, may not be sensitive enough to detect small, yet crucial variations in test weight.

Therefore, we change the functional form of the dependent variables to their original level form as a robustness check.

Also, several previous studies obtained meaningful results by using average temperature or max/min temperature (Kelly et al., 2005). Hence, we replace the three temperature variables in our preferred model with measures of average, maximum, and minimum daily temperatures for another way of checking for robustness. The modified functional form of the weather variables is as follows:

$$f(X_{jt};\beta) = \beta_1 tavg_{jt} + \beta_2 tavg_{jt}^2 + \beta_3 P_{jt} + \beta_4 P_{jt}^2, \tag{1.5}$$

and

$$f(X_{jt};\beta) = \beta_1 tmin_{jt} + \beta_2 tmax_{jt} + \beta_3 P_{jt} + \beta_4 P_{jt}^2. \tag{1.6}$$

1.7 Results

Tables 1.2 and 1.3 illustrate the results from the estimating equations from (1.1) to (1.2), assessing the effect of different levels of heat exposure on wheat yield and test weight,

respectively. The columns from (2) to (4) of each table give brief information on alternative specifications. Here we consider five specifications. The first column shows the estimates from our preferred model. In specification (2), we remove time trends and precipitation from the preferred model, so that we can see the response to weather effects directly. Specification (3) adds linear time trends and precipitation. Finally, in the specification (4), linear time trends and quadratic precipitation are included.

The findings from our preferred model indicate that low temperatures positively influence both wheat yield and quality, whereas medium and high temperatures exert a detrimental effect. Specifically, an additional twenty-four-hour exposure to medium or high temperatures during the growing season is linked to a decrease in yield by approximately -0.16% (significant at the 1% level) and -0.05% (significant at the 5% level), respectively. Conversely, cooler temperatures have a favorable impact on wheat yield. This pattern holds for wheat quality as well: medium and high temperatures lead to a decline in quality by 0.001% and 0.002%, respectively, both significant at the 1% level. Low temperatures enhance the wheat test weight. While these results appear numerically small, they show the subtle yet statistically significant trends present in the data.

Contrary to the situation with rice in Japan, our findings indicate a significant positive influence of low temperature exposure on both the quantity and quality of wheat (Kawasaki and Uchida, 2016). However, our findings demonstrate the negative impacts of elevated temperatures align with those of other previous studies(Ramsey et al., 2021, Tack et al., 2017). Additionally, they do not exhibit any strong sensitivity to precipitation and time trend. Note that our preferred model encompasses the entire growing season, from fall to early summer, which is an aggregate of each individual season. Temperature effects on wheat in the preferred model are similar in comparison to the other specifications. However, the effects of weather on wheat production differ significantly across the three seasons. Detailed

Table 1.2: Weather Effect on Wheat Yields

	Model Specifications			
	Preferred	No Trend, No Prec	L Trend, L Prec	L Trend, Q Prec
Intercept	5.0657***	4.9186***	5.1372***	5.0079***
	(0.2718)	(0.2097)	(0.2249)	(0.2319)
Year	-0.0086		-0.0038	-0.0042
	(0.0081)		(0.0030)	(0.0031)
Year Sq.	0.0001			
	(0.0001)			
Low	0.0003***	0.0001	0.0004***	0.0003***
	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Med	-0.0016***	-0.0014***	-0.0016***	-0.0016***
	(0.0001)	(0.0002)	(0.0001)	(0.0001)
High	-0.0005**	-0.0004	-0.0006***	-0.0005**
	(0.0002)	(0.0002)	(0.0002)	(0.0002)
Prec	0.0000		-0.0005***	0.0000
	(0.0001)		(0.0000)	(0.0001)
Prec Sq.	-0.0000***			-0.0000***
	(0.0000)			(0.0000)
R^2	0.4486	0.4310	0.4464	0.4485
Adjusted R^2	0.4367	0.4189	0.4346	0.4367
F-statistic	42.1082	42.5957	55.9233	47.5857

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.3: Weather Effect on Wheat Quality

	Model Specifications			
	Preferred	No Trend, No Prec	L Trend, L Prec	L Trend, Q Prec
Intercept	4.0611***	4.0666***	4.0964***	4.0653***
	(0.0393)	(0.0371)	(0.0336)	(0.0355)
Year	-0.0000		-0.0002	-0.0003
	(0.0014)		(0.0006)	(0.0006)
Year Sq.	0.0000			
	(0.0000)			
Low	0.0001***	0.0000	0.0001***	0.0001***
	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Med	-0.0001***	-0.0001**	-0.0001***	-0.0001***
	(0.0000)	(0.0000)	(0.0000)	(0.0000)
High	-0.0002***	-0.0001***	-0.0002***	-0.0002***
	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Prec	0.0000		-0.0001***	0.0000
	(0.0001)		(0.0000)	(0.0001)
Prec Sq.	-0.0000***			-0.0000***
	(0.0000)			(0.0000)
R^2	0.3208	0.2785	0.3136	0.3207
Adjusted R^2	0.3052	0.2621	0.2980	0.3052
F-statistic	81.2389	13.5704	97.0331	92.5263

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

findings can be found in Tables 1.4, 1.5, and Figure 1.6 in the appendix in Section 1.9.

We observe that weather impacts vary across the growing stages and are distinct for low, medium, and high temperatures within each stage. Given the uncertainty about which effects are predominant, we assess yield and quality impacts for a uniform temperature range of 1°C to 5°C, as illustrated in Figure 1.2.

Figure 1.2: Warming Scenarios

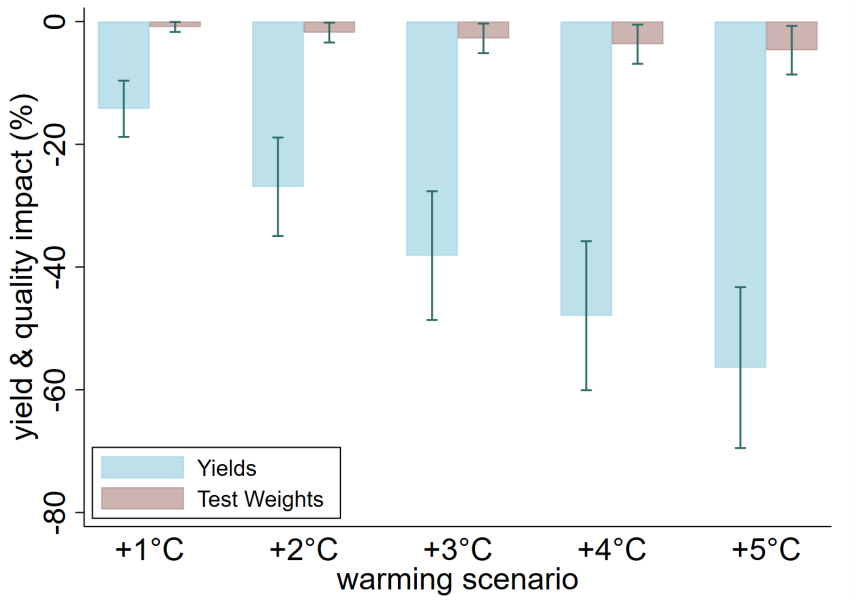

Note: Figure 1.2 illustrates the predicted impacts of warming across the entire growing stages of Fall, Winter, and Spring. The impacts are reported as the percentage change in yield (and quality) relative to the historical climate. Bars represent 95% confidence intervals, with standard errors clustered by location. The predicted negative effects become more severe as the warming scenarios increase.

In our results, increased warming leads to overall declines in both yield and test weight throughout the growing season. The negative effects consistently outweigh the positives. As temperatures rise, both the yield and quality of the crop are projected to decrease even further. Specifically, under a warming scenario of 1°C, we anticipate a decline of approximately 15% in yield. For a more severe 5°C warming scenario, the predicted drop in yield nears

60%. The detrimental effect on test weight remains under 8% across all projected scenarios. However, when we break down the growing season into its three constituent stages, as detailed in the appendix in Section 1.9, these trends show some variation. Notably, a warmer climate might, in certain contexts, have a marginally beneficial effect on wheat production, though this is not statistically significant.

Equation (1.4) examines the interaction between distinct wheat varieties and high temperatures during their growing period. This modified hierarchical model allows for the potential variability in how different wheat varieties react to high temperatures. Using the estimates, we find varieties that were first introduced in more recent years appear to have a better capacity to maintain or even increase their yield despite increased temperature conditions, as depicted by Figure 1.3. The newly entered variety has improved steadily in terms of heat resilience on the quality side as well. We also compared these variety-specific heat resilience estimates with their respective mean yields and test weights, which represent the predicted outputs under average climate changes in Figure 1.4. Our results show varieties that tend to produce higher mean yields or test weights under typical weather conditions also exhibit enhanced resilience to heat stress.

To evaluate yield relationships with various weather variables, multiple models were employed to validate initial results.

The "Yield (level)" model consistently shows a reduction in yields with a rise in average temperature compared to the preferred one. The "MinMax" model focuses on extreme temperatures and shows a significant decrease in yields with a dip in minimum temperatures. Maximum temperatures show no significant effect. These robustness checks confirm the consistency of the primary findings across different models.

Figure 1.3: Resilience and Introduction Year

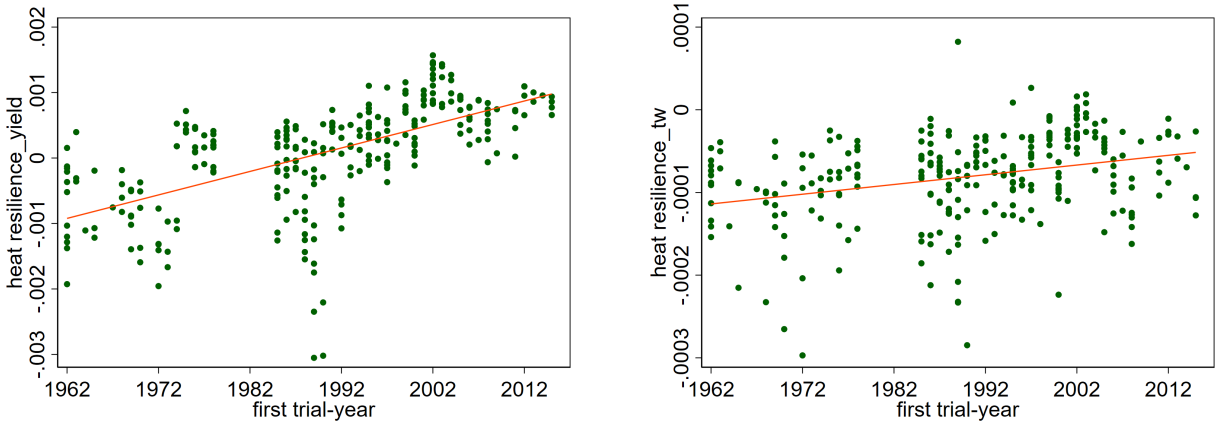

(a) Yield Effect (b) Quality Effect

Note: Figure 1.3 displays the varietal improvements in wheat resilience to high temperatures and the introduction year of the varieties. The mean performance of newly introduced varieties has improved in terms of both quantity (Panel a) and quality (Panel b).

Figure 1.4: Resilience and Mean Effect

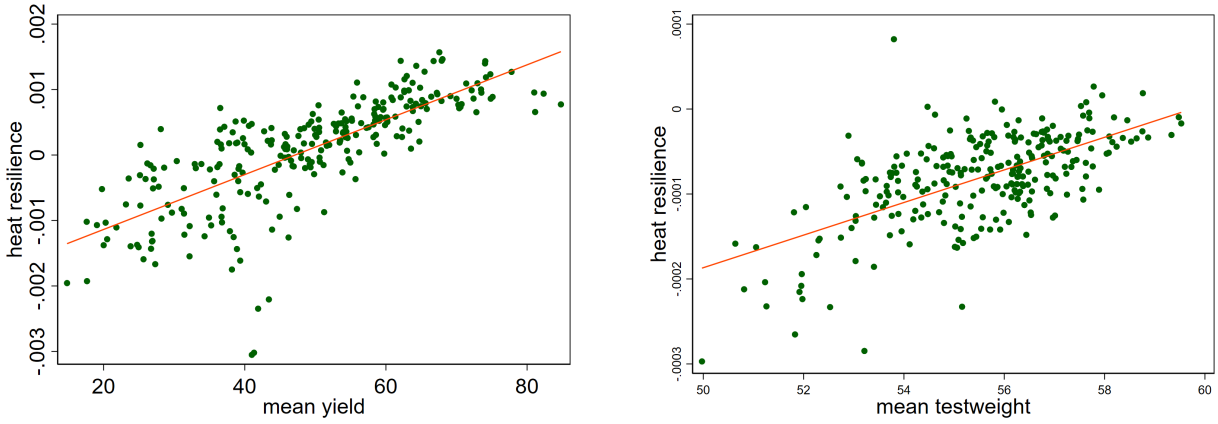

(a) Yield Effect (b) Quality Effect

Note: Figure 1.4 illustrates the improvements in wheat varieties with respect to resilience under high heat and their mean yields (or test weight). The varieties that are more resilient to high heat demonstrate a positive correlation with both yield (Panel a) and quality (Panel b).

1.8 Conclusion

Our study produces several findings about winter wheat production and its breeding efforts: The primary finding is that lower temperatures positively impact wheat, while medium and high temperatures have adverse effects. This relationship holds true across various model specifications, reinforcing the conclusion that temperature is a critical factor in wheat production.

Comparing these results to studies in other crops, such as rice in Japan, emphasize the crop-specific nature of climate impacts. Despite the complexities of these interactions, the consistency across different model setups in our study strengthens the reliability of these findings. Particularly concerning is the prediction of significant declines in both yield and quality with rising temperatures, a trend that is evident throughout the overall growing seasons of wheat, although it varies across each growth stage Another important aspect of this research is the examination of how different wheat varieties respond to high temperatures. Newer varieties show increased tolerance to heat, maintaining or even improving yield and quality under high-temperature conditions. This suggests progress in breeding wheat varieties capable of withstanding climatic challenges.

There are several limitations to our model. Firstly, we assume a fixed growing season due to the unavailability of planting date data. Consequently, the climate impacts are considered constant throughout the periods, making it challenging to discern the varied weather effects on different growth stages, which could introduce bias (Kawasaki and Uchida, 2016). Secondly, since our study does not monetarily compare the effects on yield (quantity) and quality, it remains unclear which impact is more dominant. Thirdly, the accuracy of our spatial data may be compromised when interpolating between weather stations and wheat sites.

In future research, we plan to investigate the heterogeneous effects of local weather conditions on wheat production. This approach is based on the understanding that climate change impacts vary across different geographic regions. By focusing on local weather patterns, we aim to address how specific climatic factors, such as regional temperature fluctuations, levels of precipitation, and extreme weather events, uniquely affect wheat yield and quality in various areas.

However, our study does confirm the critical impact of temperature on wheat production and the importance of developing heat-tolerant wheat varieties. These findings contribute to the development of strategies for sustaining wheat production under changing climatic conditions. Additionally, we expect to provide insights for agricultural practices and policy development, aimed at mitigating the adverse effects of global climate change on wheat crops.

1.9 Appendix

Table 1.4: Seasonal Weather Effect on Wheat Yields

	Seasons			
	Total	Fall	Winter	Spring
Intercept	5.0657***	5.3243***	4.0895***	4.6117***
	(0.2718)	(0.2384)	(0.1789)	(0.2934)
Year	-0.0086	-0.0107	-0.0023	-0.0164
	(0.0081)	(0.0104)	(0.0095)	(0.0100)
Year Sq.	0.0001	0.0000	-0.0000	0.0002
	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Low	0.0003***	-0.0009***	0.0021***	0.0004***
	(0.0001)	(0.0003)	(0.0007)	(0.0001)
Med	-0.0016***	-0.0010***	-0.0027***	-0.0021***
	(0.0001)	(0.0004)	(0.0009)	(0.0002)
High	-0.0005**	0.0002	-0.0018***	0.0107***
	(0.0002)	(0.0003)	(0.0005)	(0.0016)
Prec	0.0000	-0.0012***	0.0001	0.0006***
	(0.0001)	(0.0002)	(0.0002)	(0.0002)
Prec Sq.	-0.0000***	0.0000***	-0.0000***	-0.0000***
	(0.0000)	(0.0000)	(0.0000)	(0.0000)
R^2	0.4486	0.4197	0.4242	0.4430
Adjusted R^2	0.4367	0.4072	0.4118	0.4310
F-statistic	42.1082	13.4063	16.0404	31.0794

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.5: Seasonal Weather Effect on Wheat Quality

	Seasons			
	Total	Fall	Winter	Spring
Intercept	4.0611***	4.1060***	4.0406***	4.1433***
	(0.0393)	(0.0425)	(0.0256)	(0.0379)
Year	-0.0000	-0.0010	-0.0006	-0.0007
	(0.0014)	(0.0015)	(0.0015)	(0.0016)
Year Sq.	0.0000	0.0000	0.0000	0.0000
	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Low	0.0001***	-0.0000	0.0001	0.0001***
	(0.0000)	(0.0001)	(0.0001)	(0.0000)
Med	-0.0001***	0.0000	0.0001	-0.0003***
	(0.0000)	(0.0000)	(0.0001)	(0.0000)
High	-0.0002***	-0.0002***	-0.0003***	0.0018***
	(0.0000)	(0.0000)	(0.001)	(0.0002)
Prec	0.0000	-0.0000	0.0000	-0.0003***
	(0.0001)	(0.0000)	(0.0000)	(0.0000)
Prec Sq.	-0.0000***	0.0000	-0.0000	0.0000***
	(0.0000)	(0.0000)	(0.0000)	(0.0000)
R^2	0.3208	0.2765	0.2799	0.3391
Adjusted R^2	0.3052	0.2599	0.2634	0.3239
F-statistic	81.2389	5.0666	11.4751	112.5814

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.6: Robustness Check for Yield

	Model Specifications			
	Preferred	Yield	Avg	MinMax
Intercept	5.0657***	100.4841***	5.6416***	5.0735***
	(0.2718)	(9.0936)	(0.2379)	(0.2335)
Year	-0.0086	-0.2181	-0.0104	-0.0048
	(0.0081)	(0.2311)	(0.0089)	(0.0083)
Year Sq.	0.0001	-0.0016	0.0001	0.0000
	(0.0001)	(0.0032)	(0.0001)	(0.0001)
Low	0.0003***	0.0114***		
	(0.0001)	(0.0040)		
Med	-0.0016***	-0.0572***		
	(0.0001)	(0.0051)		
High	-0.0005**	-0.0036*		
	(0.0002)	(0.0071)		
Prec	0.0000	-0.0050*	0.0001	0.0003***
	(0.0001)	(0.0029)	(0.0001)	(0.0001)
Prec Sq.	-0.0000***	-0.0000***	-0.0000***	-0.0000***
	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Tavg			-0.0002**	
			(0.0001)	
Tavg Sq.			-0.0000***	
			(0.0000)	
Tmin				-0.0006***
				(0.0001)
Tmax				-0.0000
				(0.0000)
R-squared	0.4486	0.5387	0.4416	0.4431
Adj. R-squared	0.4367	0.5288	0.4296	0.4312
F-statistic	42.1082	65.8902	38.0699	38.8799

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.7: Robustness Check for Quality

	Model Specifications			
	Preferred	Test Weight	Avg	MinMax
Intercept	4.0611***	57.2945***	4.1933***	4.0540***
	(0.0393)	(1.8123)	(0.0386)	(0.0474)
Year	-0.0000	-0.0201	-0.0007	-0.0003
	(0.0014)	(0.0664)	(0.0016)	(0.0015)
Year Sq.	-0.0000	-0.0001	0.0000	-0.0000
	(0.0000)	(0.0009)	(0.0000)	(0.0000)
Low	0.0001***	0.0036***		
	(0.0000)	(0.0007)		
Med	-0.0001***	-0.0057***		
	(0.0000)	(0.0011)		
High	-0.0002***	-0.0088***		
	(0.0000)	(0.0012)		
Prec	0.0000	0.0015**	0.0000***	0.0001***
	(0.0000)	(0.0007)	(0.0000)	(0.0000)
Prec Sq.	-0.0000***	-0.0000***	-0.0000***	-0.0000***
	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Tavg			-0.0000**	
			(0.0000)	
Tavg Sq.			-0.0000	
			(0.0000)	
Tmin				-0.0001***
				(0.0000)
Tmax				0.0000***
				(0.0000)
R-squared	0.3208	0.3780	0.3149	0.3203
Adj. R-squared	0.3052	0.3637	0.2992	0.3048
F-statistic	81.2389	98.3757	63.9734	78.1176

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Figure 1.5: Quantity and Quality by Year

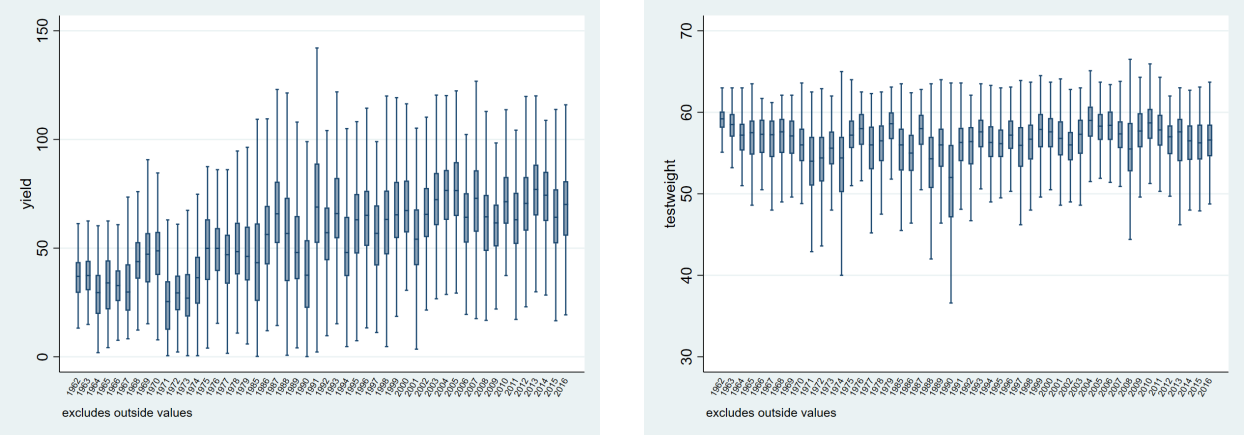

(a) Yield by Year

(b) Quality by Year

Note: These figures depict quantity and quality by year. While yield has an increasing trend over time and a large distribution, test weight has an even distribution, and is consistent over time.

Figure 1.6: Marginal Effects across Three Seasons

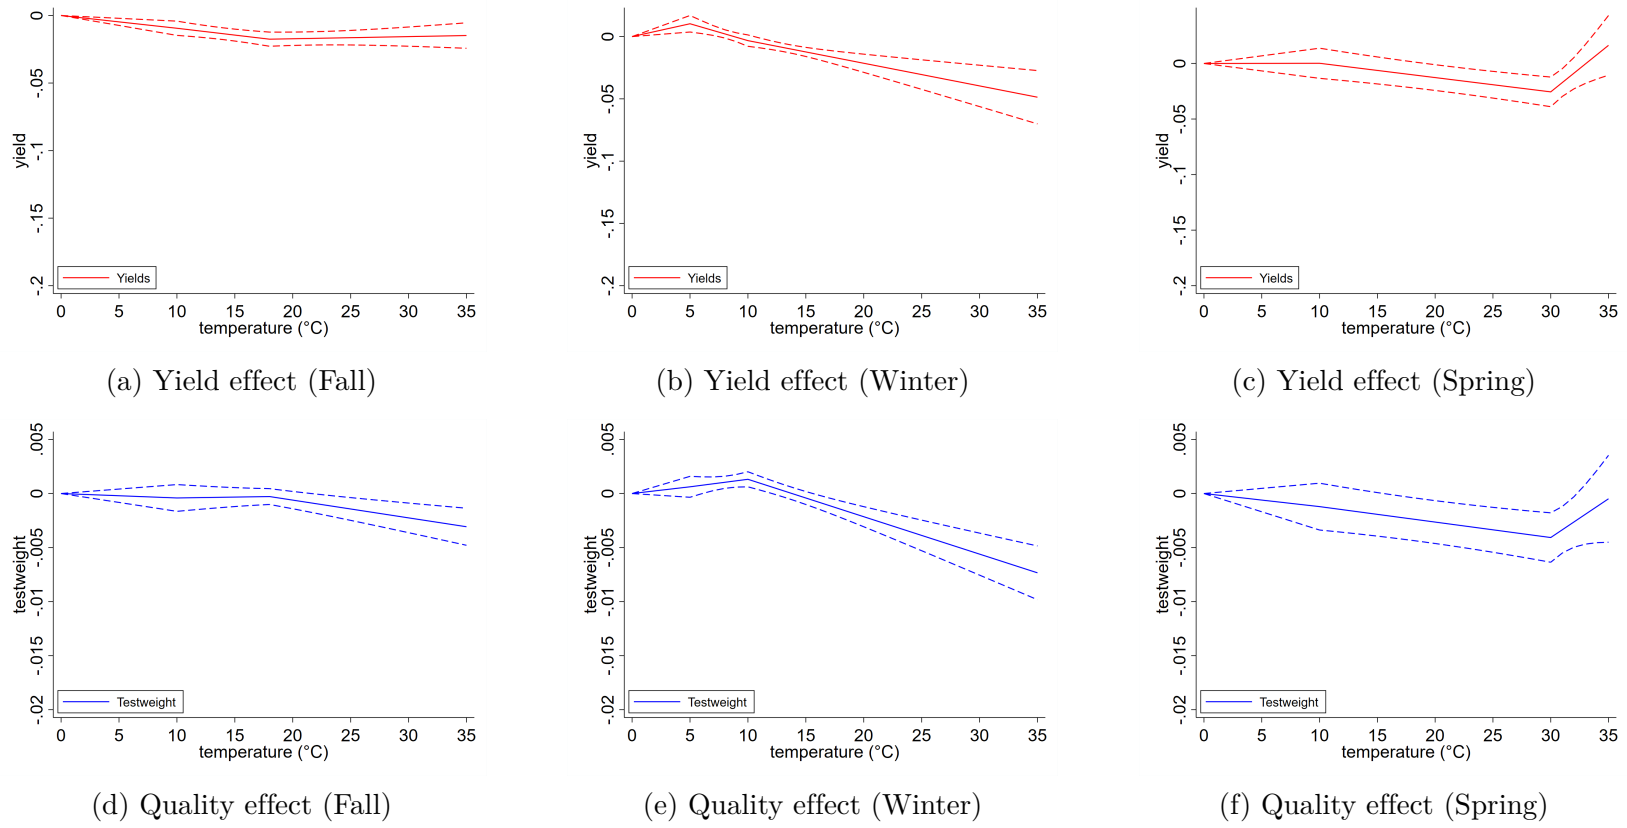

(a) Yield effect (Fall)

(b) Yield effect (Winter)

(c) Yield effect (Spring)

(d) Quality effect (Fall)

(e) Quality effect (Winter)

(f) Quality effect (Spring)

Note: Figure 1.6 illustrates the nonlinear relationship between positive temperature and yields (upper graph) as well as test weights (lower graph) during the three growing stages of Fall, Winter, and Spring. It depicts how log yield (and test weight) changes when wheat is exposed to an additional 1°C of temperature for 24 hours. The dashed lines represent 95% confidence intervals, with standard errors clustered by location.

Figure 1.7: Warming Scenario across Three Seasons

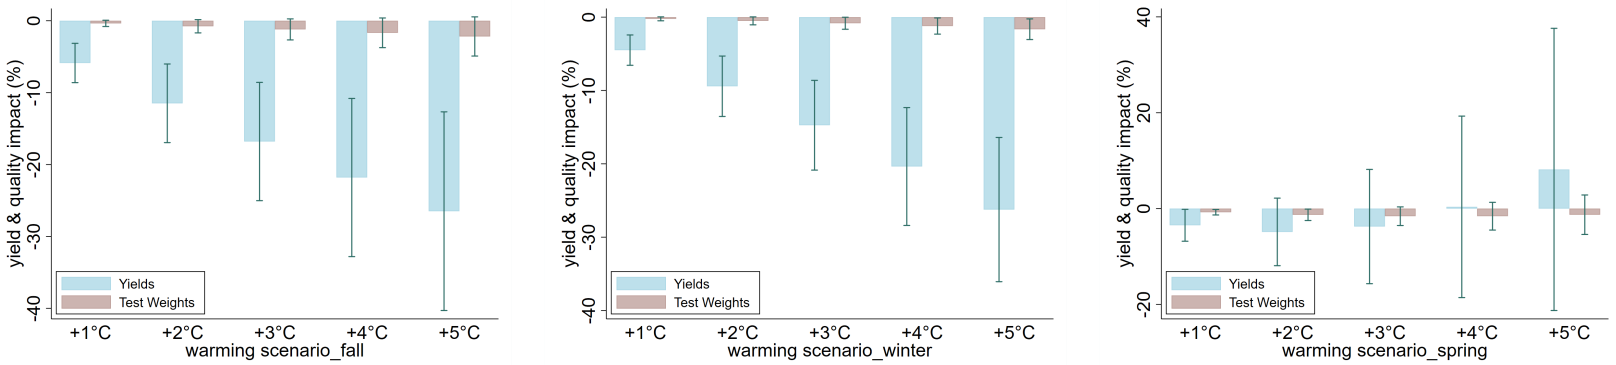

(a) Fall (b) Winter (c) Spring

Note: Figure 1.7 illustrates the predicted warming impacts during the growing stages of Fall, Winter, and Spring, respectively. The optimal thresholds are 10 and 18°C in the Fall, 5 and 10°C in the Winter, and 18 and 30°C in the Spring. Bars represent 95% confidence intervals, with standard errors clustered by location. The predicted negative effects are different across the growing stages.

Chapter 2

Semiparametric Analysis of Out-farm Migration in China

2.1 Abstract

This study investigates the nonlinear relationship between the internal migration rate and the income gap in China, a country where agriculture still forms a significant proportion of employment and contributes substantially to the economy. While previous studies have presumed a positive linear relationship between occupational migration and the sectoral income or wage gap, we challenge this assumption. We postulate that farm laborers may not immediately migrate upon reaching a certain income differential. Hence, we adopt a more flexible semiparametric model, instead of imposing strict functional form assumptions, to capture this complex relationship. We apply Robinson's (1988) double residual estimator to balance the advantages of nonparametric models with the structure of standard parametric models. Our findings indicate a non-monotonic relationship between off-farm migration rate and income gap, suggesting a potential complex interaction between these two variables. However, this alone does not provide a sufficient basis to assert the superiority of the semiparametric model over a linear one for explaining out-farm migration in China.

2.2 Introduction

The relationship between out-farm migration and the income gap has been a topic of interest in the literature, as it offers a valuable understanding of the driving forces of the development of countries. In numerous cases, a linearity between the two variables has been hypothesized suggesting that the migration rate responds proportionally to changes in the income gap – as the income gap increases, the migration rate would consistently increase, and vice versa (Barkley, 1990, Harris and Todaro, 1970, Mundlak, 1978). Although it is reasonable to assume that a larger income gap might provide a stronger incentive for agricultural employees to migrate to non-agricultural sectors in search of better economic opportunities, various factors can influence this relationship, making it more complex than a simple linear model.

Additionally, most research concentrates on occupational migration in developed countries (Barkley, 1990, Burda, 1995, Önel and Goodwin, 2014, Ramsey et al., 2023). These nations typically have a larger share of the workforce employed in the services and high-tech industries, and more formal labor markets with better social protection systems, wage regulations, and job security while developing countries often have a more significant agricultural sector and a lower share of employment in non-agricultural sectors. Hence, there still is a higher chance for individuals to migrate, and policies that promote industrialization and urbanization may continue to encourage migration from the agricultural sector to the industrial sector. Accordingly, it is crucial to consider the inter-sectoral migration in developing countries separately as their pace and dynamics may differ considerably.

This study aims to contribute to the existing literature on migration and income gap by examining the potential nonlinear nature of the relationship between out-farm migration and income gap in a developing country using semiparametric methods. By employing a flexible functional form that allows for the possibility of a more complex relationship, this

research seeks to provide new insights into the dynamics of migration patterns. Focusing on China's case, we explore whether varying patterns and drivers of inter-sectoral migration appeared along its development stage. While a vast literature on migration in China has been developed, few papers have studied migration using nationally representative data, which would be our other contribution.

The following chapters delve into the background of China's inter-sectoral migration, theoretical framework, data and methodology, empirical results, discussion and implications, and conclusion.

2.3 Literature Reviews

Identifying the determinants of migration has long been considered a crucial topic. The most extensive literature on migration is concerned with human movements at an international level, but there is also a large literature on migration considering an internal level within a country. Given the limited availability of labor, a country's economic development is partly attributed to the transition of surplus labor from the agricultural sector to the modern industrial sector, where the workforce can contribute to higher productivity levels. Although the methods and level of data used in inter-sectoral migration studies vary, economic factors are often mentioned as fundamental determinants for occupational movement.

Harris and Todaro (1970) have played a key role in conceptualizing the specific economic motivations for migration. The Harris and Todaro model does not presume that rural-urban wage differentials are the sole basis of migration, as is suggested by most migration theories. According to them, migration is a rational and intentional decision by workers that actively seek better living and working conditions. Furthermore, their analysis considers not only wage differences but also employment prospects in urban areas, thus determining migration

more by urban-rural differences in "expected" earnings than actual earnings. Migrants who expect a greater income differential may be more tolerant of high levels of unemployment in urban areas for this reason.

Empirical evidence at the aggregate level on out-farm migration factors has been provided by Mundlak (1978) and Barkley (1990). By conducting a massive cross-country analysis of 72 countries for the period from 1960 to 1970, Mundlak (1978) shows the out-migration from agriculture is determined by sectoral income differentials, the composition of the labor force, and the growth rate of the labor force.[1] Similarly, using the U.S. data from 1940 to 1985, Barkley (1990) finds that farm labor is responsive to the return ratio and the labor force ratio while less significantly so on other variables such as land value, urban unemployment rate, and government payments.

Meanwhile, the underlying assumption behind the studies discussed above is that off-farm migration rates are linearly related to income gaps. Under this assumption, labor flows immediately from areas with lower income to areas with higher income as the gap exceeds a certain level of moving cost, which is called a "Marshallian trigger," to maximize or increase their welfare and standard of living. In reality, these gaps alone do not necessarily constitute a sufficient reason for leaving one's sector. As migration is a complicated and multifaceted phenomenon, social, geographical, and individual factors may also affect its incidence and timing (Bukenya et al., 2003, Burda, 1995, Stark and Taylor, 1991). Accordingly, the decision-making process regarding occupational migration may result in non-linear migration patterns.

Some articles incorporate this nonlinearity into their analysis. In an extended cross-country study from the previous work, Larson and Mundlak (1997) estimate a nonlinear aggregate

[1] Mundlak (1978) explains in detail using income gap (average labor productivity) is a better measurement than the wage gap.

migration equation that includes a wedge between two sectors. However, the estimated value of the wedge is small and statistically insignificant.

One other approach that captures the nonlinearity is real option theory which suggests that migration may be postponed beyond the Marshallian trigger since the value of waiting may be sufficiently positive in the face of uncertainty (Dixit and Pindyck, 1994). Waiting can be advantageous in resolving uncertainty and eliminating the downside risk associated with costly changing occupations. Migration under the linear regression assumes paying an immediate cost in the expectation of future payoff. However, if the immediate cost is sunk and future returns are uncertain, the timing of migration becomes a critical consideration.

The theory of real option is first applied to occupational migration by Burda (1995) to explain the low migration rates in East to West Germany despite the large wage differentials. This article shows that individuals prefer to wait before migrating, even if the present value of the wage differential is positive. Rather, making a decision depends heavily on the existence of a community. This is because of the uncertainty and the sunk costs that are associated with the migration behavior.

Dennis and İşcan (2007) adopt the real option theory using a general equilibrium framework and emphasize that workers do not migrate during periods when the relative wage is below the threshold. These results are similar to those reported by Burda (1995). In terms of methodology, they use calibration methods and the parameters of the theoretical model in place of statistical methods to measure this threshold.

The threshold regression model for the real option approach is first applied by Önel and Goodwin (2014). This threshold model describes the jumping character or any other discontinuity in the relationship between variables. In their work using annual U.S. data from between 1948 and 2009, Önel and Goodwin (2014) observe the nonlinear threshold behavior

between wage difference and off-farm migration. However, it is unclear whether the results reflect the sunk cost or option value of waiting where people would not migrate.

Recently, Ramsey et al. (2023) raise important considerations about the validity of threshold models for representing migratory behavior based on real option theory. Using the most up-to-date data from the U.S. and Japan, compared to the previous studies, the authors empirically test for nonlinearity and find little evidence of threshold behavior in the aggregate migration equations. Their results suggest that while individuals might exhibit threshold behavior in migration decisions, these tendencies may not be as prominent when examined at an aggregate level.

In this paper, we augment the previous methodology by applying a semiparametric model to the Chinese case. Rather than suggesting a replacement, this methodological choice offers a complementary approach. The semiparametric model can identify potential nonlinear relationships in a less restrictive manner than the threshold model. This model remains compatible with real option theory, contributing to an understanding of how volatility and risk impact migration decisions, and further elucidating their role in determining the timing and pace of these transitions.

We apply this expanded methodology to the case of China, a developing country. This choice is prompted by a hypothesis that the nonlinear relationship between economic factors and out-farm migration is likely more apparent in developing countries, as opposed to their developed counterparts. This is because the factors driving migration in developing nations are often more varied and dynamic, reflecting the rapid changes these countries typically experience during their development process.

For instance, consider the cases of highly industrialized nations such as the U.S. and Japan. Their stable social and institutional structures, along with an abundance of non-agricultural

job opportunities, have led to a significant decrease in the proportion of farm workers to under 3% of the total workforce (Ramsey et al., 2023).

Among developing countries, China is a particularly intriguing case. While the transition from farm to non-farm sectors is ongoing like in many developing countries, China has unique socio-economic conditions and policies that influence its rural labor mobility. This does not imply that China is an exception among developing nations, but rather highlights the need to account for these distinct factors when studying its off-farm migration. In the next section, we provide a brief overview of these elements.

2.4 Background of Chinese Inter-sectoral Migration

2.4.1 Overview of Chinese Out-farm Migration

China has experienced rapid economic growth and urbanization over the past few decades, leading to significant changes in its social, economic, and demographic landscapes. Like many other developing countries, one of the most prominent features of this transformation is the large-scale migration of rural residents to urban areas in search of better employment opportunities and improved living standards. This phenomenon has played a crucial role in shaping China's labor market, economic development, and income distribution.

But at the same time, China stands as the world's leading agricultural economy, and its agricultural sector continues to make substantial contributions to poverty reduction and overall development. Using our data set from the National Bureau of Chinese Statistics, agriculture's contribution to China's GDP was approximately 10% in 2010. This is markedly different from the situation in developed countries like the U.S. and Japan, where agriculture accounts for only about 1% of GDP.

Even as China experiences rapid economic growth, a substantial proportion of its workforce remains engaged in agriculture. Although the percentage of labor employed in the sector decreased from 75% in 1978 to roughly 40% in 2010, the share of labor in the agricultural sector remains high as in Figure 2.1.(a). This paradox, where a large number of Chinese farmers continue to work in agriculture despite the country's rapid economic advancement, raises intriguing questions. Furthermore, it suggests a significant potential for future migration out of farming. [2]

In light of this context, our study aims to discern whether farm laborers opt to migrate as soon as a certain income differential is met. To explore this issue, we utilize Chinese data in our analysis. By incorporating the information available from the National Bureau of Statistics in China, we are able to closely examine labor trends in the country, as depicted in Figure 2.1.(a). It indicates a steady decrease in the proportion of labor involved in the farm sector, in contrast to a rise in the non-agricultural sector. This trend could be associated with ongoing industrialization and urbanization in China, where traditional farming jobs are gradually being replaced by non-farming opportunities.

However, in comparison to developed countries, China's agricultural sector continues to employ a considerable number of workers, even in recent times.

Figure 2.1.(c) presents a clear upward trend in off-farm migration over time, signifying a notable increase in individuals transitioning from the agricultural sector to non-farm employment opportunities overall. Despite the overall upward trend, it is notable to observe fluctuations in certain periods, including instances where the migration rate drops below zero. This fluctuation seems inconsistent with the persistent decrease in the labor share of the farm sector depicted in Figure 2.1.(a).

[2]However, we acknowledge that labor mobility in China is affected by governmental policies and regulations, which must be taken into account when evaluating the potential for increased migration from farming

Figure 2.1: Labor Share, Productivity and Migration Rates

(a) Labor Share of Farm and Non-farm Sector

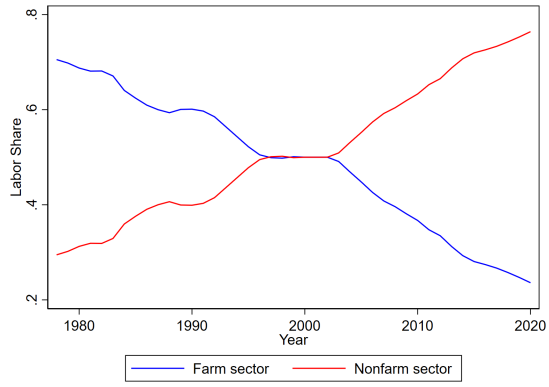

(b) Productivity of Farm and Non-farm Sector

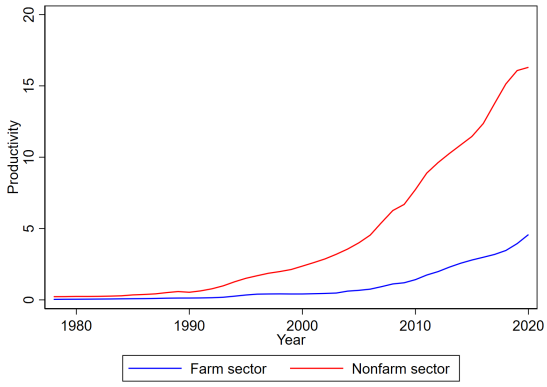

(c) Off-farm Migration Rates over time

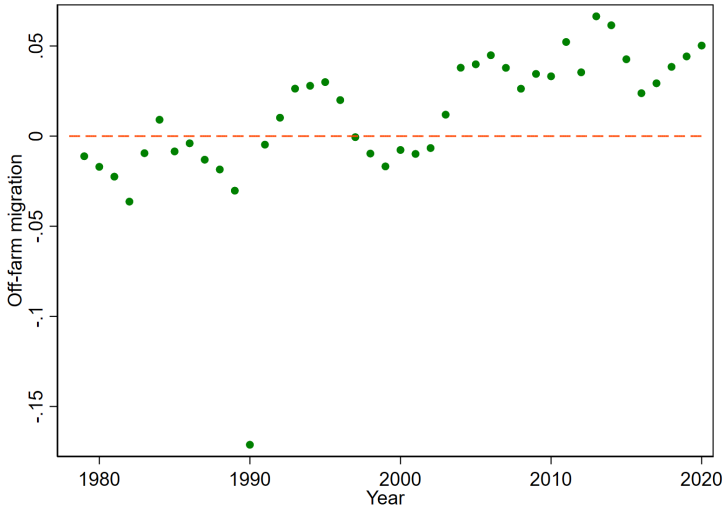

The explanation for this discrepancy lies in understanding the different measures that the two figures represent. The farm labor share in Figure 2.1.(a) is calculated as a fraction of the total workforce (comprising both farm and non-farm workers) that is employed in the farm sector. As such, even if the absolute number of farm workers increases from one year to the next (which would result in a negative off-farm migration rate in Figure 2.1.(c)), the farm worker share in Figure 2.1.(a) can still decrease if the increase in non-farm workers during the same period is proportionally greater.

On the other hand, the off-farm migration rate in Figure 2.1.(c) is calculated based on changes in the number of farm workers from year to year. This formula represents the proportional decrease in the farm labor force each year, which is interpreted as the rate of off-farm migration. Therefore, a negative migration rate means the number of farm workers in the current year is greater than in the previous year, but it does not solely imply an influx of workers from non-farm to farm jobs. Instead, it could reflect a natural increase in the farm labor force, such as from new entrants to the labor market or from previous non-workers taking up farm jobs.

Hence, it is necessary to consider the negative migration rates in Figure 2.1.(c) do not contradict the persistent decrease in the farm labor share seen in Figure 2.1.(a); instead, they reveal complexity in labor migration trends that are not captured by sectoral labor share alone.

Figure 2.1.(b) illustrates the progression of productivity disparities between the farm and non-farm sectors in China. The figure demonstrates a clear pattern of an expanding income differential over time, suggesting that incomes in the non-farm sector have consistently outpaced those in the farm sector. This expanding gap may be indicative of variations in productivity advancements and output value between the two sectors.

2.4.2 Chinese Unique Land System, Hukou

Inter-sectoral migration in China stands out due to the Hukou system, a Soviet-style household registration mechanism established in 1958. It categorizes Chinese citizens as either urban or rural, creating barriers to mobility and influencing migration patterns.

The Hukou system has historically favored urban residents, providing them with better jobs, living conditions, and services. In contrast, rural residents are often marginalized (Hao and Tang, 2015). However, the 1978 economic reforms shifted China to a market-based economy, prompting industrialization and urbanization, and increasing urban labor demands. Consequently, individuals with a rural hukou took advantage of relaxed migration rules to move to cities and boost non-agricultural employment.

However, there is an ongoing struggle to balance urban and rural interests. Welfare policies sometimes prevent rural workers from moving with families, leading to a rise in 'floating' migrants. Additionally, even with the option to change hukou status becoming easier, it remains a rare choice among rural migrants.

A key aspect of China's household system is that while agricultural workers can use land, they are not allowed to sell or transfer it. This restriction means that if they leave agriculture, they forfeit potential land income. Transitioning from a rural to an urban hukou also means losing certain rural privileges. Given that rural households often rely on land for income and receive housing plots without charge, it is clear that land ownership heavily influences their migration decisions (Hao and Tang, 2015).

Considering the Hukou system's complexities, permanent migration in China is less common than one might think. Chinese farmers tend to split their labor between sectors rather than migrate permanently. Recognizing the benefits of rural hukou including land access and relaxed family planning, many prefer the flexibility of moving between urban and rural

areas. This way they can balance higher urban wages while retaining the benefits of rural land Chen and Fan (2016). In this regard, our study assumes that their migration decisions might not be solely based on income gaps, but a more significant disparity might be needed to drive them.

2.5 Theoretical Framework

We borrow the theoretical model of Larson and Mundlak (1997) that provides a link from the household-level to the aggregate migration function. The indirect utility model for migration decisions by Larson and Mundlak (1997) provides a basis for understanding the factors influencing migration behavior. It highlights the role of income, individual and household characteristics, and prices of goods and services, as well as the costs associated with migration, in shaping migration decisions.

Consider a representative individual with age (g) and demographics (z) who works either in the non-farm sector (nf) or in the farm sector (f). The indirect utility function for sector $j \in \{nf, f\}$ can be written as

$$v(g, z, j) = v(p_j, w_j, g, z, c_j(d_j, g, z, y)), \tag{2.1}$$

where p_j is the price of consumer goods, w_j is the expected income in occupation j, and c_j is the cost of migration that is increasing in distance of other employment (d_j) and decreasing in the general level of infrastructure development. Thus, the indirect utility functions capture the benefits of migration, such as higher income, improved living conditions, and better job opportunities. However, migration decisions also involve costs, such as transportation, housing, and other relocation expenses, as well as psychological costs, such as leaving familiar

surroundings, family, and friends (Bukenya et al., 2003, Burda, 1995, Stark and Taylor, 1991).

Indirect utility is also a function of age. The migration decision is made based on lifetime expectation of utility, rather than utility at a certain age. Thus, we derive the present value of the lifetime utility of an individual (g, z, j) for a given retirement age T and time-discounting rate ρ.

$$V(g, z, j) = \int_g^T \exp(-\rho t) v(t, z, j) dt, \tag{2.2}$$

which aggregates the indirect utility from the age of g to T in occupation j.

Following Larson and Mundlak (1997), we assume that an individual migrates from the agricultural occupation to the non-agricultural occupation if the present value of lifetime utility in the former is higher than the present value of lifetime utility in the latter.

$$V(g, z, nf) > V(g, z, f). \tag{2.3}$$

To arrive at the migration function, we use a measure $h_i(f, nf)$ for the out-farm migration.

$$h_i(f, nf) = \begin{cases} 1 & \text{if } V(g, z, nf) > V(g, z, f) \\ 0 & \text{otherwise.} \end{cases} \tag{2.4}$$

That is, we have $h_i(f, nf) = 1$ only when migration to the non-farm sector from the farm sector gives a positive return for household i. Note that individuals in both sectors may migrate to the other sector depending on their lifetime utility. To measure the in-farm migration, we use a measure $h_i(nf, f)$ which is similarly defined.

Finally, we take into account both out-farm and in-farm migrations to count for the net influx of the labor force to the non-farm sector $M(f, nf)$.

$$M(f, nf) = \sum_{i=1}^{L_f} h_i(f, nf) - \sum_{i=1}^{L_{nf}} h_i(nf, f), \tag{2.5}$$

where L_f and L_{nf} are the sectoral labor force from which $\sum_{i=1}^{L_f} h_i(f, nf)$ and $\sum_{i=1}^{L_f} h_i(f, nf)$ indicate the number of out-farm migration and the number of in-farm migration, respectively.

Adding elements from the real option theory to our model, we can articulate the potential postponement of a migration decision under certain conditions. Specifically, the decision to migrate can be treated as a 'real option' that an individual may choose to exercise when the anticipated benefits outweigh the costs and uncertainties.

To this end, let's redefine our migration decision function $h_i(f, nf)$ to take into account the option to delay migration. The choice to migrate now becomes contingent not only on the difference in present value of lifetime utilities between sectors but also on a threshold level of this difference which justifies migration.

We can rewrite the migration function as follows:

$$h_i(f, nf) = \begin{cases} 1 & \text{if } V(g, z, nf) - V(g, z, f) > \theta_i \\ 0 & \text{otherwise,} \end{cases} \tag{2.6}$$

where θ_i is the individual-specific threshold of the income gap, beyond which the individual decides to migrate. This threshold can be influenced by various factors, including risk tolerance, personal and family circumstances, local labor market conditions, and individual perceptions of the costs and benefits of migration.

In our theoretical framework, we attempt to capture the essence of real option theory, where individuals have the option to postpone migration amidst uncertainty. This approach acknowledges that migration decisions are influenced by demographic factors, current in-

come, potential future income, and uncertainties, providing a comprehensive explanation for observed migration patterns. However, acknowledging that threshold methodologies have shown limited validation in previous literature, we transition from a strict threshold behavior to a semiparametric model while retaining the core concepts of real option theory.

2.6 Model

2.6.1 Data

We use annual national time series data from 1978 to 2020. Key variables of interest originate from the China Statistical Yearbook compiled by the National Bureau of Statistics (NBS).[3] Specifically, data on non-farm sector employment (L_{nf}) and farm sector employment (L_f) are derived from the '4-2 Employment and Wages' section. GDP figures for agricultural and non-agricultural industries (y_{nf}, y_f) are available in the '3-6 Value-added by Sector' section. In this dataset, the 'farm sector' encompasses all labor and income involved in agriculture, forestry, animal husbandry, and fishery. The non-agricultural sector information is deduced by subtracting agricultural labor or GDP from the total labor force or GDP, as it is not provided separately.[4]

The non-farm unemployment rate (u_{nf}) is sourced from the International Labour Organization (ILO) database. Average years of total schooling (*educ*) data are obtained from Barro and Lee (2013), which present educational attainment data from 1950 to 2015 in five-year intervals.

[3] http://www.stats.gov.cn/sj/ndsj/2021/indexeh.htm

[4] The added value of agriculture and related industries is defined as the value created by all permanent resident units in China engaged in the production activities of these industries within a certain period of time.

The data that constitutes the Chinese official statistics is gathered primarily from three categories: national economic census data, routine statistical data, and administrative data. The routine statistical data includes annual industry data for national economic accounting and annual statistical reports of relevant specialties. The administrative data contains Department Statistics from the Ministry of Agriculture and Rural Areas and trade association statistics. These varied sources of data can introduce potential inconsistencies, which could vary over time, across different regions, or among various collecting methodologies in China. Holz (2014) has raised concerns about potential falsification, measurement bias, and post-collection adjustments in Chinese official statistics. Acknowledging these challenges, the National Bureau of Statistics has implemented retrospective revisions in recent years to enhance the accuracy and consistency of the data (Bai et al., 2006). In this study, despite the potential limitations, the most recent edition of the China Statistical Yearbook (2021) was used. It provides the most up-to-date and revised data to maximize accuracy and consistency.

Nevertheless, the China Statistical Yearbook (2021) notes a significant shift in data collection practices during our period of study. Particularly after 1990, the compilation of labor force and employment statistics increasingly relied on alternative census surveys. This adjustment in data sourcing could generate outlier values, which might distort our analysis. In anticipation of this possible irregularity, we have introduced an annual dummy variable, termed 'dum90', into our empirical models. This methodological inclusion will help mitigate the potential influence of these outliers, ensuring more reliable and accurate outcomes.

Table 2.1 presents summary statistics and descriptions of the variables.

Table 2.1: Summary Statistics and Variable Description

Variable	Description	Mean	Std. Dev.
L_{nf}	Number of employees in non-farm sector (10 thousand)	35,205.47	14,497.86
L_f	Number of employees in farm sector (10 thousand)	30,583.67	5,909.237
y_{nf}	Value added by non-farm sector (100 million yuan)	220,103.5	284,417.5
y_f	Value added by farm sector (100 million yuan)	23,667.87	23,751.21
APL_{nf}	Labor productivity in non-farm sector (y_{nf}/L_{nf})	4.45	4.99
APL_f	Labor productivity in farm sector (y_f/L_f)	.97	1.22
ur	Unemployment rate in non-farm sector	3.5	.97
age	Percent of labor force $\geq$		
$educ$	Avg. years total schooling	7.39	1.08
$dum90$	Dummy variable post 90	0.68	0.47
M_t	Out farm migration rate	1.04%	3.93%
r	$\log(APL_{nf}/APL_f)$	1.57	.16
u	$\log(ur)$	1.21	.3
g	$\log(L_{nf}/L_f)$	.061	.603

2.6.2 Parametric Estimation and Results

In line with Larson and Mundlak (1997), we acknowledge that the wage rate often serves as an inadequate indicator of economic incentives because it merely forms a partial component of income and is highly susceptible to measurement errors. Therefore, we choose the income differential, as measured by GDP per worker, to be a variable indicative of economic incentives that can trigger sectoral change. Farm employment indicates the number of total workers in the agricultural sector including farm operators and hired workers.

We adopt a basic Harris-Todaro type model to analyze the migration rate (M_t) from the farm sector to the non-farm sector. The empirical Ordinary Least Squares (OLS) migration model for farm workers is as follows:

$$M_t = \alpha_0 + \alpha_1 r_{t-1} + \alpha_2 u_{t-1} + \alpha_3 g_{t-1} + \alpha_4 lv_{t-1} + \epsilon_t, \tag{2.7}$$

where r is the gap in average GDP per worker between the farm and non-farm sectors ($log(R_{nf}) - log(R_f)$), u is the log of the unemployment rate of the non-farm sector, g is the log of the ratio of employees in the non-farm sector to that in the farm sector, and lv is the log of real farmland values. Note that equation (2.7) is expressed using a lagged form to all independent variables to handle potential endogeneity problems (i.e., a causal relationship between migration and income).

The dependent variable, off-farm migration rates M_t, is defined as a percentage change in agricultural employment from one year to the next, following Barkley (1990) and Önel and Goodwin (2014).

$$M_t = \frac{(L_{f,t-1} - L_{f,t})}{L_{f,t-1}}. \tag{2.8}$$

In the context of China, where real farmland values are unavailable, we modify the model:

$$M_t = \alpha_0 + \alpha_1 r_{t-1} + \alpha_2 u_{t-1} + \alpha_3 g_{t-1} + \alpha_4 educ_{t-1} + \alpha_5 dum90 + \epsilon_t, \tag{2.9}$$

Here, *educ* represents the log of the average education period for employment from age 25 to 64. Note that the unavailability of accurate land value data in China is mainly due to the aforementioned unique property system, characterized by government land ownership and the sale of land-use rights, instead of the land itself. The *educ* variable is not a traditionally utilized measure in many existing papers on sectoral migration, except for Ramsey et al. (2023). In the context of our study, *educ* not only compensates for the lack of real farmland values data but is also expected to serve as a proxy for the effects of China's hukou system. The hukou system, which controls access to essential public services such as education and healthcare based on one's registered household location, plays a vital role in influencing intersectoral migration. Thus, shifts in the *educ* variable can be considered to indirectly reflect the changes in the impacts of the hukou system on migration decisions. However, it's important to acknowledge that *educ* captures only a portion of the hukou system's influence and could incorporate additional aspects.

To adjust for the significant revisions in the data by the different survey institutes, we include a dummy variable for 1990 ($dum90$ if $year > 1990$, and $dum90 = 0$ if $year \leq 1990$).[5] This variable accounts for potential shifts in the data that occurred in that year.

The findings from equation (2.9) are presented in Table 2.2. The OLS results indicate a coefficient of 0.0852 for the income gap, with a robust standard error of 0.044 and a t-ratio of 1.95. This t-ratio is just below the conventional threshold for statistical significance. While the results suggest that the income gap between the farm and non-farm sectors has an

[5]This methodology has been adopted from Barkley (1990) and Önel and Goodwin (2014).

effect on migration rates, the proximity of the t-ratio to the significance threshold means the relationship is not clearly established. Another notable result is that the employment ratio has a statistically significant impact on the rate of migration from farming, as evidenced by its significance at the 5% level.

Table 2.2: OLS Estimation

Variable	Coefficient	Robust s.e.	t-ratio
income gap (r)	0.0852*	0.044	1.95
employment ratio (g)	0.0906**	0.036	2.52
unemployment rate (u)	-0.0253	0.017	-1.50
education (e)	-0.2672	0.164	-1.63
dum90	0.0658	0.042	1.55
Dependent Variable: Migration rate			
Sample size: 42			
Durbin–Watson d-statistic = 1.964979			

Upon extending our model to incorporate the square of the income gap as an extra regressor, both income coefficients lose their statistical significance. This pattern persists even with the inclusion of the cubed income term alongside the income gap and its square in the model; all three income coefficients continue to lack statistical significance (as shown in Table 2.3). This outcome suggests two observations:

1. Increasing the complexity of the model from a quadratic to a cubic relationship does not lead to significant improvements.

2. A linear relationship may not adequately capture the impact of income on migration, pointing to the presence of potential nonlinearities.

Table 2.3: Estimates from Quadratic and Cubic Models

Variable	Coefficient	Robust s.e.	t-ratio
Quadratic model			
income gap	-1.26	1.05	-1.20
income gap^2	0.41	0.33	1.27
Cubic model			
income gap	-11.10	11.33	-0.98
income gap^2	6.51	6.77	0.96
income gap^3	-1.25	1.34	-0.93
Dependent variable: Migration rate			
Same regressors as above besides income			
Durbin–Watson d-statistic(Quadratic) = 1.999327			
Durbin–Watson d-statistic(Cubic) = 2.023515			

To account for these potential nonlinearities, we adopt the semiparametric approach proposed by Robinson (1988). This approach allows the shape of the relationship between income and migration to be determined by the data itself, offering more flexibility compared to rigid linear or polynomial models. Also, we only consider income as a potential candidate for a nonlinear independent variable.

2.6.3 Semiparametric Partially Linear Regression Model

Our study aims to discern relationships between dependent and independent variables without enforcing stringent functional form assumptions. Although this approach provides enhanced flexibility, it comes with its own set of challenges. Specifically, nonparametric models necessitate larger sample sizes to achieve similar statistical power as their parametric coun-

terparts. However, considering our research's sample size is limited to 42, solely depending on nonparametric models is impractical. We compromise this by employing Robinson's semiparametric double residual estimator (Robinson, 1988), which combines the benefits of nonparametric models with the structured approach of standard parametric models. In this model, the regressors specify the conditional mean of the dependent variable as linear in some parameters, while potentially nonlinear in others. This flexibility means we can capture the influence of factors like the income gap, which might have a nonlinear impact on migration without needing to specify the exact functional form of this relationship.

Our semiparametric partially linear model is formulated as:

$$\mathbf{y} = \mathbf{X}\beta + \mathbf{m}(\mathbf{z}) + \epsilon, \tag{2.10}$$

where the regressors $(\mathbf{X}, \mathbf{z})$ define the conditional mean of $\mathbf{y}$ as linear in $\mathbf{X}$, but it could be nonlinear in $\mathbf{z}$. Specifically, applying the conditional expectation on $\mathbf{z}$ yields the following:

$$\mathbb{E}(\mathbf{y} \mid \mathbf{z}) = \mathbb{E}(\mathbf{X}\beta \mid \mathbf{z}) + \mathbb{E}(\mathbf{m}(\mathbf{z}) \mid \mathbf{z}) + \mathbb{E}(\epsilon \mid \mathbf{z}) \tag{2.11}$$

$$= \mathbb{E}(\mathbf{X} \mid \mathbf{z})\beta + \mathbf{m}(\mathbf{z}). \tag{2.12}$$

Subtracting equation (2.12) from equation (2.10) gives:

$$\mathbf{y} - \mathbb{E}(\mathbf{y} \mid \mathbf{z}) = (\mathbf{X} - \mathbb{E}(\mathbf{X} \mid \mathbf{z}))\,\beta + \epsilon, \tag{2.13}$$

where $(\mathbf{y} - \mathbb{E}(\mathbf{y} \mid \mathbf{z}))$ is the conditional residual from the regression of $\mathbf{y}$ on $\mathbf{z}$ and $(\mathbf{X} - \mathbb{E}(\mathbf{X} \mid \mathbf{z}))$ is the conditional residual from the regression of $\mathbf{X}$ on $\mathbf{z}$. In this sense, equation (2.13)

is the residual regression. Thus, the estimated coefficient vector is

$$\hat{\beta} = \Big[(\mathbf{X} - \mathbb{E}(\mathbf{X} \mid \mathbf{z}))^T \cdot (\mathbf{X} - \mathbb{E}(\mathbf{X} \mid \mathbf{z}))\Big]^{-1} \cdot ((\mathbf{X} - \mathbb{E}(\mathbf{X} \mid \mathbf{z}))^T \cdot (\mathbf{y} - \mathbb{E}(\mathbf{y} \mid \mathbf{z})), \tag{2.14}$$

with variance

$$\mathrm{Var}(\hat{\beta}) = \sigma_\epsilon^2 \Big[(\mathbf{X} - \mathbb{E}(\mathbf{X} \mid \mathbf{z}))^T \cdot (\mathbf{X} - \mathbb{E}(\mathbf{X} \mid \mathbf{z}))\Big]^{-1}, \tag{2.15}$$

where σ_ϵ^2 is the variance of the error term ϵ. We can finally fit the nonlinear relation between $\mathbf{y}$ and $\mathbf{z}$ using the estimated parameter of β in the following nonparametric model.

$$\mathbf{y} - \mathbf{X}\hat{\beta} = \mathbf{m}(\mathbf{z}) + \epsilon. \tag{2.16}$$

This approach from Robinson (1988) suggests estimating $\mathbb{E}(\mathbf{y} \mid \mathbf{z})$ and $\mathbb{E}(\mathbf{x} \mid \mathbf{z})$ first by the Nadaraya-Watson (NW) kernel regression (Nadaraya, 1964, Watson, 1964). These NW estimators then replace the parameters in equation (2.14) to obtain $\hat{\beta}$.

In our model, $\mathbf{X}\beta$ represents the parametric part of the equation (2.10). $\mathbf{m}(\cdot)$ indicates the nonparametric part which permits the use of any form of the unknown function $\mathbf{m}$. $\mathbf{z}$ is the variable that affects the dependent variable but whose functional form is unknown, which is the income gap in our case, and ϵ is the error term. From equation (2.10), a consistent estimation of β can be obtained without the explicit modeling of $\mathbf{m}(\mathbf{z})$ by estimating $E(\mathbf{y}|\mathbf{z})$ and $E(\mathbf{X}|\mathbf{z})$ through a nonparametric regression method and subsequently substituting these estimations into the model in equation (2.16). The function $\mathbf{m}(\mathbf{z})$ can be estimated nonparametrically by regressing $(\mathbf{y} - \mathbf{X}\hat{\beta})$ on $\mathbf{z}$.

2.7 Results and Discussion

2.7.1 Semiparametric Results

Table 2.4 presents the semiparametric estimates of the parameter β, including the same explanatory variables as the OLS estimation, except for income differential. For comparison, both linear coefficients and their respective t-ratios are reported. Overall, the semiparametric estimates align closely with their linear counterparts: Specifically, the coefficient for employment ratio (g) is statistically significant at the 5% level, suggesting a positive relationship between the employment ratio and the migration rate. This is consistent with the findings from our linear model, further confirming the robustness of this relationship.

The coefficient for the education variable demonstrates a significant negative relationship with migration rate in both the linear and semiparametric models. This indicates an inverse correlation between education levels and off-farm migration rates, which may initially appear counterintuitive.

Table 2.4: Estimates from the Semiparametric Model

Variable	Partial linear		Linear	
	Coefficient	t-ratio	Coefficient	t-ratio
employment ratio (g)	0.0898**	2.7	0.0906**	2.53
unemployment rate (u)	-0.0223	-0.84	-0.0253	-0.89
education (e)	−0.294**	-2.6	−0.2672**	-2.23
dum90	0.0792***	3.35	0.0658***	2.72
Dependent Variable: Migration rate				
Sample size: 42				

Figure 2.2: Semiparametric Results

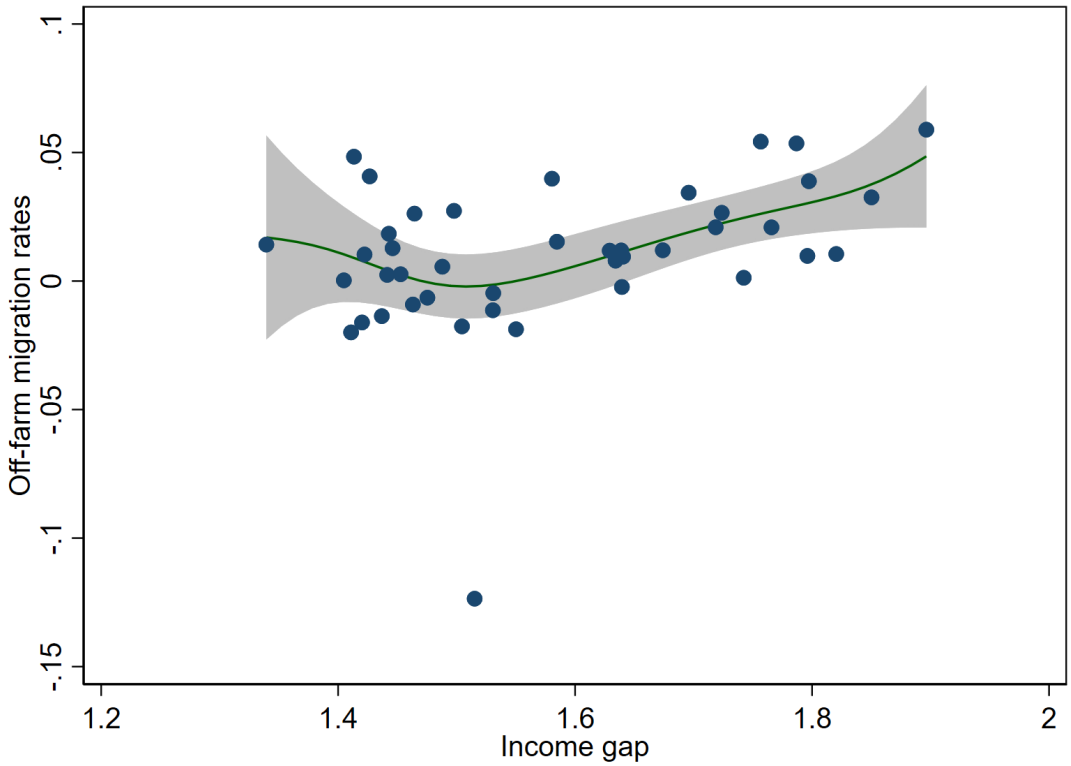

Several potential explanations could account for this outcome. Firstly, individuals with

higher levels of education may find improved opportunities within their local sectors, thereby reducing the need to migrate for work in the case of China. Secondly, there could be barriers in the non-agricultural sector that deter individuals with higher education from migrating. These might include a competitive job market, high living costs, or other socio-cultural challenges. Interestingly, the hukou system in China plays a crucial role in shaping these dynamics.

The estimates on the $dum90$ variable in both models show the considerable influence of data revision in 1990 while the coefficients for the unemployment rate are not statistically significant. Note that this model does not include an intercept parameter, as it can be absorbed into the nonparametric component $\mathbf{m}(\mathbf{z})$, where income serves as the variable $\mathbf{z}$ in equation (2.16).

The most significant difference between the two models lies in the treatment of the income gap variable. The OLS model identified a less significant (at the 10% level) linear relationship, whereas the semiparametric model uncovered a subtle, non-linear association between the income gap and migration rate.

The results of our semiparametric model are visualized in Figure 2.2. Here, the income gap ($\mathbf{z}$) is graphed on the horizontal axis, and the off-farm migration rates are on the vertical axis. This graph demonstrates that migration rates initially appear to decrease and then increase as the income gap between the farm and non-farm sectors widens.

Our findings suggest that at smaller income gaps (below approximately 1.5), the difference in economic opportunities is not substantial enough to prompt rural workers to migrate, resulting in low migration rates. Only if the income gap becomes sufficiently large, the incentives for farm labors to migrate to urban sectors begin to outweigh the factors that initially suppressed migration. At this stage, the migration rate would start to increase

again in China.

Furthermore, this observed nonlinear relationship aligns with the insights from the real option theory presented in our theoretical framework. The theory suggests that individuals may decide to postpone migration until the income gap reaches a particular level. In the context of our research findings, when the log of the income gap reaches around 1.5 in China, the expectation on higher income from migration starts to become significantly appealing. It is at this point that individuals decide to migrate, or in other words, "exercise their option" to migrate. This interpretation connects with the theory and illustrates the practical considerations influencing the timing of migration decisions.

2.7.2 Specification Test

To assess whether the nonparametric component of a semiparametric empirical model can be effectively approximated by a polynomial functional form, we utilize a specification test based on Hardle and Mammen (1993)'s approach. This method involves fitting the dataset to both a polynomial and a nonparametric model, followed by a comparison using a goodness-of-fit test. If the nonparametric model proves to be significantly superior to the polynomial model, it indicates that the polynomial approximation is inadequate, and the nonparametric component is essential for accurately capturing the intricate relationships among the variables.

The proposed test statistic, denoted as T, is defined as:

$$T_n = N\sqrt{h}\sum_{i=1}^{N}\left\{\hat{f}(z_i) - \hat{f}(z_i,\theta)\right\}^2 \pi(\cdot). \tag{2.17}$$

To determine the critical values, Hardle and Mammen (1993) recommend employing the

wild bootstrap technique for simulating test statistic values. The wild bootstrap generates a collection of resampled data by resampling the residuals from the initial model fit and subsequently re-estimating the test statistic. This procedure is iteratively performed to produce a distribution of simulated test statistics, from which critical values and p-values can be derived.

- H_0: Parametric and non-parametric fits are not different
- H_1: Parametric and non-parametric fits are different

Hence, if the null hypothesis (H_0) cannot be rejected, it suggests that the differences between the fits are not statistically significant, implying that using a parametric model, such as polynomial adjustment, could be as effective as a non-parametric approach for this case.

In other words, the inability to reject the null hypothesis suggests that the polynomial model is a satisfactory representation of the nonparametric component, and the semiparametric model may not be necessary to describe the complex relationships between the variables accurately.

Below are the results of the Hardle and Mammen (1993) test when the polynomial adjustment tested is of order 1, 2, or 3 respectively. These outputs do not present the results concerning the parametric part since they are the same as in the output presented above. When we compare it with a polynomial adjustment of degree 1, the null is rejected which suggests that the semiparametric model is more appropriate to be used and there can be a nonlinear linkage between the off-farm migration rate and the income differences between the sectors. But when we compare it with a polynomial adjustment of degree 2 or 3, the null is no longer rejected, which means that instead of a semiparametric model, a pure parametric model with a polynomial fit of degree 2 or 3 of income gap could be used.

2.7.3 Robustness Check: Controlling for Heteroskedasticity

We conduct robustness checks to ensure that our results are not overly sensitive to specific model specifications or estimation techniques. One potential issue in regression analyses is heteroskedasticity, which occurs when the variability of the error term is not constant across all levels of the explanatory variables. Heteroskedasticity, if present and unaccounted for, can lead to inefficient estimates and incorrect standard errors. This, in turn, could yield misleading inferences about the relationships we are studying.

In the context of our semiparametric model, heteroskedasticity could potentially bias the estimation of the parameters in the parametric part of our model, β, and the nonparametric function $\mathbf{m}(\mathbf{z})$. As such, we find it necessary to perform a robustness check to determine whether our model is susceptible to heteroskedasticity.

To do so, we adjust our semiparametric model to control for potential heteroskedasticity. This involves using heteroskedasticity-consistent standard errors, a common approach to dealing with this issue. It is crucial to note that our primary interest lies in the consistency of the estimates, which is not affected by heteroskedasticity. However, the efficiency and the standard errors of our estimates, which are important for hypothesis testing, could be affected.

Upon controlling for heteroskedasticity, we find that the results are nearly the same as those from our original model. This reinforces the robustness of our original results, as they were not significantly influenced by the presence of heteroskedasticity. The only noticeable difference was a slight decrease in the t-ratios, which is expected given that the heteroskedasticity-consistent standard errors are typically larger than the traditional ones.

Therefore, while the core findings remain unchanged, this analysis illustrates our results are robust to this potential issue.

Table 2.5: Controlling Heteroskedasticity

Variable	Partial linear		Heteroskedasticity controlled	
	Coefficient	t-ratio	Coefficient	t-ratio
employment ratio (g)	0.0898**	2.7	0.0898**	2.64
unemployment rate (u)	-0.0223	-0.84	-0.0223	-0.84
education (e)	-0.294**	-2.6	-0.294*	-1.81
dum90	0.0792***	3.35	0.0792*	1.81
Dependent Variable: Migration rate				
Sample size: 42				

Figure 2.3: Semiparametric Results after controlling Heteroskedasticity

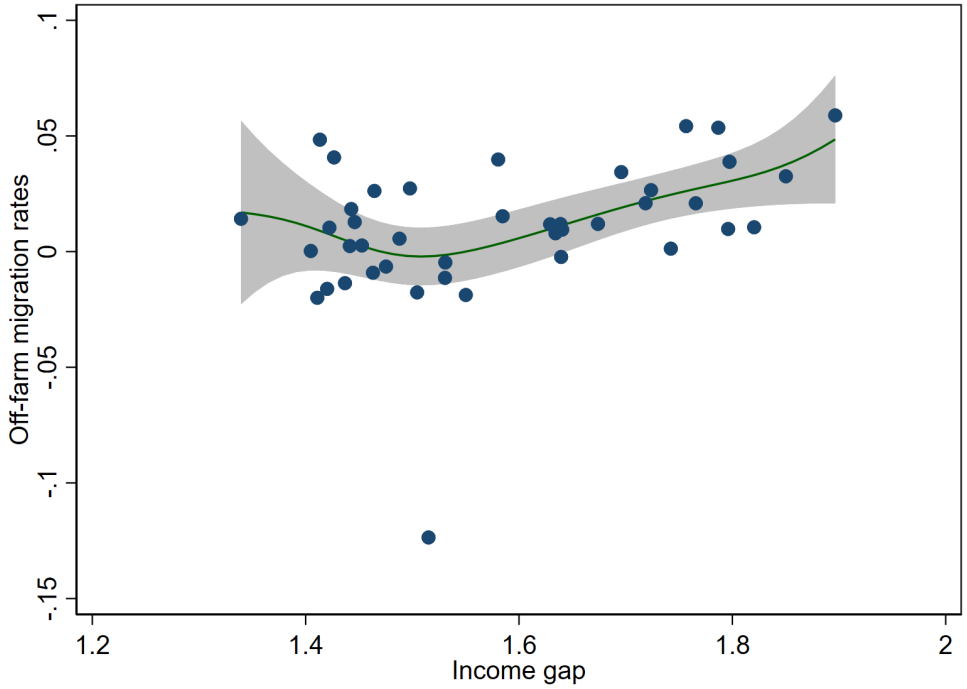

2.7.4 Remarks

Our research is inevitably faced with certain limitations, particularly about the data used for our analysis. One of the key limitations is the lack of data related to the hukou system in our analysis. The hukou system is a household registration policy in China that serves as a form of internal passport system, controlling population movement particularly from rural to urban areas. This system has a significant impact on migration patterns, as it can pose a significant barrier to migration.

Due to the potential influences, the absence of hukou-related data in our analysis might limit the accuracy of our findings. Our nonlinear relationship between migration rate and income gap might be influenced by the hukou system, and excluding this factor could potentially bias our results.

Secondly, the classification of employment sectors and geographic regions in the Chinese data presents challenges that need to be acknowledged. The distinctions between farm and non-farm sectors, and between rural and urban areas, are not always clear-cut in the available data. Often, these categories are mixed or combined, which can lead to potential misinterpretations. For instance, there are instances of rural residents engaged in non-farm sector jobs and urban dwellers working in the farm sector. These crossover cases are not always adequately accounted for or distinguished in the data, posing a potential limitation to the precision of our analysis.

The intricacies of labor dynamics and regional variances in China are complex and multilayered, making it challenging to categorize individuals strictly into 'farm', 'non-farm', 'rural', or 'urban.'

Furthermore, the aforementioned hukou system in China can complicate matters. It is possible that a worker who officially holds a rural hukou can be employed in a non-agricultural

job in an urban area, or vice versa. This complexity may not be entirely reflected in the available data, leading to potential discrepancies in our analysis.

Therefore, future research should aim to incorporate hukou-related data into the analysis to provide a more comprehensive understanding of the dynamics of migration in China. This could involve obtaining data on individuals' hukou status, as well as the policies and practices that affect the ability to change hukou status. Additionally, qualitative research can be conducted to explore individuals' experiences and perceptions of the hukou system, and how it influences their decision to migrate.

2.8 Conclusion

This study attempts to investigate the relationship between out-farm migration and income gap in China, with a particular focus on the potential nonlinearities in the relationship. Drawing on the conventional linear concept as a starting point, the study employed a semi-parametric approach to provide a more flexible analysis of the relationship between these two variables.

The results of the OLS model delivered useful insights into the factors influencing migration rates, however, the semiparametric model added another layer of complexity, capturing the nonlinear relationship between the income gap and migration rate that the OLS model failed to represent.

While our specification tests suggest a polynomial relationship, our main findings of the semiparametric analysis reveal a subtle nonlinear relationship between out-farm migration and the income gap in China, which deviates from the pattern proposed by previous studies (Barkley, 1990, Larson and Mundlak, 1997, Mundlak, 1978). As the income gap increases,

one might expect migration rates to rise as well. However, in a non-monotonic relationship in our Chinese case, migration rates do not simply increase as the income gap widens. Rather, there is a level of income disparity at which migration rates begin to decrease. This may be attributed to factors such as improved conditions in rural areas, the presence of economic opportunities outside of urban environments, or even the challenges and costs associated with migration.

The findings of this research align with our initial theoretical hypothesis - the real option theory, which postulates that individuals often delay their decision to migrate. However, it is important to note that it would be an oversimplification to assume that this theory is universally applicable to all countries or is exclusively relevant to developing nations. We initially hypothesized that developing countries, with their extensive opportunities for migration from agricultural to non-agricultural sectors, might exhibit a more pronounced manifestation of this theory. Yet, the complexities we encountered during our research suggest a more complicated picture.

Defining a country as 'developed' or 'developing' is not easy, and these terms may not fully capture the complex dynamics of inter-sectoral migration and income gaps. Instead, we propose focusing on the developmental stage or urbanization level of a country for future research, as these may provide a more accurate reflection of migration dynamics. For instance, even in developed countries, during their transitional period from agricultural to non-agricultural sectors (i.e., the 1960s), might demonstrate a more pronounced nonlinear pattern of migration. The real options theory offers valuable insight and was well explained in our study, but future research will benefit from a more targeted focus on specific developmental stages, and levels of urbanization for other cases.

Furthermore, the limitations encounter in this study, such as the lack of relevant data (i.e., hukou), underline the need for more detailed investigations. Future research could benefit

from employing a more localized approach, perhaps studying specific regions within a country to get more specific insights into the dynamics of inter-sectoral migration. These local-based or more detailed studies would help refine our understanding of the relationship between income gaps and migration decisions in the case of China.

Chapter 3

Enhanced Salience of Nonlinear Pricing and Energy Conservation

3.1 Abstract

Demand management has become an important strategy for utilities to address electricity production shortfalls and intermittent issues resulting from the transition to renewable energy. Many utilities apply increasing block tariffs to prevent pricing low-income households out of basic electricity access, while simultaneously discouraging wasteful overconsumption by high-income households. However, consumers' limited understanding of or attention to such complex pricing systems can result in private and socially sub-optimal behavior, rendering nonlinear pricing ineffective for energy conservation. This project aims to determine whether enhancing the understanding of energy use and nonlinear electricity pricing can help households respond to marginal pricing, thereby increasing energy conservation. We are conducting a large-scale experiment covering 30,000 users of a recently launched mobile app from a state-owned electric utility company in Vietnam. One treatment group receives a real-time app display of estimated daily marginal prices, while the other views their total estimated bills to date. All treatments last for a minimum of six months. To assess and compare the persistent effects of providing high-frequency nonlinear price and total cost information on consumer behavior, we will gather electricity billing data from the participants

for one year before and after the experiment.

3.2 Introduction

The energy sector in low- and middle-income countries (LMICs) faces two pressing challenges: accommodating the increasing energy demand for economic growth and limiting carbon emissions from primarily coal-based electricity generation. To curb air pollution levels and enhance electricity reliability, utilities have been prompting households and businesses to lower their electricity usage. Demand management has progressively become a crucial strategy for utilities to address shortcomings in electricity production and intermittency issues from the transition to renewable energy. Many utilities employ nonlinear pricing or escalating block tariffs to prevent pricing low-income households out of basic electricity access while simultaneously discouraging wasteful overconsumption by high-income households. However, growing evidence suggests that households respond to average prices rather than marginal prices due to their misunderstanding of or lack of attention to complex nonlinear pricing (Ito, 2014). In some cases, households may mistakenly perceive the marginal price as the price they must pay for all consumption units, resulting in underconsumption and significant welfare losses (Shaffer, 2020). Therefore, the energy conservation benefits of nonlinear pricing are not well-realized. Consequently, the design of nonlinear pricing systems and efforts to enhance consumers' comprehension of such pricing structures are crucial in realizing their potential benefits.

In this project, we have partnered with EVN Hanoi, a state-owned electric utility company in Vietnam that is the sole provider of electricity in the city of Hanoi, the second-largest city in Vietnam. We leverage its newly launched app for a large-scale randomized experiment. This experiment provides real-time price and total cost information via the app's display

to a sample of residential customers in Hanoi. We aim to study whether an enhanced understanding of energy use and nonlinear electricity pricing can prompt households to respond to marginal pricing and consequently save more electricity. Specifically, we examine three questions:

1. Does frequent information about electricity prices help households respond more effectively to marginal prices and alter their electricity consumption?

2. Which type of information has a stronger impact on consumer behavior - marginal electricity prices or total electricity bills?

3. How do households' responses to daily information about prices and bills evolve over time?

The project contributes to the existing literature in several ways. First, inattention and imperfect information are acknowledged as contributors to excessive resource usage, and many studies suggest that nudges – for example. regular updates on pricing and usage – could foster energy conservation (Allcott and Mullainathan, 2010, Gilbert and Graff Zivin, 2014, Tiefenbeck et al., 2018). Our study enhances this concept by integrating these nudges into a nonlinear pricing context, to discern whether frequent updates on electricity usage and its corresponding nonlinear pricing can induce a more effective response to marginal prices, thereby reducing overall electricity consumption. Stojanovski et al. (2020) align with our study's goal, demonstrating that a one-time, in-person information provision can motivate households to decrease their electricity usage under a nonlinear pricing schedule. Unlike their approach, we provide households with frequent updates on electricity usage and prices over a prolonged period. Conversely, Jessoe and Rapson (2014) found that price salience did not affect electricity conservation, even with regular information updates. Our study delves

deeper by offering information about a nonlinear pricing structure and probing how such detailed pricing salience impacts customers' energy consumption behaviors.

Second, this project involves the use of digital tools, such as a mobile app and smart meter data, to provide real-time energy, price, and billing information. The digitalization of our everyday lives simplifies access to information and can help individuals overcome bias and align their actions more closely with their preferences (Tiefenbeck et al., 2018). Our intervention employs this technology, addressing the information cost issues associated with nonlinear pricing.[1] With the recently launched app, communication with customers has become much easier and cheaper. Moreover, many customers now have smart meters, enabling utility companies to share frequent consumption data with them. This data provides immediate feedback that assists customers in their conservation efforts. Therefore, this study benefits from the availability and practicality of digital data, contributing to a more comprehensive examination of both the intervention and energy consumption patterns.

Third, this project examines the effects of price salience over varied durations. Gilbert and Graff Zivin (2014) discovered that consumers are most responsive to electricity prices immediately after receiving their bills. However, this responsiveness diminishes over time as the information becomes less prominent. This 'Recency Effect' of information provision is also highlighted by Allcott and Rogers (2014). Therefore, it is crucial to evaluate both the short-term and long-term impacts of interventions, as well as how these impacts differ over diverse durations. Assuming the price is dependent on the total electricity use in a billing cycle and not on the day of use, we are keen to observe *i*) whether households "temporarily" reduce consumption during the last few days of the billing cycle when the in-app displayed electricity price is at its peak and *ii*) whether households balance their consumption over a

[1] Ito (2014) says that the information cost of understanding the marginal price of electricity is likely to be substantial in terms of *i*) monthly utility bills do not provide relevant information on the nonlinear structure of the pricing, and *ii*) difficulties in monitoring cumulative consumption during a billing month.

billing cycle after "permanently" comprehending the price structure.

Fourth, further research into energy consumption behavior, especially in the context of developing countries, is urgently needed to address the dual issues of providing sufficient energy to promote economic development while limiting emissions from the energy sector. This is especially true for research employing technology such as smart meters almost exclusively focused on high-income countries with little understanding of such information and technology in low- and middle-income countries where these investments are nonetheless being pursued. Our project seeks to bridge this gap by centering our study on Hanoi, a city experiencing rapid economic growth, rising middle-class prosperity, and an increasing demand for electricity-intensive appliances. The city's issues with high per capita electricity consumption and energy constraints, resulting in frequent blackouts, make it representative of urban areas in developing countries expected to experience a surge in energy use. With Hanoi's context embodying the energy issues faced by many developing countries, the insights gained from this study have the potential to inform and shape sustainable growth policies beyond Vietnam.

EVN Hanoi serves the entire city of Hanoi, encompassing 2.8 million customers, the majority of whom are equipped with smart meters that record and transmit consumption data to the company's database daily. Among the 2.8 million smart meter users, 30% have installed the app, offering us a substantial sample size. We randomly select a sample of 45,000 households and split them equally into one control group and two treatment groups. Every household, including those in the control group, receives information about their past and present daily electricity usage via the electric utility company's app. The first treatment group is given daily updates about the estimated electricity prices for their last unit of consumption in each billing cycle, through the app's interactive display. The second treatment group receives information about their projected up-to-date electricity bill, also through the app's interactive

display. By randomizing the groups, we are able to estimate the average treatment effect of an enhanced understanding of nonlinear pricing on consumers' energy consumption.

The remainder of this paper is structured as follows. Section 3.4 provides background information about the research setting; Section 3.5 outlines our experimental design. In Section 3.6, we detail the data collection process. Our empirical analysis and results are presented in Section 3.7, and finally, concluding remarks in Section 3.8.

3.3 Conceptual Model

We outline a simple of consumer demand to analyze the influence of our interventions on the electricity consumption of households. Household i gets utility from electricity services s_i (with unit price p_s) and from a numeraire consumption good x_i. We assume that households have quasilinear utility:

$$u(s_i, x_i) = v_i(s_i) + x_i, \tag{3.1}$$

where v_i is an increasing and strictly concave function ($v' > 0$ and $v'' < 0$).[2] The household faces a budget constraint, $I_i = s_i p_s + x_i$, where I_i is the full income of household i. Household i maximizes utility in Equation (3.1) subject to the budget constraint that the demand for energy services is derived from the equation: $v'(s_i) = p_s$, so the demand for electricity services is $s_i = (v')^{-1}(p_s)$.

Electricity services cannot be purchased directly but are produced in a household production framework using electricity e_i according to the production function $s_i = \eta e_i$, where η is a parameter that converts units of electricity to units of energy services or η measures

[2]Quasilinear utility function is a standard use in the literature (e.g., Ito (2014)).

the energy efficiency level.[3] Given the energy efficiency level η, the production requires $s_i p_s - e_i p_e = 0$ where p_e is the unit price of electricity. Thus, the price of services s_i is $p_s = p_e/\eta$. Demand for electricity is thus:

$$e_i = (v')^{-1}\left(\frac{p_e}{\eta}\right). \tag{3.2}$$

The electricity demand function is decreasing in p_e as v' is a decreasing function (because v is strictly concave).

We denote δ_{p_e} is the consumer's misperception about the electricity price p_e, which measures a bias from the actual price $\tilde{p}_e$. Thus, we have $\tilde{p}_e + \delta_{p_e} = p_e$. If households overestimate the electricity price then $\delta_{p_e} > 0$. If households underestimate the electricity price then $\delta_{p_e} < 0$.

Without misperception ($\delta_{p_e} = 0$), the optimal energy use is

$$e_i^* = (v')^{-1}\left(\frac{\tilde{p}_e}{\eta}\right). \tag{3.3}$$

If $\delta_{p_e} > 0$, then household under-consumes electricity $e_i < e_i^*$ as:

$$e_i = (v')^{-1}\left(\frac{\tilde{p}_e + \delta_{p_e}}{\eta}\right) < (v')^{-1}\left(\frac{\tilde{p}_e}{\eta}\right) = e_i^*. \tag{3.4}$$

If $\delta_{p_e} < 0$, then household overconsumes electricity $e_i > e_i^*$ as:

$$e_i = (v')^{-1}\left(\frac{\tilde{p}_e + \delta_{p_e}}{\eta}\right) > (v')^{-1}\left(\frac{\tilde{p}_e}{\eta}\right) = e_i^*. \tag{3.5}$$

Both under and over-consumption of electricity can lead to welfare losses. Our key intervention is to enhance the salience of the electricity price and thus to reduce $|\delta_{p_e}|$ to 0. Thus,

[3]For example, electricity services are measured in electric vehicle miles per travelled, η is the efficiency (miles per kWH), and electricity is measured in kWh.

if the household's existing $\delta_{p_e} < 0$, then enhancing the salience of the electricity price will lead to reduction in energy use. However, if the household's existing $\delta_{p_e} > 0$, then enhancing the salience of the electricity price will lead to increases in energy use. Thus, the impacts of our interventions on energy use are ambiguous and require empirical investigation. Additionally, our basic static model lacks the capacity to depict the dynamics of improving the understanding of the electricity pricing. As a result, we will empirically explore both the impacts of our interventions and their dynamics in the subsequent sections.[4]

3.4 Research Setting

3.4.1 Background of Vietnam's Energy Use

Vietnam is experiencing a surge in energy consumption due to various influential factors. The country's significant economic growth has led to an increase in household incomes, which in turn is contributing to the emergence of a middle class with changing lifestyle preferences. This economic shift has led to a marked increase in the ownership and use of electronic devices, including computers, smartphones, and televisions. The hot and humid weather that is characteristic of Vietnam further drives the demand for cooling appliances such as air conditioners, particularly during the wet seasons. The humid conditions also lead to the use of dehumidifiers and other appliances aimed at improving indoor comfort, further adding to energy consumption. As energy consumption continues to rise, Vietnam faces the twin challenges of meeting the growing energy demands and coping with energy constraints, which manifest into rolling blackouts and intermittent electricity access, a scenario common

[4]In our upcoming endeavors, we aim to develop a dynamic and structural model. We intend to leverage experimental data to back out misconception parameters, which may vary across diverse customer types. Understanding these parameters will enhance our ability to clarify empirical results and assess the welfare implications of our interventions more effectively.

in many developing countries.

Vietnam has implemented a nonlinear pricing structure that comprises six price tiers determined by consumption quantity as shown in Table 3.1. This pricing structure is universal across the country and is operated by the state-owned Vietnam Electricity (EVN) group and its subsidiaries. Its purpose is to incentivize consumers to manage their electricity use wisely, thereby allowing lower-income households to access electricity at more affordable rates while placing the burden of higher prices on excessive consumption. This approach has the potential to be an effective instrument in addressing the aforementioned challenges in terms of assigning higher prices to higher consumption levels. However, understanding this structure can be challenging for customers, as noted by Ito (2014). EVN frequently faces customer complaints since they do not understand why their bills are so high. Given the difficulty and lengthy process of changing tariffs politically, improving customer understanding of multi-tiered pricing systems becomes crucial.

Our sample population will be households in Hanoi, the capital and second-largest city of Vietnam. With its expanding economy, a diverse population ranging from low-income to wealthy individuals, and energy challenges that reflect broader national issues, Hanoi serves as a relevant and informative representative sample.[5] A subsidiary of the EVN, EVN Hanoi, serves the city, providing approximately 2.8 million customers. Notably, 90% of these customers are equipped with smart meters, which send consumption data to the company's database daily, and 30% have adopted the utility's app, facilitating direct communication between the company and its users. This tech-inclined segment who already installed the app forms our study's sample.

Given these attributes, our dataset from Hanoi is indicative of urban areas where energy

[5]The trends in energy consumption put the city at the top of the nation in terms of electricity consumption per person.

consumption is expected to address the grid balancing problem, reduce power outages, and achieve a reduction in emissions. Thus, this project will address the sustainable growth policy challenges faced by Vietnam and other developing countries that might follow suit.

Table 3.1: Nonlinear Pricing Structure Implemented by EVN Hanoi, Vietnam

Tier	**Electricity usage (kWh)**	**Marginal electricity cost (Dong/kWh)**
1	0-50	1,728
2	51-100	1,786
3	101-200	2,074
4	201-300	2,612
5	301-400	2,919
6	401-	3,015

3.4.2 Sampling and Randomization Process

We obtain monthly consumption data for 700,000 households from the EVN Hanoi utility company, covering the period from August 2022 to July 2023. These households represent all app users of EVN Hanoi. To ensure data quality, we exclude households with more than 15% missing daily data from May 2023 to July 2023.[6] We also exclude households that had participated in other experiments.

Our study focuses on evaluating the impact of marginal price salience on energy consumption. In our experimental context, the highest price tier applies when monthly energy use exceeds 400 kWh. Conversely, the lowest price tier is applicable when consumption is at

[6]The absence of daily data is primarily due to temporary technical issues, which include meter failures, errors during data transfer, storage, and backups. This could also be attributed to households discontinuing services, taking long-term vacations, or the fact that the homes are rental or vacation units.

or below 50 kWh. As such, households consuming significantly more than 400 kWh/month or significantly less than 50 kWh/month will not be affected by our marginal price intervention. To enhance the statistical power of our study, we limit our sample to households with monthly energy usage ranging from 10 to 1000 kWh. Please note that less than 1% of all 700,000 app users consume less than 10 kWh monthly, and only 2% consume 1,000 kWh or more. We have also excluded households with missing data for any month over a twelve-month span, reducing our dataset to about 400,000 households. As illustrated in Figure 3.1, a comparison of the histograms of all 700,000 app users to those of the 400,000 users after exclusion shows a strong similarity. This suggests that, despite the removals, the data sample remains representative.

Figure 3.1: Distribution of Monthly Electricity Use among App Users

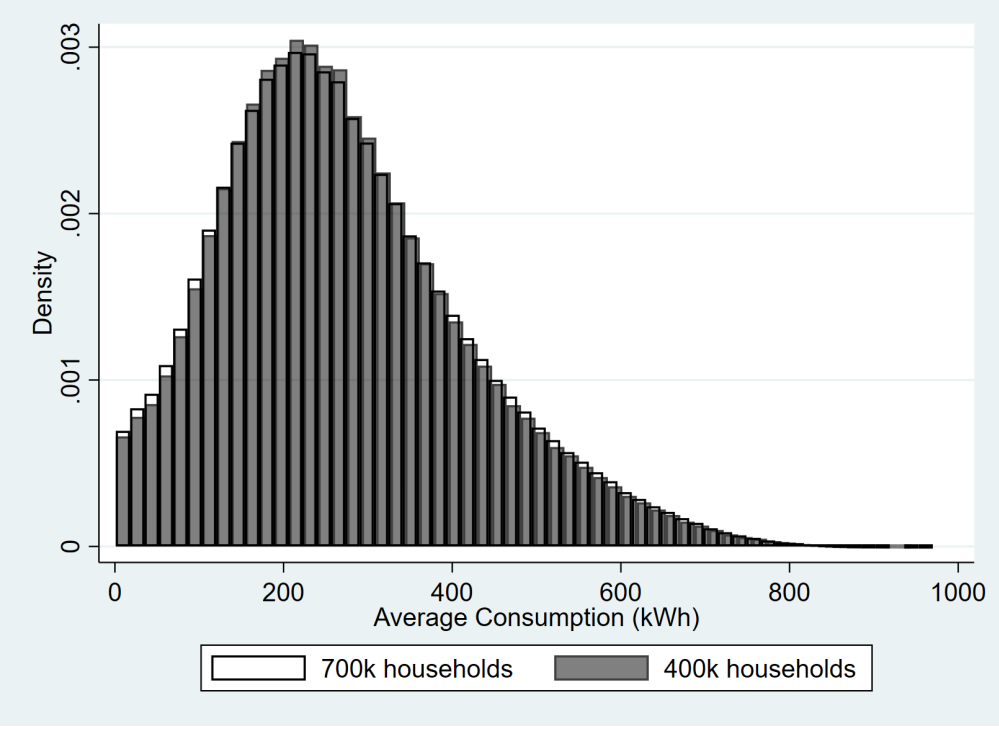

To finalize our sample, we randomly select 30,000 households in the initial pool of 400,000. To check if the mean and distribution of monthly electricity use among app users are consistent

between the final sample of 30,000 and the rest of the original 400,000 households, we perform a balance test. This test involves conducting an OLS regression, comparing the 30,000 selected households against the remaining 370,000. The dependent variables in this analysis are the average monthly electricity consumption for each household from August 2022 to July 2023. Figure 3.2 illustrates that both the 370,000 unselected and the 30,000 selected households exhibit nearly identical average monthly consumption over the analyzed period. Furthermore, Table 3.2 presents results indicating that the distribution of these two groups is similar, suggesting no significant differences between the selected and unselected samples. In terms of potential attrition, we anticipate minimal withdrawal given that there is no cost for households to remain in the study.

Figure 3.2: Mean Comparisons between the rest of the sample (370K) and 30K

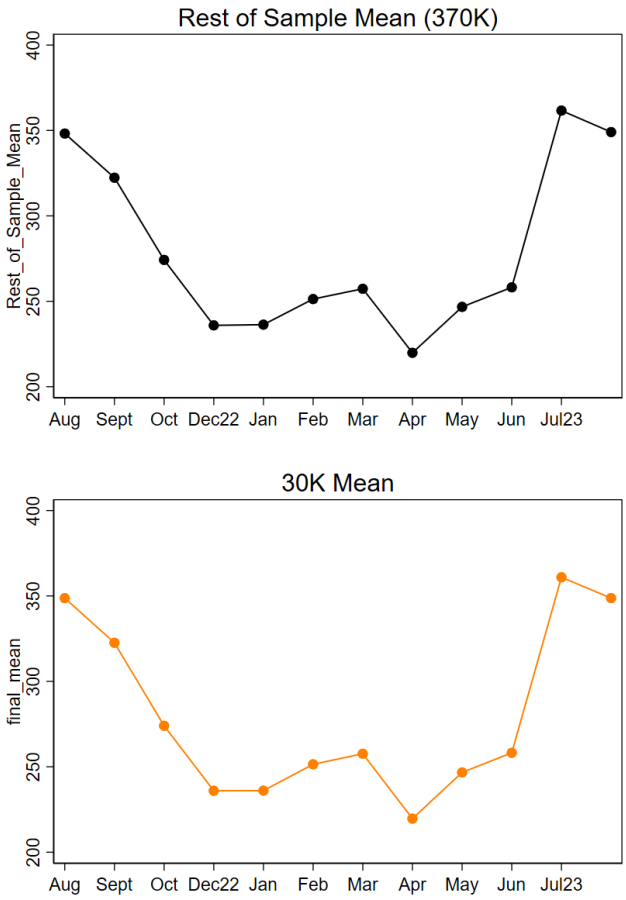

Note: Comparison of average monthly energy use (August 2022 - July 2023) between unselected 370K and selected 30K households from the original 400K sample after the randomizing process. Coefficients and P-values suggest no significant differences between the groups.

3.5 Experimental Design

The experiment is conducted on a subset of EVN Hanoi's customers, who have installed the utility's app. The app displays each household's daily consumption, and it compares customers' current and past consumption levels. We collaborate with I-Com, EVN Hanoi's app developer, to develop functionality on the mobile app to assign treatments to different

Table 3.2: Kolmogorov-Smirnov Tests between the rest of the sample (370K) and 30K

Month	370K vs 30K
Aug22	.708
Sep22	.242
Oct22	.747
Nov22	.553
Dec22	.944
Jan23	.916
Feb23	.910
Mar23	.650
Apr23	.953
May23	.728
Jun23	.461
Jul23	.515
Average	.868

Note: An observation is a household. All variables are defined at the household level. The second column reports the combined corrected p-values of a Kolmogorov-Smirnov equality-of-distributions test.

groups and to provide households with our designed and customized messages and information. Specifically, we take the universe of households who have installed the app as our population interest and randomly select and assign 30,000 households into a control and two treatment groups. The first and second treatment groups will receive real-time pricing and real-time billing information, respectively. The intervention strategies applied to each group are summarized in Table 3.3.

Control Group: This group is not be provided with any interactive display for daily electricity prices and up-to-date bills. Figure 3.3.(a) shows the app interface of the control

group.

Real-time Price Treatment: During the six-month treatment period from September 30, 2023, to February 28, 2024, households receive information regarding their estimated electricity prices for each unit of consumption in their current billing cycle via the app's interactive display. However, this study focuses on a shorter timeframe of three weeks, specifically from September 30 to October 22, 2023, due to constraints in data availability. Figure 3.3.(b) presents an example of the redesigned interactive display, visually representing the nonlinear pricing system, as compared to the original interface shown in Figure 3.3.(a). The graph in Figure 3.3.(b) highlights the primary intervention, allowing customers to visualize their current energy consumption along with the corresponding marginal price on a real-time six-tier graph. The design's objective is to help participants better understand this pricing structure.

Real-time Bill Treatment: In this treatment, similar to the price proficiency treatment, households will receive information regarding their estimated, up-to-date electricity usage and bill each day of their current billing cycle through the app's interactive display. The app's interactive display will remain active during the study period. Figure 3.3.(c) showcases the interactive interface for this treatment. Unlike treatment group 1, this treatment provides users with a simple daily display of total energy costs. Households in this group can view a semicircular gauge on their app, visually representing energy consumption. As energy usage escalates over time, the gauge fills up, transitioning its color from blue (signifying low consumption) to red (signifying high consumption).

Figure 3.3: App Interface Variations: Control Group vs. Treatment Groups

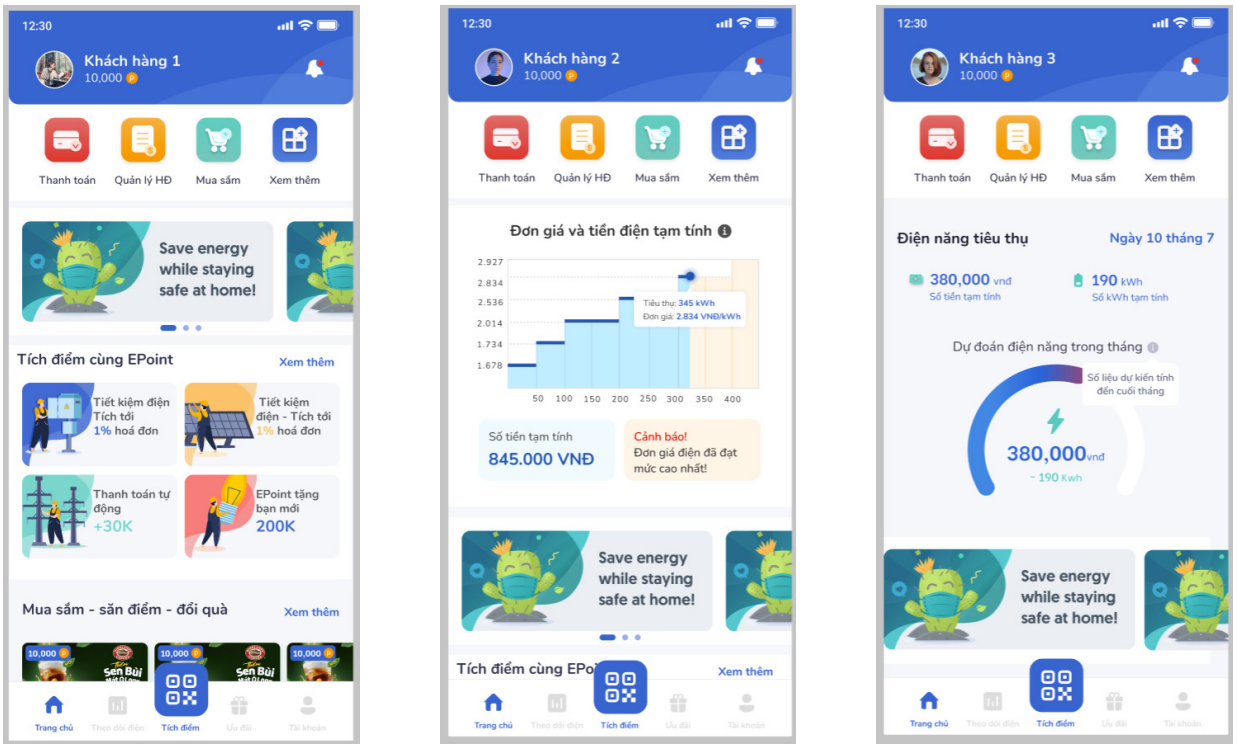

(a) Control (b) Treatment 1 (c) Treatment 2

Note: The control group (Panel (a)) does not have access to price or usage information on the app's main page. In contrast, the interface for treatment group 1 (Panel (b)) has been redesigned for a more visually intuitive understanding of the nonlinear pricing system. The app display for treatment group 2 (Panel (c)) offers users a vivid day-by-day visual of their accumulating energy costs.

3.6 Data

3.6.1 Variables

The data for this study is collected through a partnership with EVN Hanoi, a state-owned electric utility company in Vietnam, and I-Com, the app developer of EVN Hanoi. Specifically, we construct a household panel dataset by combining the following data sources to conduct our empirical analysis. Our data spans from October 01, 2022 to October 22, 2023

Table 3.3: Intervention Strategies by Groups

	Intervention	
Group	Participants are to be provided marginal price info on their app display for 6 months	Participants provided total cost info on their app display for 6 months
Control (10K households)	-	-
Treatment 1 (10K households)	X	-
Treatment 2 (10K households)	-	X

Note: An "X" in a cell indicates treatment received, and a dash ("-") indicates no treatment. In this experiment, all groups, including the control group receive information on daily energy consumption and bill-to-date in numeric form. While the intervention spans six months, our dataset currently includes only the three weeks following the intervention.

(the intervention start day is September 30, 2023). We use historical consumption data dating back 12 months prior to the intervention and continue to track electricity consumption. This monitoring period is anticipated to last approximately a year, this paper utilizes three weeks of available post-intervention data for the analysis.

The main outcome variable is electricity consumption at the household level. The smart meters provide an accurate measurement of energy usage and are capable of recording consumption data at regular intervals (e.g., monthly or daily). The dataset is obtained from EVN Hanoi, and includes historical data prior to the intervention, as well as consumption data throughout and after the treatment period.

Next, in collaboration with I-Com, we collect data on the usage of the mobile app by the participating households. This includes information on the frequency and duration of app usage, as well as the push notification receipt. This data allows us to analyze the engagement

levels of the households with the app and assess the effectiveness of the app in improving the households' understanding of their energy consumption and nonlinear pricing. We also obtain billing data from EVN Hanoi for the participating households, which includes information on the billed electricity consumption, the corresponding charges, and payment records for each household. Using this information, we can analyze the impacts of the intervention on on-time bill payments and delinquency rates, as well as validate consumption data obtained from smart meters.

3.6.2 Balance Check

We assess the balance between the treatment and control groups by examining household historical electricity consumption data. More precisely, we analyze the average daily electricity consumption for each month leading up to the intervention, spanning from from October 1, 2022, to September 29, 2023, as part of our balance checks.

We verify the balance of our assignment using the following regression:

$$y_i = \beta_0 + \sum_{k \in \{1,2\}} 1\{treatment_i = k\}\beta_k + \varepsilon_i, \tag{3.6}$$

where y_i denotes either the historical average daily electricity consumption of household i by month or its overall monthly consumption during the pre-intervention period. $treatment_i$ is an indicator variable that represents whether household i is assigned to treatment 1 or treatment 2, and ε_i signifies the error term. Note that this regression excludes indicators for the control group. In our balance analysis of this mean comparison, we report estimates for the coefficients of β_k, their standard errors, and the p-value from a joint test of the statistical significance of all coefficients on the treatment indicators.

Additionally, we apply the Kolmogorov-Smirnov test to continuous variables, such as electricity consumption, to compare distributions between each treatment group and the control group. In this test, the null hypothesis assumes that the distributions across the groups are identical. A detailed final balance check for these groups is presented in Table 3.4. This table reflects the findings from Equation (3.6). January 2023 stands out, with treatment 2 showing a coefficient of 0.189 and a conspicuously low P-value of 0.008, suggesting a substantial deviation in consumption compared to the control group. In contrast, higher P-values in other months indicate a more uniform consumption pattern among the control and the treatment groups during those periods. Figure 3.4 graphically represents the coefficients in Table 3.4, visualizing the similarities in groups. In this figure, the average daily consumption for each month is depicted for all three groups. Across all periods, the patterns for the three groups are nearly identical except for treatment 2 in January 2023. The Kolmogorov-Smirnov test in Table 3.5 results also indicate that the distribution of consumption data across the control and treatment groups is statistically similar for all considered months, which supports the mean comparison balance test findings. This test reveals that the similarity extends beyond the mean to higher moments of the distribution. For most months, the P-values are relatively high, suggesting no significant evidence to reject the null hypothesis of identical distributions between the groups. Still, in January 2023, the P-value is 0.07 for control versus treatment 2 shows relatively low but not sufficiently small to indicate a significant distributional differences. overall, we can say that the groups in the study display not just similar trends, but remarkably identical characteristics. Table 3.6 shows the distribution of consumption across tiers in our study sample.

Table 3.4: Balance analysis: mean comparisons across control and treatment groups

Month	Control Mean	Treatment1		Treatment2		F-test P-value
		Coeff.	P-value	Coeff.	P-value	
Oct22	7.952	-.065	.308	.014	.829	.429
Nov22	7.855	-.048	.442	.058	.371	.262
Dec22	8.146	-.067	.332	.044	.535	.289
Jan23	8.199	.12	.086	.189	.008	.026
Feb23	8.12	.086	.2	.073	.289	.388
Mar23	7.881	-.017	.786	.021	.742	.84
Apr23	8.38	.06	.372	.019	.785	.66
May23	10.889	.008	.922	.041	.622	.874
Jun23	11.71	-.132	.143	-.081	.379	.336
Jul23	13.435	-.119	.256	-.002	.985	.438
Aug23	11.038	-.063	.485	-.014	.877	.767
Sep23	10.416	-.113	.186	-.01	.905	.35

Note: An observation is a household. All variables are defined at the household level. Columns 3-6 report the coefficients and p-values from OLS regression on indicators for treatment groups 1 and 2. Column 7 reports the p-value from a joint test of the statistical significance of two indicators.

Figure 3.4: Mean Comparisons across control and treatment groups

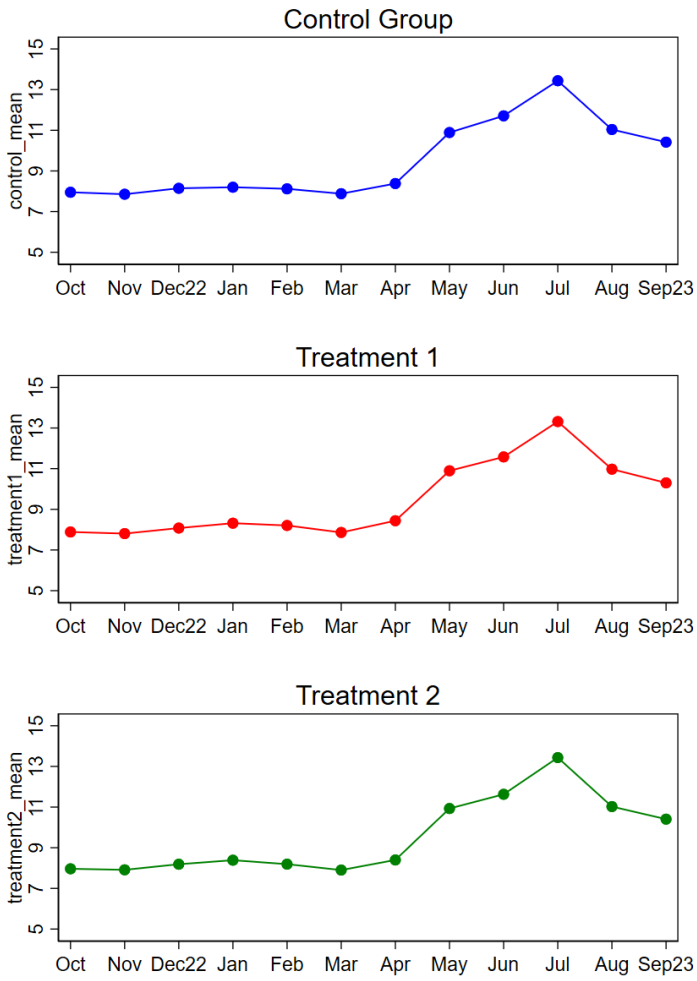

Note: Comparison of average daily energy use by month (October 2022 - September 2023) across control, treatment1 and treatment2. Coefficients and P-values suggest no significant differences between the groups except for treatment2 in January 2022.

Table 3.5: Kolmogorov-Smirnov Tests across control and treatment groups

Month	**Control vs Treatment1**	**Control vs Treatment2**
Oct22	.331	.635
Nov22	.435	.996
Dec22	.49	.591
Jan23	.204	.07
Feb23	.087	.156
Mar23	.75	.697
Apr23	.57	.956
May23	.691	.459
Jun23	.081	.529
Jul23	.375	.989
Aug23	.25	.672
Sep23	.14	.952

Note: An observation is a household. All variables are defined at the household level. The second column reports the combined corrected p-values of a Kolmogorov-Smirnov equality-of-distributions test.

Table 3.6: Distribution of Consumption in Tiers

Tier	Energy usage (kWh)	Control		Treatment1		Treatment2		Total	
		Obs.	%	Obs.	%	Obs.	%	Obs.	%
1	0-50	276	1.84	308	2.05	297	1.98	881	1.96
2	51-100	861	5.74	866	5.77	840	5.60	2,567	5.70
3	101-200	3,639	24.26	3,690	24.60	3,710	24.73	11,039	24.53
4	201-300	4,504	30.03	4,393	29.29	4,424	29.49	13,321	29.60
5	301-400	2,912	19.41	2,969	19.79	2,891	19.27	8,772	19.49
6	401-600	2,355	15.70	2,339	15.59	2,402	16.01	7,096	15.77
Above 6	601-801	437	2.91	419	2.79	425	2.83	1,281	2.85
	801-	16	0.11	16	0.11	11	0.07	43	0.10
Total		10,979	100.00	10,617	100.00	9,730	100.00	31,326	100

Note: This table displays the distribution of consumption across tiers after the randomization, showing the number of observations and their respective percentage share within each tier and group. The electricity utility in Vietnam (EVN) officially operates on a six-tier pricing system; any usage beyond this is categorized as "above 6".

3.7 Empirical Analysis

3.7.1 Treatment Effects

In this experiment, we estimate the average treatment effect of two distinct interventions. Treatment group 1 is exposed to the salience of nonlinear pricing, while treatment group 2 receives emphasis on the salience of billing. All of our participants have already downloaded the app and have been randomly allocated to one of the three groups. The previous balance test results confirm that this random allocation ensures that the groups are similar, except for the specific treatments they receive. To measure these treatment effects, we start by performing a regression analysis on cross-sectional data:

$$y_i = \beta_0 + \sum_{k \in \{1,2\}} 1\{treatment_i = k\}\beta_k + X_i'\gamma + \varepsilon_i, \tag{3.7}$$

where y_i is the average daily consumption or aggregate electricity consumption of household i during the study period. A dummy variable $treatment_i$ indicates the assignment of household i to treatment group k. X_i is a vector of control variables that captures other factors influencing electricity consumption (one specification includes no covariates, another specification includes the covariates of either the historical average of daily electricity consumption by month, or the total monthly electricity use during a year prior to the intervention) ; ε_i is the error term clustered at the household level.

In this specification, the parameters β_k capture the average effect of the respective treatments on electricity consumption. For example, β_1 would measure the average change in electricity consumption for households in treatment group 1 relative to the control group, while β_2 would capture the corresponding effect for treatment group 2.

We analyze electricity energy use in three forms: levels, a logarithm, and using inverse hy-

perbolic transformation. Using the logarithmic form is beneficial as it ensures the outcomes are non-negative, preventing nonsensical negative predictions in the context of energy use. Additionally, due to likely variations and skewness in household energy consumption, the log transformation helps normalize the distribution and lessen the impact of outliers. This also reduces the variance in outcomes, enhancing the statistical power of our analysis. It allows us to detect smaller effects and increases the chances of identifying significant relationships between independent variables and energy use. However, the log transformation poses a challenge with zero consumption, as the logarithm of zero is undefined. This could cause analytical complications for households with zero electricity consumption. To address this, we employ an inverse hyperbolic transformation for comparison, which shares the log transformation's properties but can handle zero values. Additionally, we examine energy consumption in its untransformed state, in raw levels.

The results from Equation (3.7) the are displayed in Table 3.7. The dependent variable is defined as either the average energy consumption during the experiment (from September 30, 2023, to October 22, 2023) or the total energy consumption in this period. Columns (1) to (4) in the table present the dependent variable in its original levels. Columns (5) to (8) show the dependent variable in the form of natural logarithms for both average and total consumption. Columns (9) to (12) apply the Inverse Hyperbolic Sine (IHS) transformation to address issues with zero values. All models include day fixed effects. The even-numbered columns incorporate controls for the household's average daily energy consumption in each month prior to the experiment (from October 1, 2022, to September 29, 2023), while the odd-numbered columns do not include these controls.

Table 3.7: Treatment effects: Cross-Sectional variation

	Dependent Variable											
	Level				Log				Inv. Hyperbolic Sine			
	Avg. Cons. (kWh)		Total Cons. (kWh)		Avg. Cons. (kWh)		Total Cons. (kWh)		Avg. Cons. (kWh)		Total Cons. (kWh)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Treatment1	-0.055	-0.092	-2.397	-3.438	-0.009	-0.021	-0.017	-0.029	-0.011	-0.021	-0.024	-0.039
	(0.079)	(0.074)	(1.789)	(1.699)	(0.011)	(0.012)	(0.012)	(0.012)	(0.011)	(0.011)	(0.015)	(0.017)
Treatment2	0.040	-0.058	-1.011	-3.468	0.008	-0.01	0.002	-0.018	0.002	-0.019	-0.014	-0.046
	(0.080)	(0.074)	(1.811)	(1.684)	(0.012)	(0.012)	(0.012)	(0.012)	(0.011)	(0.011)	(0.015)	(0.017)
Controls	N	Y	N	Y	N	Y	N	Y	N	Y	N	Y
N	720,061	447,074	720,498	447,442	708,860	440,036	708,860	440,036	720,061	447,074	720,498	447,442
Mean	9.349	9.470	210.491	212.743	2.015	2.033	5.125	5.14	2.696	2.712	5.725	5.737
Test	0.236	0.656	0.446	0.986	0.128	0.348	0.123	0.35	0.271	0.893	0.532	0.688

Note: Standard errors are clustered at the household level in parentheses. The dependent variables are the average daily consumption over the study period and their aggregate totals. Analyses are conducted in level, log, and IHS forms. Controls are based on historical average daily or monthly consumption pre-intervention (Oct 1, 2022 - Sep 29, 2023). "Mean" represents the sample's average dependent variable value. "Test" indicates the p-value from a two-sided F-test comparing coefficients of treatments 1 and 2. Analysis restricted to the study period.

Table 3.8: Treatment effects: panel

	Level				Log				Inverse Hyperbolic Sine			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Treatment1	-1.833	-1.908	-0.104	0.155	-0.199	-0.211	-0.009	0.002	-0.089	-0.089	-0.020	0.032
	(0.057)	(0.062)	(0.078)	(0.085)	(0.006)	(0.007)	(0.009)	(0.009)	(0.008)	(0.012)	(0.013)	(0.011)
Treatment2	-1.773	-1.673	-0.044	0.390	-0.184	-0.207	0.006	0.007	-0.087	-0.043	-0.018	0.078
	(0.055)	(0.049)	(0.079)	(0.068)	(0.007)	(0.007)	(0.009)	(0.009)	(0.009)	(0.010)	(0.012)	(0.013)
HH FE	N	Y	N	Y	N	Y	N	Y	N	Y	N	Y
Time FE	N	N	Y	Y	N	N	Y	Y	N	N	Y	Y
N	4,636,248	4,636,248	4,636,248	4,636,248	3,999,702	3,999,702	3,999,702	3,999,702	4,636,248	4,636,248	4,636,248	4,636,248
Mean	10.748	10.748	10.748	10.748	2.285	2.285	2.285	2.285	2.578	2.578	2.578	2.578
Test	0.446	0.009	0.446	0.009	0.102	0.668	0.101	0.590	0.873	0.000	0.873	0.000

Note: Standard errors are clustered at the household level in parentheses. Specifications vary from no fixed effects to including either household or date fixed effects, or both. Analyses are conducted in level, log, and IHS forms. "Mean" indicates the average value of the dependent variable in the sample. "Test" presents the p-value from a two-sided F-test, testing the null hypothesis that treatments 1 and 2 have identical coefficients. All periods are considered.

Overall, the findings indicate a general reduction in energy consumption, particularly in treatment group 1 compared to the treatment group 2. However, the coefficients are small and statistically insignificant. Our results indicate the estimated impact increase with the inclusion of covariates inferring the covariates have chosen appropriately. The table also reports the "Mean" of the dependent variable in the sample and the p-values from a two-sided F-test. The hypothesis tested is whether treatment group 1 and the group 2 have equal coefficients. The results from "Test" provides insufficient evidence to reject the null hypothesis that the intervention's effects are the same across groups except for Columns (10) and (12). These Columns are the specifications using Inverse Hyperbolic Sine while controlling for date fixed effects and for both of household and date fixed effects, respectively.

Table 3.8 presents an analysis of a similar approach but differs in its application to a panel data structure, including the periods before and after the intervention. The dataset spans from October 1, 2022, to October 22, 2023, offering an overview of the treatment's average effect over the entire duration covered. The dependent variable is now daily energy consumption (kWh), instead of using the average or total consumption observed over the study period. The model incorporates fixed effects for both households and time, ranging from models with no controls to those with two fixed effects, resulting in four distinct specifications for each analysis. To address potential correlations within the dataset, standard errors are clustered at the individual customer level. The study has 4,636,248 observations, revealing an average daily use of 10.748 (kWh) in the level specification, 2.285 in the log, and 2.578 in the IHS specification. For both treatment groups, the level specification generally indicates a decrease in daily energy use. However, this effect demonstrates variability upon the inclusion of household and time fixed effects in the model. Such variations suggest that energy use could either decrease or increase under different conditions, depending on the specific model specification. When examining the log and IHS forms, the impact of

the treatments appears more ambiguous, with numerous coefficients are close to zero. This indicates that each of the treatments' influence on the percentage change in energy use is relatively minimal than what is captured in the level specification.

There are several key considerations to account for the observed null or ambiguous results from the findings in 3.8. Firstly, in reference to the studies by Ito (2014) and Shaffer (2020), the ambiguity in the signs of effects may stem from the variability in individual responses to electricity bills or prices. This variability might lead to a null result, suggesting that either people do not respond at all, or they respond in opposite directions. Such divergent responses align with the findings in Ito (2014) and Shaffer (2020)'s research, which demonstrated that consumer behavior varies based on their perceptions and understanding of electricity pricing. Secondly, the issue of statistical power must be considered. Our current challenge is the limited length of the post-intervention period compared to the pre-intervention phase, which includes a full year of data, due to the ongoing stage of the trial. Hence, drawing a conclusion on the treatment effect at this stage might be premature. Finally, acknowledging the potential for dynamic and heterogeneous effects is essential. The effects of interventions can affect by the different conditions of energy use and the diverse household characteristics in the study. These heterogeneous effects will be explored in more detail in the following section.

To further our understanding of the treatment effect over time, we adopt a panel analysis approach. The following regression equation is used for this purpose:

$$y_{i,t} = \beta_0 + \sum_{k\in\{1,2\}} \sum_{\tau=0}^{t-1} 1\{treatment_i = k\} \cdot 1\{t = \tau\}\beta_{k,\tau} + X_i'\gamma + \varepsilon_{i,t}, \tag{3.8}$$

where $y_{i,t}$ represents the energy consumption of household i at time t. Initially, our analysis focuses solely on the period of the study. In this context, t could represent various time

frames like days, weeks, months, or even years since the study's start. Considering our dataset includes data for only 22 days post-intervention, we estimate the effects in terms of weeks.

The results derived from Equation (3.8) allow us to examine the treatment effect over time, utilizing the coefficients $\beta_{k,\tau}$ obtained from the regression analysis. Figure 3.5 graphically represents these results, plotting the treatment effects across the time frame of the study period, specifically weeks. In this analysis, the variable t, representing days after treatment, ranges from 1 to 22, More specifically, the days after treatment span September 30, 2023, to October 22, 2023.

The upper row of Figure 3.5 illustrates the variation in energy usage between treatment group 1 and the control group over this period. The lower row contrasts treatment group 2 with the control group. The first column of plots presents data in its original level form, the second in log, and the third uses IHS transformation. These graphical representations include both date fixed effects and household effects.

The results indicate that treatment group 1 experienced a decrease in energy consumption following the intervention compared to the control group. The magnitude of this decrease is greatest in the first week, with a lesser reduction observed in the second week, and nearly returning back to the original level in the third week. Treatment group 2 demonstrates an initial decrease in energy consumption in the first week following the app change. However, this is rapidly reversed, leading to an increase in energy consumption in the subsequent weeks 2 and 3, compared to the control group.

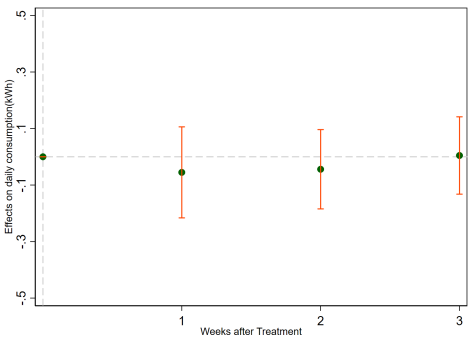

Treatment 1 vs Control (level)

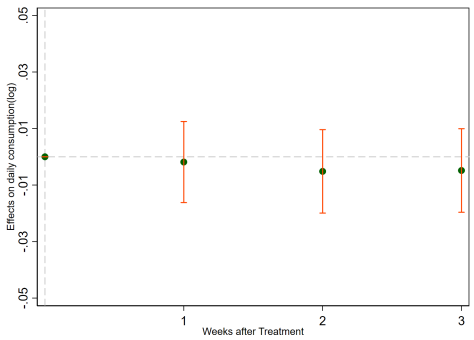

Treatment 1 vs Control (log)

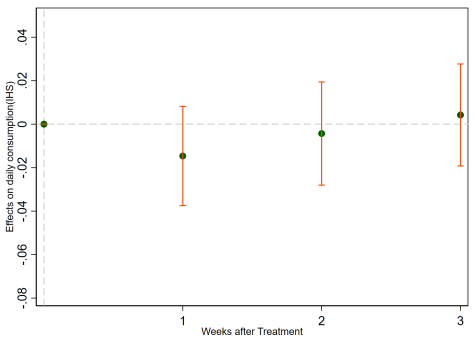

Treatment 1 vs Control (IHS)

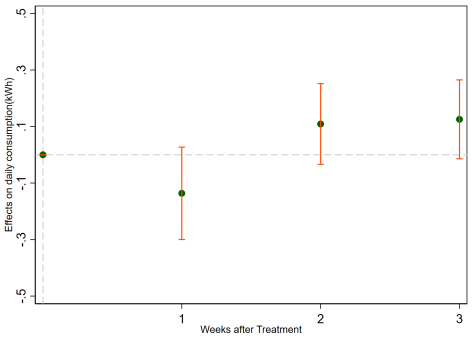

Treatment 2 vs Control (level)

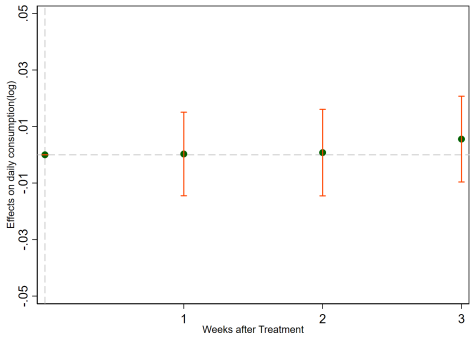

Treatment 2 vs Control (log)

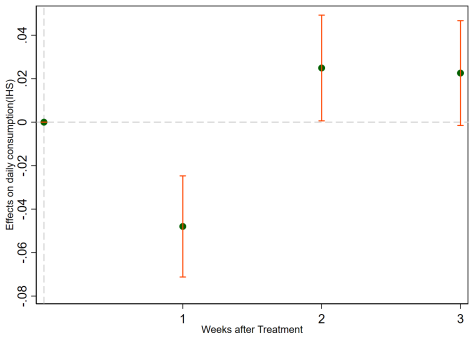

Treatment 2 vs Control (IHS)

Figure 3.5: Time Effect by Week (study period)

Note: Standard errors are clustered by household, with each observation representing a day for a specific household. Figures display the differential energy use (in kWh, log, and IHS) of treatment groups versus the control group over time, measured by indicators at the weeks after treatment. All models include date and household fixed effects.

3.7.2 Heterogeneous Treatment Effects

To capture these heterogeneous treatment effects, we modify our regression model by adding interaction terms for specific variables. The model is represented as:

$$y_i = \beta_0 + \sum_{k\in\{1,2\}} 1\{treatment_i = k\}b_{1,k} + \sum_{k\in\{1,2\}} 1\{treatment_i = k\}Z_i b_{2,k} + X_i'\gamma + \epsilon_i, \quad (3.9)$$

where y_i is the electricity consumption at the household level, and Z_i includes covariates such as the number of days between billing cycles. ε_i is the error term clustered at the household level. Here, the impact of treatment k on the treated, given the covariate Z, is $b_{1,k} + b_{2,k}Z$.

In the cross-sectional analysis, we introduce four dummy variables representing the weeks when each household's billing cycle begins. For example, a household with a billing start day of 5 is categorized in the first week, while one starting on 29 falls in the fourth week. This is important because the timing of the billing cycle relative to our intervention's start date (September 30) may influence the treatment's effectiveness. We investigate whether the treatment effect differs based on which week of the month a household's billing cycle starts. Next, in our panel approach, we consider the specific day of the billing cycle. This is to understand how consumption patterns might vary over the billing cycle and how these patterns interact with the treatment effects. Our analysis first examines the study period only and later extends to include both pre-intervention and post-intervention phases.

Further, we will add day of reminders, app usage frequency, and seasonal or holiday influences as well. These factors could affect household responsiveness to the intervention.

In the analysis of confined to the study period and employing a cross-sectional data, we first explore how the timing of a household's billing cycle start date within a month influences energy consumption. This aspect is important since the app's display change on September

30 may interact with where a household is in its billing cycle, potentially influencing their behavior. For instance, households at the start of their billing cycle in early September might see higher usage numbers on the intervention day, whereas those in the final week might see lower figures.

We include an interaction term in our model to test this hypothesis. The billing start day is categorized into four weekly segments: week 1 (1st to 7th of the month), week 2 (8th to 15th), week 3 (16th to 23rd), and week 4 (24th to end of the month). The regression results displayed in the Table 3.9, allow us to understand how the impact of these interventions varies across different weeks, with week 4 serving as the base category.

For treatment 1, the estimated average treatment effect indicates a negative change in electricity consumption. The coefficient for treatment 1, $\beta_{1,1}$, is the overall impact of this intervention, averaged across all weeks of the month. In the specified model, the coefficient for the interaction term $T1 * W1$, denoted as $\beta_{2,1}W1$, quantifies the deviation of treatment1's effect in week 1 from its average effect across all weeks, $\beta_{1,1}$. Therefore, to interpret the total effect of treatment 1 in week 1, it needs to be computed as the sum $\beta_{1,1} + \beta_{2,1}W1$. This sum reflects the cumulative impact of treatment 1 when billing commences in week 1, relative to the baseline of week 4.

In all specifications, the total effect of treatment 1 in billing week 1 is positive, while the effects in weeks 2 and 3 are negative, with week 3 exhibiting greater negative values. This indicates that, given our dependent variable is energy consumption and the goal is energy savings, the intervention is most effective in week 3, moderately effective in week 2, then the baseline week 4, and least effective in week 1, where it actually leads to an increase in energy consumption. Treatment 2 also shows a similar pattern, with smaller magnitudes. These results suggest that the timing of the intervention, specifically the billing week, influences its effectiveness in reducing energy consumption. However, the results are not statistically

significant.

In the next, we employ a panel data structure to examine the heterogeneous effects associated with different stages of the billing cycle. Specifically, the objective here is whether the impact of our experimental intervention vary depending on the number of days elapsed since the last billing payment. The timing within the billing cycle might be more relevant than calendar days for influencing user behavior in the sense that when users access the app, their energy usage or costs are displayed based on their unique billing cycle.

To facilitate this, we categorize the start day of each household's billing cycle. For instance, if a household's billing cycle starts in which day of the month, this day is coded as '1' in our billing cycle calculations. Our findings are visually represented in Figures 3.6 and 3.7. Figure 3.6 focuses solely on the study period, while Figure 3.7 considers data from both before and after the intervention.

Across both analyses and treatment groups, we observe a consistent pattern that households tend to exhibit more effective energy saving behavior during the mid-range of their billing cycle compared to either the beginning or the end of the cycle. This results infers that the timing within the billing cycle also may affect differently on how households respond to reducing energy consumption.

Table 3.9: Heterogenous Effects: Billing Start Timing

	Level		Log		Inverse Hyperbolic Sine	
	avg use	agg use	avg use	agg use	avg use	agg use
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment1	-2.064	-45.052	-0.282	-0.275	-0.251	-0.181
	(1.528)	(35.165)	(0.222)	(0.223)	(0.215)	(0.223)
T1xW1	2.331	50.605	0.314	0.296	0.309	0.265
	(1.531)	(35.238)	(0.222)	(0.223)	(0.215)	(0.223)
T1xW2	1.943	42.129	0.266	0.256	0.223	0.132
	(1.529)	(35.190)	(0.222)	(0.223)	(0.215)	(0.223)
T1xW3	1.440	30.923	0.176	0.166	0.167	0.104
	(1.537)	(35.369)	(0.223)	(0.224)	(0.216)	(0.225)
Treatment2	-1.043	-25.257	-0.104	-0.109	-0.079	-0.019
	(1.462)	(33.034)	(0.210)	(0.211)	(0.204)	(0.211)
T2xW1	1.701	38.933	0.207	0.209	0.185	0.151
	(1.466)	(33.124)	(0.211)	(0.211)	(0.204)	(0.211)
T2xW2	0.993	22.721	0.099	0.100	0.064	-0.017
	(1.463)	(33.062)	(0.211)	(0.211)	(0.204)	(0.211)
T2xW3	0.459	10.618	0.017	0.019	-0.010	-0.092
	(1.472)	(33.257)	(0.212)	(0.212)	(0.205)	(0.214)
N	720,061	720,498	708,860	708,860	720,061	720,498
Mean	9.349	210.491	2.015	5.125	2.696	5.725

Note: Standard errors are clustered at the household level. The dependent variables are the average daily consumption over the study period and their aggregate totals. A cross-sectional analysis assesses the heterogeneous effects based on billing timing. The main effects of treatments 1 and 2 are reported separately, with interaction coefficients also important for interpreting the overall results. Week 4 is the baseline.

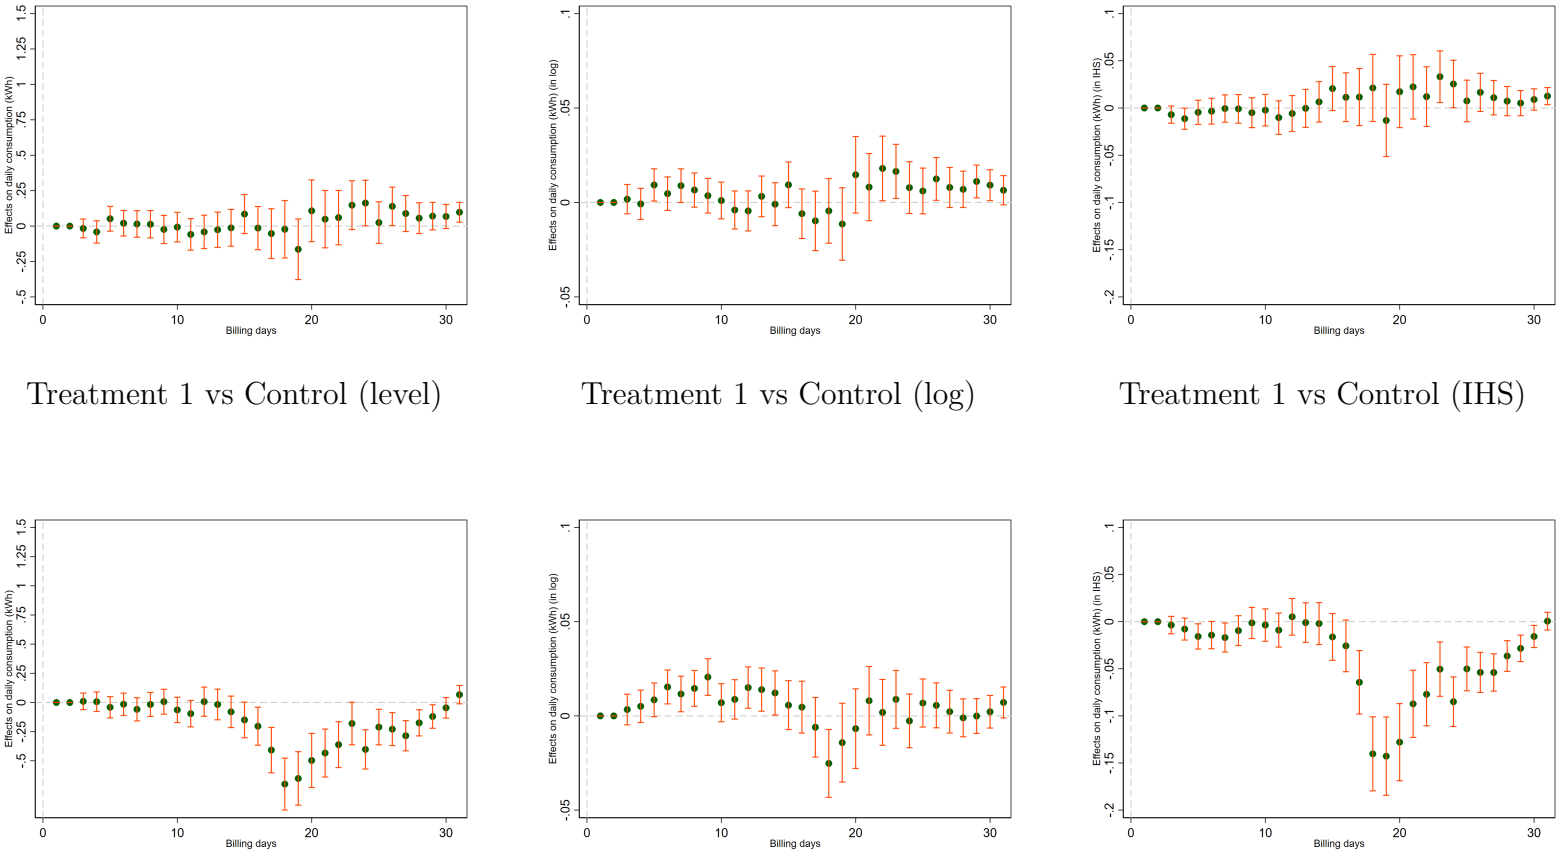

Treatment 1 vs Control (level) | Treatment 1 vs Control (log) | Treatment 1 vs Control (IHS)

Treatment 2 vs Control (level) | Treatment 2 vs Control (log) | Treatment 2 vs Control (IHS)

Figure 3.6: Heterogeneous Effect: Billing Cycle (study period only)

Note: Standard errors are clustered by household, with each observation representing a day for a specific household. Figures display the differential energy use (in kWh, log, and IHS) of treatment groups versus the control group across the billing cycle. All models include date and household fixed effects, covering study period only.

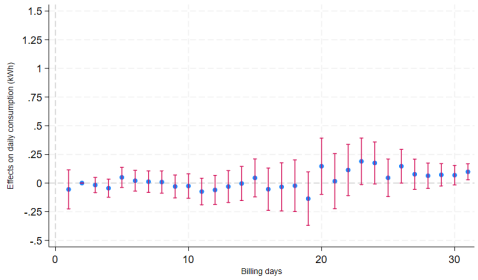

Treatment 1 vs Control (level)

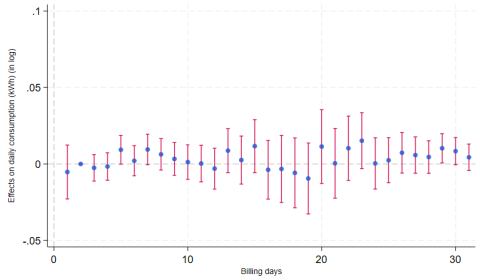

Treatment 1 vs Control (log)

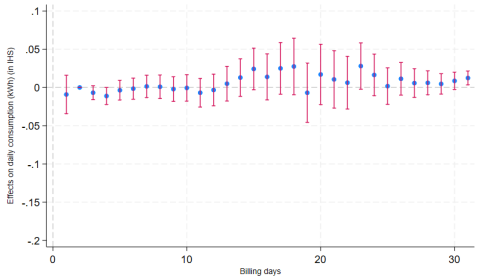

Treatment 1 vs Control (IHS)

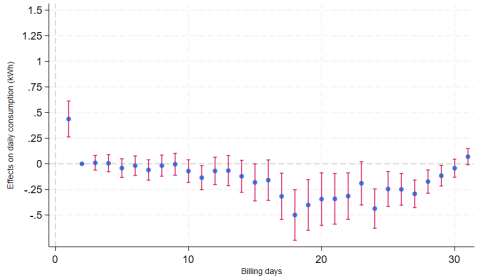

Treatment 2 vs Control (level)

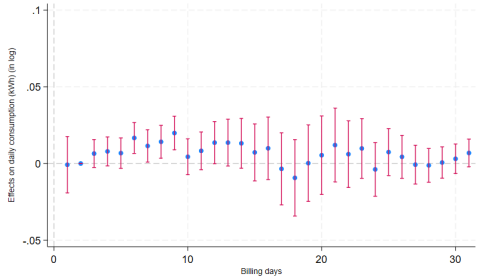

Treatment 2 vs Control (log)

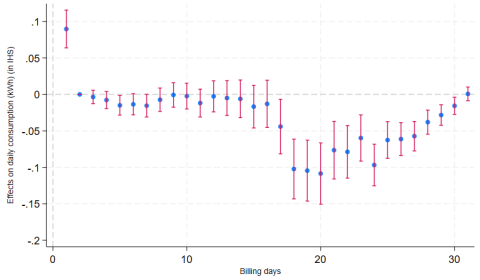

Treatment 2 vs Control (IHS)

Figure 3.7: Heterogeneous Effect: Billing Cycle (all periods)

Note: Standard errors are clustered by household, with each observation representing a day for a specific household. Figures display the differential energy use (in kWh, log, and IHS) of treatment groups versus the control group across the day of billing cycle. All models include date and household fixed effects, covering all periods.

3.8 Conclusion

In this study, we investigate the role of real-time price information in improving consumer responses to nonlinear pricing systems and encouraging efficient energy use. Our research focuses on three key questions: the effects of frequent electricity price updates, determining which type of information—marginal prices, total bills, or average prices—has the most significant impact on consumers, and how household responses evolve over time.

We partnered with EVN Hanoi, Vietnam's state-owned electric utility, to conduct a large-scale randomized controlled trial with 30,000 households in Hanoi. The study design includes a control group and two treatment groups. The first treatment group receives daily app updates on estimated prices for their most recent electricity use. The second group gets updates on their cumulative electricity bill through the same app. Our data, ranging from October 1, 2022, to October 2023, intervention starting from September 30, 2023, shows a noticeable yet not statistically significant link between reduced energy consumption and the treatments, especially treatment 1, the primary focus of our study. We also observe that the intervention's potential effectiveness is influenced by time-dependent factors and variables such as the billing cycle. However, the overall results do not show consistent significance across different model specifications, and the effect sizes are generally small.

The limited duration of the intervention period's data collection may be a factor in these inconclusive findings, as it hinders our ability to fully capture and understand the intervention's impact. The short duration of our study may not fully represent the long-term behaviors and responses of households to the intervention.

As we plan to collect more data over longer periods in future research, we expect to address these limitations and provide a more comprehensive analysis. An expanded dataset covering a longer post-intervention period will better equip us to robustly assess the intervention's

effectiveness and understand the factors influencing energy consumption behavior. While our current findings provide valuable initial insights, they also point to the importance of ongoing data collection and analysis for a more complete understanding of the intervention's impact.

Ultimately, our research seeks to contribute to the existing body of research by combining behavioral nudges with nonlinear pricing strategies and using digital tools for real-time price information. We also aim to examine the temporal effects of these interventions in a rapidly developing country with increasing energy demands. By giving households a clear picture of their consumption costs, we hope to encourage more sustainable energy usage patterns and inform policy development in the sector.

Bibliography

H Allcott and S Mullainathan. Energy. behavior and energy policy. *Science*, 327(5970): 1204–1205, 2010.

Hunt Allcott and Todd Rogers. The short-run and long-run effects of behavioral interventions: Experimental evidence from energy conservation. *American Economic Review*, 104 (10):3003–3037, 2014.

Chong-En Bai, Chang-Tai Hsieh, and Yingyi Qian. The return to capital in china. *Brookings Papers on Economic Activity*, 37(2):61–102, 2006.

Andrew P. Barkley. The determinants of the migration of labor out of agriculture in the United States, 1940–85. *American Journal of Agricultural Economics*, 72(3):567–573, 1990.

Robert J. Barro and Jong Wha Lee. A new data set of educational attainment in the world, 1950–2010. *Journal of Development Economics*, 104:184–198, 2013.

James Obadiah Bukenya, Peter V Schaeffer, and Tesfa G Gebremedhin. The effect of wage differentials and regional job growth on migration: A case of west virginia. *West Virginia University*, 2003.

Michael C Burda. *Migration and the option value of waiting*. Institute for International Economic Studies, Stockholm University, 1995.

Marshall Burke and Kyle Emerick. Adaptation to climate change: Evidence from us agriculture. *American Economic Journal: Economic Policy*, 8(3):106–140, 2016.

Marshall Burke and David Lobell. *Food Security and Adaptation to Climate Change: What Do We Know?*, pages 133–153. Springer Netherlands, Dordrecht, 2010.

Chuanbo Chen and C Cindy Fan. China's hukou puzzle: Why don't rural migrants want urban hukou? *China Review*, 16(3):9–39, 2016.

T Curtis and NG Halford. Food security: the challenge of increasing wheat yield and the importance of not compromising food safety. *Annals of Applied Biology*, 164(3):354–372, 2014.

Benjamin N. Dennis and Talan B. İşcan. Productivity growth and agricultural out-migration in the United States. *Structural Change and Economic Dynamics*, 18(1):52–74, 2007.

Avinash K. Dixit and Robert S. Pindyck. *Investment under Uncertainty*. Princeton University Press, Princeton, N.J, 1994.

Ben Gilbert and Joshua Graff Zivin. Dynamic salience with intermittent billing: Evidence from smart electricity meters. *Journal of Economic Behavior & Organization*, 107:176–190, 2014.

Pu Hao and Shuangshuang Tang. Floating or settling down: The effect of rural landholdings on the settlement intention of rural migrants in urban china. *Environment and Planning A*, 47(9):1979–1999, 2015.

Wolfgang Hardle and Enno Mammen. Comparing nonparametric versus parametric regression fits. *Annals of Statistics*, pages 1926–1947, 1993.

John R. Harris and Michael P. Todaro. Migration, unemployment and development: A two-sector analysis. *American Economic Review*, 60(1):126–142, 1970.

Carsten A. Holz. The quality of China's GDP statistics. *China Economic Review*, 30: 309–338, 2014.

IPCC. Impacts, adaptation, and vulnerability. *Climate Change: Synthesis Report*, 2014.

IPCC. The physical science basis. *Climate Change: Synthesis Report*, 2021.

Koichiro Ito. Do consumers respond to marginal or average price? evidence from nonlinear electricity pricing. *American Economic Review*, 104(2):537–563, 2014.

Katrina Jessoe and David Rapson. Knowledge is (less) power: Experimental evidence from residential energy use. *American Economic Review*, 104(4):1417–1438, 2014.

Kentaro Kawasaki and Shinsuke Uchida. Quality matters more than quantity: asymmetric temperature effects on crop yield and quality grade. *American Journal of Agricultural Economics*, 98(4):1195–1209, 2016.

David L Kelly, Charles D Kolstad, and Glenn T Mitchell. Adjustment costs from environmental change. *Journal of Environmental Economics and management*, 50(3):468–495, 2005.

Yoko Kusunose, Jairus J Rossi, David A Van Sanford, Phillip D Alderman, James A Anderson, Yuan Chai, Maria K Gerullis, SV Krishna Jagadish, Pierce A Paul, Jesse B Tack, et al. Sustaining productivity gains in the face of climate change: A research agenda for us wheat. *Global change biology*, 29(4):926–934, 2023.

Donald Larson and Yair Mundlak. On the intersectoral migration of agricultural labor. *Economic Development and Cultural Change*, 45(2):295–319, 1997.

David B Lobell and Gregory P Asner. Climate and management contributions to recent trends in U.S. agricultural yields. *Science*, 299(5609):1032–1032, 2003.

David B Lobell and Christopher B Field. Global scale climate–crop yield relationships and the impacts of recent warming. *Environmental Research Letters*, 2(1):014002, 2007.

David B Lobell, Marshall B Burke, Claudia Tebaldi, Michael D Mastrandrea, Walter P Falcon, and Rosamond L Naylor. Prioritizing climate change adaptation needs for food security in 2030. *Science*, 319(5863):607–610, 2008.

David B Lobell, Wolfram Schlenker, and Justin Costa-Roberts. Climate trends and global crop production since 1980. *Science*, 333(6042):616–620, 2011.

Nathaniel B Lyman, Krishna SV Jagadish, L Lanier Nalley, Bruce L Dixon, and Terry Siebenmorgen. Neglecting rice milling yield and quality underestimates economic losses from high-temperature stress. *PloS One*, 8(8):e72157, 2013.

Noah Miller, Jesse Tack, and Jason Bergtold. The impacts of warming temperatures on us sorghum yields and the potential for adaptation. *American Journal of Agricultural Economics*, 103(5):1742–1758, 2021.

Satoshi Morita, Hiroshi Wada, and Yuji Matsue. Countermeasures for heat damage in rice grain quality under climate change. *Plant Production Science*, 19(1):1–11, 2016.

Yair Mundlak. Occupational migration out of agriculture: A cross-country analysis. *Review of Economics and Statistics*, 60(3):392, 1978.

Elizbar A Nadaraya. On estimating regression. *Theory of Probability & Its Applications*, 9 (1):141–142, 1964.

Lawton Lanier Nalley, Andrew P Barkley, and Forrest G Chumley. The impact of the kansas wheat breeding program on wheat yields, 1911–2006. *Journal of Agricultural and Applied Economics*, 40(1379-2016-112782):913–925, 2008.

National Bureau of Statistics of China. China Statistical Yearbook 2021. http://www.stats.gov.cn/sj/ndsj/2021/indexeh.htm, 2021. Accessed: 2023-10-09.

NDSU Agriculture Communication. Grain yield not related to test weight. https://www.ag.ndsu.edu/news/newsreleases/2012/aug-20-2012/grain-yield-not-related-to-test-weight, 2012. Accessed: 2023-10-09.

Lia Nogueira, Joel Michalski, Thomas L Marsh, and Vicki McCracken. Welfare implications of washington wheat breeding programs. *Journal of Agricultural and Applied Economics*, 47(2):147–174, 2015.

Elizabeth Nolan and Paulo Santos. The contribution of genetic modification to changes in corn yield in the United States. *American Journal of Agricultural Economics*, 94(5): 1171–1188, 2012.

Gülcan Önel and Barry K. Goodwin. Real options approach to inter-sectoral migration of U.S. farm labor. *American Journal of Agricultural Economics*, 96(4):1198–1219, 2014.

Ariel Ortiz-Bobea and Richard E Just. Modeling the structure of adaptation in climate change impact assessment. *American Journal of Agricultural Economics*, 95(2):244–251, 2013.

Shaobing Peng, Jianliang Huang, John E Sheehy, Rebecca C Laza, Romeo M Visperas, Xuhua Zhong, Grace S Centeno, Gurdev S Khush, and Kenneth G Cassman. Rice yields decline with higher night temperature from global warming. *Proceedings of the National Academy of Sciences*, 101(27):9971–9975, 2004.

A Ford Ramsey, Jesse B Tack, and Maria Balota. Double or nothing: Impacts of warming on crop quantity, quality, and revenue. *Journal of Agricultural and Resource Economics*, 47(1):1–22, 2021.

A Ford Ramsey, Tadashi Sonoda, and Minkyong Ko. Intersectoral labor migration and agriculture in the United States and Japan. *Agricultural Economics*, 54(3):364–381, 2023.

Peter M Robinson. Root-N-consistent semiparametric regression. *Econometrica*, 56(4):931–954, 1988.

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

Blake Shaffer. Misunderstanding nonlinear prices: Evidence from a natural experiment on residential electricity demand. *American Economic Journal: Economic Policy*, 12(3): 433–461, 2020.

Aaron M Shew, Jesse B Tack, Lawton L Nalley, and Petronella Chaminuka. Yield reduction under climate warming varies among wheat cultivars in South Africa. *Nature Communications*, 11(1):1–9, 2020.

Guanming Shi, Jean-Paul Chavas, and Joseph Lauer. Commercialized transgenic traits, maize productivity and yield risk. *Nature Biotechnology*, 31(2):111–114, 2013a.

Wenjiao Shi, Fulu Tao, and Zhao Zhang. A review on statistical models for identifying climate contributions to crop yields. *Journal of Geographical Sciences*, 23(3):567–576, 2013b.

Oded Stark and J. Edward Taylor. Migration incentives, migration types: The role of relative deprivation. *Economic Journal*, 101(408):1163, 1991.

Ognen Stojanovski, Gordon W. Leslie, Frank A. Wolak, Juan Enrique Huerta Wong, and Mark C. Thurber. Increasing the energy cognizance of electricity consumers in mexico: Results from a field experiment. *Journal of Environmental Economics and Management*, 102:102323, 2020.

Jesse Tack, Andrew Barkley, and Lawton Lanier Nalley. Estimating yield gaps with limited data: An application to United States wheat. *American Journal of Agricultural Economics*, 97(5):1464–1477, 2015.

Jesse Tack, Jane Lingenfelser, and SV Krishna Jagadish. Disaggregating sorghum yield reductions under warming scenarios exposes narrow genetic diversity in us breeding programs. *Proceedings of the National Academy of Sciences*, 114(35):9296–9301, 2017.

Verena Tiefenbeck, Lorenz Goette, Kathrin Degen, Vojkan Tasic, Elgar Fleisch, Rafael Lalive, and Thorsten Staake. Overcoming salience bias: How real-time feedback fosters resource conservation. *Management Science*, 64(3):1458–1476, 2018.

Tauã Magalhães Vital, Sandy Dall'erba, William Ridley, and Xianning Wang. What do the 235 estimates from the literature tell us about the impact of weather on agricultural and food trade flows? *Global Food Security*, 35:100654, 2022.

Thomas J Voon and Geoff W Edwards. Research payoff from quality improvement: the case of protein in australian wheat. *American Journal of Agricultural Economics*, 74(3): 564–572, 1992.

Ruixue Wang, Roderick M Rejesus, Jesse B Tack, Joseph V Balagtas, and Andy D Nelson. Quantifying the yield sensitivity of modern rice varieties to warming temperatures: Evidence from the philippines. *American Journal of Agricultural Economics*, 104(1):318–339, 2022.

Reiner Wassmann, SVK Jagadish, Sigrid Heuer, A Ismail, E Redona, Rachid Serraj, RK Singh, G Howell, H Pathak, and K Sumfleth. Climate change affecting rice production: the physiological and agronomic basis for possible adaptation strategies. *Advances in Agronomy*, 101:59–122, 2009.

Geoffrey S Watson. Smooth regression analysis. *Sankhyā: The Indian Journal of Statistics, Series A*, 26(4):359–372, 1964.

Jeffrey M Wooldridge. *Introductory econometrics: A modern approach*. Cengage Learning, 2015.