

UC Davis

UC Davis Electronic Theses and Dissertations

Title

Essays on Market Adjustment to Regulation, Infrastructure, and Environmental Constraints

Permalink

<https://escholarship.org/uc/item/6p43n6qq>

ISBN

9798189226656

Author

Wang, Yuhan

Publication Date

2026-06-25

Peer reviewed|Thesis/dissertation

Essays on Market Adjustment to Regulation, Infrastructure, and Environmental Constraints

By

YUHAN WANG
DISSERTATION

Submitted in partial satisfaction of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

Agricultural and Resource Economics

in the

OFFICE OF GRADUATE STUDIES

of the

UNIVERSITY OF CALIFORNIA

DAVIS

Approved:

Richard J. Sexton, Co-Chair

Aaron Smith, Co-Chair

Kevin Novan

Committee in Charge

2026

Abstract

This dissertation studies how producers, consumers, and markets adjust to policy, infrastructure, and environmental constraints. The three chapters address consumption-based production standards in U.S. agricultural markets, electricity service quality and household appliance adoption in Sub-Saharan Africa, and weather-driven supply dynamics in U.S. fresh produce markets. A common theme is that observed market and welfare outcomes are shaped by structural features that aggregate data alone do not reveal.

The first chapter develops a partial equilibrium framework to analyze consumption-based production standards, regulations that condition market access on how goods are produced rather than where. The model features vertically differentiated products, heterogeneous producer compliance costs, and trade between the regulating jurisdiction and external suppliers. I calibrate the model to California’s Proposition 12 and the U.S. shell egg market and simulate the welfare effects of the regulation. The baseline calibration implies that the regulation raises cage-free egg prices nationwide, reduces conventional production, and increases California’s reliance on out-of-state imports. Consumer welfare losses concentrate in California, producers outside California gain from higher prices and expanded demand, and California producers lose on net because conversion costs and exit among conventional producers more than offset gains to incumbent cage-free producers. A Monte Carlo analysis over 10,000 perturbed parameter draws confirms that the direction of these effects is robust to parameter uncertainty. The framework applies more broadly to consumption-based standards such as the European Union’s Deforestation Regulation and other state-level animal welfare mandates.

The second chapter asks whether the quality of electricity service, conditional on grid connection, shapes household appliance adoption in Zambia. Using household-level data from the World Bank Multi-Tier Framework survey, I estimate the effects of electricity quality, defined as achieving the highest reliability and availability tiers, on aggregate appliance capacity and on ownership of 30 individual appliances ranging from 3 to 2,000 watts. For urban households, I estimate OLS and Probit models. For rural households, where service quality may be correlated with unobserved characteristics, I instrument for quality using each household’s distance to the nearest distribution substation, exploiting voltage decline along distribution lines. Improved quality increases aggregate appliance capacity by 308 watts in urban areas and 798 to 1,329 watts in rural areas. The effect concentrates on appliances whose value depends on continuous power, such as refrigerators and electric irons, while appliances used in brief flexible episodes or constrained by purchase cost show little response. For the roughly 1.2 million Zambian households already connected to the grid, distribution infrastructure investments that improve service reliability can generate welfare gains by enabling adoption of appliances households currently forgo.

The third chapter constructs a weekly dataset of USDA truck shipment volumes, FOB shipping point prices, and gridded PRISM weather data to study weather-driven supply and price dynamics in U.S. lettuce markets between 2010 and 2026. Standard yield-based approaches to estimating weather effects are difficult to apply because high-frequency yield data do not exist for fresh vegetables. I develop a dual-layer empirical strategy that combines market-level aggregation for prediction with district-level panel variation for identification. Three results emerge. First, simple linear models outperform tree-based machine learning methods such as XGBoost and Random Forest in out-of-sample price prediction, reflecting strong autoregressive persistence and an approximately linear data-generating process conditional on lagged prices. Second, weather effects on prices are present but attenuated in aggregate data: in district-level panels, each additional freezing degree day raises weekly prices by \$0.85, while this effect is diluted in market-level regressions. Third, the results are robust across lettuce varieties, including Romaine, and across alternative validation strategies. The chapter shows that downstream market outcomes can be used to recover weather-driven supply

responses in settings where direct yield data are unavailable.

The three chapters share an emphasis on how external constraints translate into observed market and household outcomes through mechanisms that aggregate data alone tend to obscure. In the first chapter, the welfare consequences of Proposition 12 depend on the share of domestic production within the regulating jurisdiction and on the distribution of producer compliance costs. In the second chapter, electricity quality acts as a binding constraint on appliance adoption that operates largely independently of household income and affects different appliance categories in different ways. In the third chapter, weather shocks shape lettuce supply and prices through district-level disruptions that are obscured in aggregate data by the strong persistence of prices. Each chapter addresses a distinct policy question, but together they show that the welfare effects of regulation, infrastructure investment, and environmental risk cannot be evaluated from aggregate data alone, and that a careful look at the underlying structure of production, demand, and supply is needed to understand who bears the resulting costs.

Acknowledgements

I am deeply grateful to my dissertation committee, Richard Sexton, Aaron Smith, and Kevin Novan, for their guidance, encouragement, and patience throughout my doctoral studies. I appreciate the time and care they invested in training me to think and work as an economist.

To Rich, my co-chair, thank you for the encouragement that gave me the confidence to begin this PhD and for the steady mentorship that shaped how I approach research and academic life. I am grateful to Aaron, my co-chair, for the opportunity to work with large datasets and apply econometric methods to questions I genuinely cared about. I thank Kevin for his patience, his thoughtful engagement with my work, and the many conversations that challenged and inspired me.

I also owe a great deal to my cohort and friends in the Agricultural and Resource Economics program at UC Davis, especially Zhiran Qin, Mengfei Zhou, Wenjie Zhan, and Hyun Jung Lee. From shared coursework during the first two years to countless conversations about research and life since, their friendship has been one of the best parts of my PhD. We shared the difficulties of academic research and kept each other going through the harder moments, and I am grateful to have gone through this experience alongside them.

Finally, I am deeply indebted to my family for their unwavering love, sacrifices, and belief in me throughout this long journey. Their patience, encouragement, and quiet support across many years and many miles made this dissertation possible. They celebrated every small step forward and stood by me through every difficult moment. I dedicate this work to them.

Contents

List of Figures	xi
List of Tables	xiii
Introduction	1
1 Consumption-Based Production Standards and Producer Adjustment: Evidence from California’s Proposition 12	6
1.1 Introduction	7
1.2 Policy Context: Consumption-Based Production Standards	11
1.2.1 Subnational Examples: United States	11
1.2.2 National and Supranational Examples: European Union	13
1.2.3 Heterogeneous Compliance Costs: Empirical Evidence	15
1.2.4 When Domestic Production Matters	16
1.3 Literature Review	17
1.4 Economic Framework	20
1.4.1 Baseline Model: No Regulation	21
1.4.2 Post-Regulation Equilibrium	27
1.4.3 Social Welfare	32
1.5 Model Calibration	35
1.5.1 Application: California’s Egg Market	35

1.5.2	Calibration Strategy	37
1.5.3	Fixed Parameters and Constraints	37
1.5.4	Data	38
1.5.5	Estimation Procedure	38
1.5.6	Model Fit	39
1.6	Simulation Results	41
1.6.1	Prices and Quantities	41
1.6.2	Elasticities	43
1.6.3	Trade and Market Integration	44
1.6.4	Consumer Welfare	46
1.6.5	Producer Surplus	48
1.6.6	Overall Welfare Effects	49
1.6.7	Comparison with Empirical Evidence	51
1.7	Robustness Analysis	56
1.7.1	Design	56
1.7.2	Results	57
1.8	Conclusion	62
	References	65
	Appendix	70
1.9	Appendix	71
1.9.1	Comparative Statics	71
1.9.2	Compensating Variation	73
1.9.3	Producer Surplus	78
2	Beyond Access: How Does Electricity Quality Affect the Adoption of Energy-Consuming Appliances?	80
2.1	Introduction	81

2.2	Background	83
2.2.1	Electricity Quality and Appliance Usage	83
2.2.2	Energy Use in Zambia	84
2.3	Literature Review	85
2.4	Data	87
2.4.1	Household Survey Data	87
2.4.2	Grid Data	89
2.4.3	Climate Data	89
2.5	Estimation Strategy	94
2.5.1	Model	94
2.5.2	Urban	98
2.5.3	Rural	98
2.5.4	Appliance-Level Models	100
2.6	Empirical Results	100
2.6.1	Impacts of Electricity Quality	100
2.6.2	Household Appliance Ownership	105
2.6.3	Comparison to Electricity Access	108
2.6.4	Heterogeneity by Income and Gender	114
2.7	Conclusion	116
	References	118
3	Weather Shocks, Supply Dynamics, and Price Persistence: Evidence from U.S. Lettuce Markets	122
3.1	Introduction	123
3.2	Literature Review	124
3.3	Market Background	126
3.4	Data	128

3.4.1	Sources	128
3.4.2	Production Districts	128
3.4.3	Variable Construction	129
3.4.4	Summary Statistics	130
3.5	Methodology and Empirical Strategy	131
3.5.1	Dual-Layer Design	131
3.5.2	Market-Level Prediction Model	132
3.5.3	Change Model	133
3.5.4	District-Level Panel and Identification	133
3.5.5	Validation Strategy	134
3.6	Results	135
3.6.1	Market-Level Price Prediction	135
3.6.2	Price Persistence and Hidden Weather Effects	136
3.6.3	Mechanism: Weather, Supply, and Price	138
3.6.4	Volume Prediction	140
3.7	Robustness	141
3.7.1	Romaine Lettuce	141
3.7.2	Expanding vs. Rolling Window	142
3.7.3	Coverage and Prediction Error	143
3.7.4	The 2022 Extreme Year	144
3.7.5	Feature Stability	145
3.7.6	Missing Price Prediction	146
3.8	Conclusion	146
	References	148
	Appendix	151
3.9	Appendix	152
3.9.1	Variable Definitions	152

3.9.2	Full OLS Regression Output	153
3.9.3	Out-of-Sample RMSE by Year	154

List of Figures

1.1	Utility as a function of income under vertical differentiation.	24
1.2	Change in income threshold within the jurisdiction (California)	47
1.3	Change in income threshold outside the jurisdiction	47
1.4	Change in conversion cost threshold within the jurisdiction (California)	49
1.5	Change in conversion cost threshold outside the jurisdiction	50
1.6	California egg production, 2013–2024. Data from USDA NASS, California Poultry Reports.	53
1.7	Cage-free share of U.S. laying hen flock, 2016–2024. Data from USDA AMS, Cage- Free Shell Egg Reports.	54
1.8	Distribution of Pre-Regulation Equilibrium Moments	59
1.9	Distribution of Post-Regulation Price and Quantity Changes	60
1.10	Distribution of California Import Share for Restricted Eggs	60
1.11	Distribution of Welfare Outcomes	61
2.1	Zambia: Electricity Access and Quality Distribution	93
2.2	Zambia: Köppen Climate Classification and Sample Distribution	95
2.3	Distribution of Distance to the Nearest Substation, Electricity Access and Quality	99
2.4	Average Marginal Effects on Household Appliance Ownership: Percentage Points Change	109
2.5	Average Marginal Effects on Household Appliance Ownership: Percentage Change	110

3.1	Seasonal patterns in shipment volume and prices by production district.	127
3.2	Market-level prices: actual and OLS predictions.	137
3.3	Feature importance in price models.	138
3.4	Price coverage ratio and absolute prediction error.	144

List of Tables

1.1	Consumption-based production standards: role of domestic production	17
1.2	Individual Compensating Variation by Behavioral Transition	33
1.3	Calibration Targets	37
1.4	Calibrated Parameters	39
1.5	Impacts of Proposition 12 on Egg Prices, Quantities, and Price Premium	43
1.6	Demand and Supply Elasticities Before and After Proposition 12	44
1.7	Cross-Price Elasticities Before and After Proposition 12	45
1.8	Trade Effects of Proposition 12 on California Egg Market	46
1.9	Consumer Welfare (CV) Effects of Proposition 12	48
1.10	Producer Surplus (PS) Effects of Proposition 12	51
1.11	Monte Carlo Parameter Perturbation	57
1.12	Sign Consistency across 10,000 Monte Carlo Draws	57
1.13	Welfare Outcomes across Monte Carlo Draws (% of MV_0)	58
2.1	Distribution of sampled households	88
2.2	Data Description for Selected Variables	90
2.3	Descriptive statistics	91
2.4	Appliances by load level	92
2.5	Balance Table: Good vs. Poor Electricity Quality	97
2.6	Urban: Impacts of Electricity Quality	101

2.7	Rural: Impacts of Electricity Quality	103
2.8	Rural: First Stage Estimation	104
2.9	Appliance Baseline Ownership and Average Marginal Effects of Electricity Quality	106
2.10	Urban: Impacts of Electricity Access	112
2.11	Rural: Impacts of Electricity Access	113
2.12	Urban: Heterogeneity in Quality Effects (Interaction Models)	115
2.13	Heterogeneity in Quality Effects by Income and Gender (Split Sample)	116
3.1	Summary statistics for market-level data, 2010–2026	131
3.2	Out-of-sample performance for price models	136
3.3	District-level effects of weather on prices	139
3.4	Out-of-sample performance for shipment volume models	141
3.5	Robustness: iceberg vs. romaine lettuce	141
3.6	Robustness to expanding and rolling window validation	142
3.7	Prediction error by coverage ratio	143
3.8	2022 extreme year: model performance comparison	145
3.9	Coefficient stability across sub-periods (market-level OLS)	145
3.10	Missing price prediction by district	146
A1	Variable definitions and data sources	152
A2	Full OLS regression: market-level price model	153
A3	Out-of-sample RMSE by year (\$, expanding window)	154

Introduction

Markets and households increasingly face external constraints that shape production, consumption, and market outcomes. These constraints take many forms, including regulation, infrastructure, and environmental shocks, but they share a common economic role: changing the conditions under which households and firms make decisions.

Governments now regulate not only how goods are produced within their borders but also which goods can enter their consumer markets, conditioning market access on production practices. Electricity grids in developing countries have expanded substantially over the past two decades, but reliability and voltage quality vary widely across connected households, raising the question of whether access alone is sufficient for welfare gains. Climate change and weather variability place increasing pressure on agricultural supply chains, particularly for perishable specialty crops where small disruptions in production translate quickly into price volatility.

In each of these settings, the relevant economic question is not whether the constraint exists, but how the agents subject to it actually adjust their behavior.

This dissertation examines three such settings, each studying a different kind of constraint and a different population of decision makers. The first chapter analyzes consumption-based production standards in U.S. agricultural markets, using California's Proposition 12 and the U.S. shell egg market as the empirical application. The second chapter studies how electricity service quality affects household appliance adoption in Zambia, using rich household-level survey data on grid access, reliability, and appliance ownership. The third chapter examines how weather shocks propagate into supply and prices in U.S. lettuce markets, using high-frequency shipment and price data from

California and Arizona production districts. The three settings span different policy environments, geographies, and econometric strategies.

A common thread across the three chapters is that observed market outcomes are shaped by structural features that aggregate data alone do not reveal. The first chapter shows that the welfare effects of consumption-based regulation depend on heterogeneity in producer compliance costs and on the share of domestic production in the regulated market. The second chapter shows that the welfare gains from grid infrastructure depend on service quality among already connected households, a margin separate from access itself. The third chapter shows that weather shocks affect agricultural prices through localized district-level disruptions that are obscured in aggregate data by spatial averaging and price persistence. In each case, looking beneath aggregate outcomes is necessary to understand the economic mechanisms through which these constraints operate.

The first chapter develops a partial equilibrium framework to analyze consumption-based production standards, regulations that condition market access on how goods are produced rather than where. Recent state and supranational policies, including California’s Proposition 12 and the European Union’s Deforestation Regulation, illustrate this approach. Proposition 12, which took effect in January 2022, banned the sale of conventionally produced eggs in California. Although the policy applies only within California, its economic consequences extend beyond the state because California accounts for a meaningful share of national egg consumption while producing only a small share of the national flock. Existing work on Proposition 12 and related regulations has focused primarily on consumer willingness to pay or on aggregate price effects. Less attention has been paid to two features that, the chapter argues, are central to the welfare consequences of these regulations: heterogeneity in producer compliance costs and the role of domestic production within the regulating jurisdiction.

To address these gaps, the model features vertically differentiated products, heterogeneous producer compliance costs, and trade between the regulating jurisdiction and external suppliers. I calibrate the model to the U.S. shell egg market and simulate the welfare effects of Proposition 12. The baseline calibration implies that the regulation raises cage-free egg prices nationwide, reduces

conventional production, and substantially increases California’s reliance on out-of-state cage-free imports. Consumer welfare losses concentrate in California. Producers outside California gain from higher prices and expanded demand, while California producers lose on net because conversion costs and exit among conventional producers more than offset gains to incumbent cage-free producers. A Monte Carlo analysis over 10,000 perturbed parameter draws confirms that the direction of these effects is robust to parameter uncertainty. Although the empirical application focuses on eggs, the framework applies more broadly to consumption-based standards across jurisdictions and commodities.

The second chapter turns from regulation to infrastructure and asks whether the quality of electricity service, conditional on grid connection, shapes household appliance adoption in Zambia. A large literature in development economics studies the effects of grid access on household welfare, but much less attention has been paid to electricity quality among households that are already connected. Many newly connected households in Sub-Saharan Africa experience frequent outages, voltage fluctuations, and limited hours of supply, and there is growing evidence that these reliability constraints matter for whether households actually translate grid access into welfare gains. The chapter asks whether quality, defined separately from access, is itself a binding constraint on appliance adoption among connected households, and which appliances are most sensitive to it.

Using household-level data from the World Bank Multi-Tier Framework survey in Zambia, I estimate the effects of electricity quality, defined as achieving the highest reliability and availability tiers, on aggregate appliance capacity and on ownership of 30 individual appliances ranging from 3 to 2,000 watts. For urban households, I estimate OLS and Probit models. For rural households, where service quality may be correlated with unobserved household characteristics, I instrument for quality using each household’s distance to the nearest distribution substation, exploiting the fact that voltage declines with distance along distribution lines. Improved quality increases aggregate appliance capacity by 308 watts in urban areas and 798 to 1,329 watts in rural areas. The effect concentrates on appliances whose value depends on continuous power, such as refrigerators and electric irons, while appliances used in brief flexible episodes or constrained by high purchase cost

show little response. For the roughly 1.2 million Zambian households already connected to the grid, distribution infrastructure investments that improve service reliability can yield welfare gains by enabling adoption of appliances that households currently forgo due to poor power quality.

The third chapter turns to environmental constraints and studies weather-driven supply responses in U.S. lettuce markets, a setting in which the standard yield-based approach to estimating climate effects on agriculture is difficult to apply. Most existing work on weather and agricultural production relies on county-level annual yield data, but high-frequency yield data do not exist for fresh vegetables such as lettuce, which are produced on a continuous, rolling basis rather than through a single annual harvest. At the same time, lettuce supply chains are highly weather sensitive: freeze events, heavy rain, and excessive heat all disrupt local production and translate into observable price spikes in downstream markets. The chapter asks whether downstream market outcomes, specifically shipment volumes and FOB prices, can be used to recover weather-driven supply responses in the absence of direct yield measures.

I construct a weekly dataset that integrates USDA truck shipment volumes, FOB shipping point prices, and gridded PRISM weather data for six lettuce production districts in California and Arizona between 2010 and 2026. I develop a dual-layer empirical strategy that combines market-level aggregation for prediction with district-level panel variation for identification. Three results emerge. First, simple linear models outperform tree-based machine learning methods such as XGBoost and Random Forest in out-of-sample price prediction, reflecting strong autoregressive persistence and an approximately linear data-generating process conditional on lagged prices. Second, weather effects on prices are present but attenuated in aggregate data: in district-level panels, each additional freezing degree day raises weekly prices by \$0.85, while this effect is diluted in market-level regressions. Third, the results are robust across lettuce varieties, including Romaine, and across alternative validation strategies. The chapter shows that downstream market outcomes can be used to recover weather-driven supply responses in settings where direct yield data are unavailable.

Taken together, the three chapters show that the economic effects of regulation, infrastructure investment, and environmental shocks depend critically on the margins through which households

and firms adjust to external constraints. In the first chapter, producer entry, exit, and conversion decisions determine the incidence of consumption-based regulation across regions. In the second chapter, the benefits of electrification depend not only on access itself but also on whether service quality is sufficient to support appliance adoption. In the third chapter, weather shocks propagate into downstream prices through short-run disruptions in geographically concentrated production systems. Each chapter addresses a distinct policy question. Considered together, they suggest that evaluating these policies requires close attention to the mechanisms through which producers, consumers, and markets respond to the constraints they face.

Chapter 1

Consumption-Based Production Standards and Producer Adjustment: Evidence from California's Proposition 12

Abstract

Governments increasingly condition market access on how goods are produced, requiring that all products sold within a jurisdiction meet specific production standards regardless of origin. This paper develops a partial equilibrium model to study the welfare effects of such consumption-based production standards when the regulating jurisdiction is both a consumer and a producer of the regulated commodity. The model features vertically differentiated products, heterogeneous producer compliance costs, and trade between the regulating jurisdiction and the rest of the market. I illustrate the model using a calibration based on California's Proposition 12, which banned the sale of conventionally produced eggs beginning in January 2022, and the U.S. shell egg market. In the baseline calibration, I find that the regulation raises cage-free egg prices nationwide, reduces conventional production, and substantially increases California's dependence on out-of-state cage-free imports. Consumer welfare losses are concentrated in California, while producers outside California gain from higher prices and expanded demand. California producers lose on net because conversion costs and exit among conventional producers more than offset gains to existing cage-free producers. Total welfare declines despite an increase in industry rev-

enue. A Monte Carlo analysis over 10,000 perturbed parameter draws confirms that the quantity reallocation, consumer welfare losses, and net welfare decline are robust, while the price-level and producer surplus effects are more sensitive to parameter values. The implications are broadly consistent with observed market outcomes since 2022. Although the empirical application focuses on eggs, the framework applies to consumption-based standards more generally, including the European Union’s Deforestation Regulation and state-level animal welfare mandates, wherever a jurisdiction conditions market access on production practices for goods it both consumes and produces.

1.1 Introduction

Governments increasingly regulate not only how products are produced within their borders, but also how products sold within their jurisdiction are produced, regardless of origin. These consumption-based production standards have emerged relatively recently as a policy instrument in areas such as animal welfare, environmental protection, and climate policy. Instead of regulating only domestic production practices, these policies condition market access on compliance with specific production requirements. Producers must satisfy those requirements if they wish to sell into the regulated market, even when production occurs outside the jurisdiction.

This matters economically because consumption-based standards operate differently from traditional production regulations. Conventional production mandates apply only to firms operating within the jurisdiction. In integrated markets, such policies raise costs for domestic producers while leaving external suppliers unaffected, potentially shifting production toward regions with weaker standards. By contrast, when a jurisdiction restricts the sale of noncompliant goods, and in some cases also their production, all suppliers seeking access to that market must meet the same production requirements. The policy therefore extends local standards beyond the jurisdiction’s borders and reduces, but does not eliminate, the competitive disadvantage that domestic producers would otherwise face.

Governments adopt such policies for two related reasons. First, they allow jurisdictions to

extend local preferences, such as concerns about animal welfare or environmental protection, to goods produced elsewhere but sold within the jurisdiction. Second, they prevent domestic producers from being undercut by external suppliers who face no comparable requirements. However, even consumption-based standards do not fully equalize the position of domestic and external producers. Domestic producers must comply or exit the market entirely, losing access to all buyers. External producers who choose not to comply lose access only to the regulated market but can continue serving buyers elsewhere at unchanged cost. This asymmetry means that the regulation may disproportionately burden domestic producers, a mechanism that is central to the analysis developed in this paper.

Although recent discussion in the United States has focused on state-level regulations, particularly California’s Proposition 12, the underlying mechanism is not limited to subnational policy. Similar approaches appear at national and supranational levels. The European Union’s Deforestation Regulation (EUDR), for example, prohibits both the sale and the export of products linked to deforestation, covering commodities such as wood, soy, beef, coffee, and cacao. Proposition 12 bans both noncompliant production within California and sales of noncompliant products in the state. Because in-state production of noncompliant eggs is itself prohibited, the question of whether California producers could sell noncompliant products elsewhere does not arise. The EUDR takes a different approach: it is silent on production practices, but prohibits both domestic sale and export of noncompliant goods. The legal paths differ, but the economic consequence for domestic producers is the same. In both cases, producers who cannot meet the standard lose their access to markets. The economic effects of such policies depend critically on whether the regulating jurisdiction also produces the regulated commodity.

For commodities such as coffee or cacao, which the European Union does not produce domestically, the regulation primarily affects imports. In that case the policy functions largely as a market-access restriction affecting foreign suppliers. The situation differs for commodities such as soy or beef that are produced within the European Union. In those markets the regulation also affects domestic producers, who may need to modify production practices, incur conversion costs,

or exit production. As a result the policy can reshape the structure of the domestic industry in addition to altering trade flows.

California’s Proposition 12 illustrates the same distinction at the state level. The regulation applies to several livestock products, including pork and eggs, but the economic effects differ because California’s production shares vary across commodities. California produces less than 1% of the pork it consumes, so the pork regulation operates mainly through imports. For eggs, California accounts for roughly 12% of national consumption and about 3–4% of national production, meaning the regulation affects both interstate trade and in-state production decisions.

Another important feature of these policies is that compliance costs differ across producers. Farms vary in existing infrastructure, scale, access to capital, and current production practices, so that the cost of meeting a given standard ranges widely. Some producers already operate close to the regulated standard, while others must undertake substantial investments or exit entirely. [Lee, Sexton and Sumner \(2023\)](#), for example, document large variation in conversion costs among hog farms facing California’s housing requirements. Similar variation arises in pesticide regulations and environmental standards.

Despite the growing use of consumption-based production standards, much of the existing literature focuses primarily on consumer willingness to pay for regulated attributes or assumes homogeneous compliance costs. Less attention has been given to how heterogeneous producers adjust when access to a major consumer market is conditioned on production practices. In particular, the role of domestic production within the regulating jurisdiction has received relatively little attention, even though it may fundamentally shape the welfare consequences of such policies.

This paper develops a partial equilibrium framework to analyze consumption-based production standards in integrated markets with heterogeneous producers and trade. The framework highlights how the effects of these policies depend on two key factors: the share of production located within the regulating jurisdiction and heterogeneity in compliance costs across producers. These features determine which producers adopt compliant technologies, which redirect production to other, unregulated markets, and how production and trade flows adjust across regions. The

paper also contributes on the demand side by explicitly characterizing heterogeneous consumer responses within and outside the regulating jurisdiction. The model distinguishes between consumers who continue purchasing the same product, those who switch across product types, and those who exit the market altogether.¹ This decomposition provides a tractable way to link policy-induced price changes to aggregate demand responses and welfare outcomes. To my knowledge, this type of decomposition has received limited attention in models of consumption-based production standards.

I apply the framework to California’s Proposition 12 regulation of shell eggs, which took effect in January 2022. The egg market provides a useful empirical setting because California accounts for a large share of consumption and a smaller, but still substantial share, of production in a nationally integrated market. This creates substantial trade linkages with producers in other states while also leaving a meaningful domestic industry that must adjust to the regulation.

To illustrate the model’s mechanisms, I calibrate it to the U.S. shell egg market using pre-policy data on prices, quantities, and trade flows, and simulate the effects of Proposition 12. The model implies welfare losses for consumers in the regulating jurisdiction as egg prices increase and conventional eggs disappear from the market. Producers within the jurisdiction also lose on net despite higher market prices, because conversion costs and exit among conventional producers outweigh gains to existing cage-free producers. Producers outside the jurisdiction benefit from higher prices and increased market share, capturing rents created by the regulation. Overall, total welfare declines even though industry revenue increases. I compare these implications with observed market outcomes in the years following implementation of Proposition 12, using data from 2022 to 2024, as a qualitative benchmark.

Although the empirical application focuses on California eggs, the framework applies more broadly to any setting where a jurisdiction conditions market access on production practices. The core mechanisms, including sorting of heterogeneous producers, asymmetric adjustment across regions, and the interaction between domestic production and trade, arise whenever such policies are adopted.

¹I refer to these groups as incumbents, converters, and exiters throughout the paper.

The remainder of the paper proceeds as follows. Section 2 provides institutional context on consumption-based production standards. Section 3 reviews the related literature. Section 4 develops the theoretical framework. Section 5 describes the calibration strategy. Section 6 presents simulation results. Section 7 examines robustness. Section 8 concludes.

1.2 Policy Context: Consumption-Based Production Standards

As discussed in the introduction, consumption-based production standards condition market access based on how goods are produced, extending local requirements to all suppliers regardless of origin. This section reviews the main examples of such policies and documents the empirical evidence on compliance cost heterogeneity that motivates the modeling framework developed later in the paper.

1.2.1 Subnational Examples: United States

Several U.S. states have adopted consumption-based regulations governing livestock production practices. California provides the most prominent example. California’s animal welfare regulations developed in several stages. Proposition 2, approved by voters in 2008, established minimum space requirements for egg-laying hens, breeding pigs, and calves raised for veal. The measure applied only to production within California and did not specify precise space requirements, which created uncertainty regarding implementation and compliance.

Assembly Bill 1437, enacted in 2010, clarified the earlier measure and extended restrictions to eggs sold in California, regardless of where they were produced. This represented California’s first explicit consumption-based animal welfare standard.

Proposition 12, approved by voters in 2018 and implemented beginning in January 2022, replaced these earlier laws with clear and specific requirements.² The regulation applies to three

²California Health and Safety Code §§ 25990–25994. Section 25990(a) prohibits the confinement of covered animals in a manner inconsistent with the standards, and Section 25990(b) prohibits the sale of products from animals so confined. The full text is available at https://leginfo.legislature.ca.gov/faces/codes_displayText.xhtml?lawCode=HSC&division=20.&title=&part=&chapter=13.8.&article=.

livestock products: shell eggs, pork from breeding pigs, and veal. For egg production, the law requires cage-free housing systems providing at least 216 square inches of usable floor space per hen. For pork production, breeding pigs must be housed with at least 24 square feet of space per animal, and similar confinement standards apply to veal calves. Products produced under systems that do not satisfy these requirements cannot be sold in the California market.

The regulation attracted national attention because of its implications for interstate commerce. In *National Pork Producers Council v. Ross* ([Supreme Court of the United States, 2023](#)), the U.S. Supreme Court upheld Proposition 12 against a Commerce Clause challenge, confirming that states may condition market access on production practices even when compliance requires adjustments by out-of-state producers.

California’s production structure differs sharply across the commodities covered by the law, which leads to different economic effects across markets. In the egg industry, California produced approximately 4.3 billion eggs per year from roughly 15 million hens as of 2018, making it one of the larger egg-producing states in the United States ([U.S. Department of Agriculture, National Agricultural Statistics Service, 2019](#)). However, this represents only about 3–4% of national egg production, while California accounts for roughly 12% of national egg consumption. The state is therefore a large net importer of eggs, but still maintains a meaningful domestic egg industry.

The situation differs for pork. California produces less than 1% of the pork it consumes, meaning that the regulation primarily affects imports from other states. As a result, the economic mechanisms differ across commodities covered by Proposition 12: the egg market involves both domestic production adjustments and trade effects, whereas the pork market operates primarily through trade.

Other states have adopted similar consumption-based animal welfare regulations. Massachusetts voters approved Question 3 in 2016, which restricts sales of eggs and pork produced using certain confinement systems. Because Massachusetts produces very little pork or eggs domestically, the regulation primarily affects imports from other states. Michigan has enacted its own ban on the sale of eggs from caged hens. Michigan has a larger egg industry than Massachusetts, so its regula-

tion affects both in-state production and imports, though Michigan’s share of national egg output remains modest.

This remains an area of active policy debate. Constitutional challenges to state confinement laws continue in federal courts, and the U.S. House Committee on Agriculture included language in its 2026 Farm Bill that would preempt state-level confinement standards for most livestock products in interstate commerce, though eggs are currently excluded from that provision.

1.2.2 National and Supranational Examples: European Union

The European Union’s Deforestation Regulation (EUDR), adopted in 2023, provides a clear example of consumption-based standards at the supranational level. The EUDR restricts the sale within the EU of products associated with deforestation after December 2020. Covered commodities include wood, soy, beef, coffee, cacao, rubber, and palm oil. Firms placing these products on the EU market must demonstrate, through due diligence procedures, that their supply chains are not linked to recent deforestation. Implementation has been postponed several times, with the current deadline set for December 2026 for large and medium operators and June 2027 for micro and small operators.

A key legal distinction is that the EUDR regulates market access rather than production practices directly. The regulation is silent on whether deforestation may occur within the EU; instead, it prohibits both the sale and the export of deforestation-linked products. Under Article 3 of the regulation, covered products “shall not be placed or made available on the market or exported” unless they satisfy the deforestation-free requirements.³ An EU producer who cleared forest to grow soy could therefore not sell that soy domestically or export it to third countries. The legal paths differ from Proposition 12: the EUDR restricts both sale and export but is silent on production practices, whereas Proposition 12 directly bans noncompliant production within California but is silent on exports. In practice, however, the effect on domestic producers is similar, because

³Regulation (EU) 2023/1115 of the European Parliament and of the Council of 31 May 2023 on deforestation-free supply chains, Article 3. See https://environment.ec.europa.eu/topics/forests/deforestation/regulation-deforestation-free-products_en.

under either approach producers of noncompliant goods face severely limited commercial outlets and must either adopt compliant practices or exit. Despite this difference in legal mechanism, the EUDR shares the core structure of a consumption-based standard: it conditions market access on production practices, affects both domestic and foreign suppliers, and generates heterogeneous compliance costs.

As with Proposition 12, the economic effects depend on whether the EU also produces the regulated commodity. For coffee and cacao, which the EU does not grow, the EUDR operates primarily through trade: foreign producers must comply or lose market access, while EU consumers face higher prices. For soy and beef, which the EU does produce, the regulation also reshapes domestic industry. Some EU producers already satisfy the deforestation-free requirements and stand to benefit from reduced import competition. Others face substantial adjustment costs. Beef farmers who source Brazilian soy for feed, for instance, must switch to costlier EU-grown alternatives or undertake expensive supply-chain verification. Producers with high compliance costs may exit entirely. The regulation thus sorts domestic producers based on their existing practices, with very different welfare implications than in the coffee case.

Other EU Examples

The European Union prohibited conventional caged hen housing within its member states beginning in 2012 and has considered extending animal welfare standards to imported products. Because the EU also produces eggs domestically, such policies affect both trade and internal production decisions.

A similar logic applies to the EU's Carbon Border Adjustment Mechanism, adopted in 2023, which extends carbon pricing to imports of products such as steel, cement, fertilizer, aluminum, and electricity. Since the European Union also produces these goods, the policy affects both foreign suppliers and domestic industry. Compliance costs differ across firms depending on technology, fuel source, and capital vintage.

1.2.3 Heterogeneous Compliance Costs: Empirical Evidence

A key driver of the welfare effects discussed above is that compliance costs vary substantially across producers. This subsection reviews the empirical evidence on that heterogeneity.

Evidence from Animal Confinement Regulations

[Lee, Sexton and Sumner \(2023\)](#) provide detailed evidence on heterogeneous conversion costs for hog production under California’s Proposition 12. They show that about 30% of North American sow operations already use group housing, but with per-sow space that does not meet Proposition 12’s 24-square-foot requirement. For these farms, compliance requires increasing space allowances from roughly 20 to 24 square feet per sow, raising fixed costs by about 20%, with incremental capital costs of approximately \$3 per pig. Compliant farms also face higher variable costs of about \$2 per pig due to reduced efficiency in breeding, farrowing, and nursing, yielding total incremental costs of roughly \$5 per head. Farms that rely on conventional gestation stalls face substantially higher conversion costs because they must rebuild housing infrastructure entirely.

[Lee, Sexton and Sumner \(2023\)](#) assume conversion costs are heterogeneous across farms, uniformly distributed over a range from low to high. In their model, only farms with sufficiently low conversion costs find it profitable to serve the regulated market, while higher-cost farms continue producing under conventional practices for the unregulated market.

Evidence from Cage-Free Egg Production

[Sumner et al. \(2008\)](#) estimate that cage-free egg production is at least 20% more costly than conventional cage production. The cost difference reflects higher feed costs, higher mortality, more labor-intensive management, and larger capital requirements per hen.

These figures are averages, however, and actual costs vary substantially across farms. Some producers already operated cage-free systems before California’s earlier animal welfare regulations and therefore faced little or no conversion cost. Others would have needed complete barn reconstruction. More recent work by [Caputo et al. \(2023a\)](#) also shows that access to capital is an

important constraint, suggesting that otherwise profitable conversion may still be infeasible for credit-constrained producers.

Evidence from Other Regulatory Contexts

Similar patterns arise in other regulatory settings. Compliance costs associated with pesticide regulations may depend on crop type, soil conditions, pest pressure, and existing pest management practices. Environmental regulations also generate heterogeneous costs across firms due to differences in technology, fuel source, and capital vintage. In each case, existing production conditions shape the cost of compliance.

The evidence from animal confinement, egg production, pesticide regulation, and environmental standards collectively suggests that wide dispersion in compliance costs is the norm rather than the exception. This empirical regularity motivates the heterogeneous cost structure in the model developed below.

1.2.4 When Domestic Production Matters

The examples reviewed above highlight a central but underexplored feature of consumption-based production standards. While these policies are often viewed as extending regulatory requirements to all suppliers, they do not impose symmetric constraints. Domestic producers must either comply with the regulation or exit the regulated market, as their production is tied to the jurisdiction. In contrast, external producers who do not comply can reallocate output to unregulated markets and continue production. The domestic producer has no such fallback.

This asymmetry in adjustment margins has not been systematically analyzed in the existing literature and is central to understanding the equilibrium and welfare effects of such policies.

The implications of this mechanism depend critically on the jurisdiction's role as a producer, not just as a consumer. Table 1.1 summarizes the key cases.

When the jurisdiction produces little of the regulated good, the policy operates mainly through trade and functions as a market-access restriction affecting external suppliers. When the jurisdiction

Table 1.1: Consumption-based production standards: role of domestic production

Policy	Commodity	Domestic production
EUDR	Coffee, cacao	None (trade effects only)
EUDR	Soy, beef	Significant (trade + domestic adjustment)
CA Prop 12	Pork	Minimal (<1% of consumption)
CA Prop 12	Eggs	Moderate ($\approx 3\%$ of national output)
MA Question 3	Eggs, pork	Minimal

is also a significant producer, the asymmetry in adjustment becomes more consequential: domestic producers must bear the cost of conversion or exit, while external producers retain the option to serve alternative markets. As a result, the regulation reshapes domestic industry structure in addition to altering trade flows.

The framework developed in Section 1.4 is designed to capture this asymmetry and the interaction between domestic adjustment and trade reallocation.

1.3 Literature Review

This paper sits at the intersection of several literatures: trade effects of product standards, regulation with heterogeneous compliance costs, vertical differentiation models, and the empirical work on California’s egg market. I review each in turn.

Product and process standards have long been recognized as barriers to trade. Sanitary and technical measures significantly reduce bilateral trade flows, particularly for developing-country exporters (Disdier, Fontagné and Mimouni, 2008). Whether such standards are welfare-improving or merely protectionist depends on context. Beghin, Maertens and Swinnen (2015) argue that standards can be trade-creating when they resolve information asymmetries, while Swinnen and Vandemoortele (2011) emphasize that standards on imported products often serve both legitimate regulatory goals and protectionist interests simultaneously. The welfare calculus is similarly ambiguous for border carbon adjustments: Fischer and Fox (2012) and Böhringer, Carbone and Rutherford (2012) both find that such adjustments reduce emissions leakage but shift costs onto trade partners.

[Marette and Beghin \(2010\)](#) reach a more optimistic conclusion, arguing that mandatory standards on imports can improve welfare even when they carry protectionist side effects, provided they address genuine externalities.

Most of this literature focuses on national or supranational policy. The subnational case has received less attention, despite the growing importance of state-level regulation in the United States. [Sumner \(2017\)](#) surveys U.S. state and local regulation of farm practices. [Carter and Scheitrum \(2021\)](#) go further, showing that when a state as large as California imposes production standards, the effects propagate through interstate markets, raising wholesale egg prices nationally. The present paper formalizes this mechanism in an equilibrium framework.

The environmental economics literature has established that compliance costs vary widely across firms, and that this heterogeneity matters for both efficiency and distribution. Uniform standards are inefficient when abatement costs differ ([Lichtenberg and Zilberman, 1988](#)). Empirically, the pattern is consistent: Clean Air Act regulations reduced manufacturing activity unevenly across sectors ([Greenstone, 2002](#)), and emissions trading in Southern California generated gains that depended critically on cross-firm variation in abatement costs ([Fowlie, Holland and Mansur, 2012](#)).

In agriculture, the sources of cost heterogeneity are concrete: farm infrastructure, scale, existing production practices, and access to capital. Dairy farms in California’s San Joaquin Valley face widely varying costs of complying with air quality regulation ([Zhang, 2018](#)). For hog producers subject to Proposition 12, conversion costs depend heavily on whether farms had already adopted group housing ([Lee, Sexton and Sumner, 2023](#)). Among egg producers, access to capital appears to be the binding constraint for many operations considering cage-free conversion ([Caputo et al., 2023a](#)). These findings motivate the continuous cost distribution in my model.

The demand side of the model builds on the vertical differentiation framework of [Shaked and Sutton \(1982\)](#), where consumers with heterogeneous willingness to pay sort across products of different quality. The theoretical foundations go back to [Mussa and Rosen \(1978\)](#). In the context of quality regulation, [Rønne \(1991\)](#) shows that minimum quality standards can intensify competition and raise welfare by eliminating the lowest-quality products, while [Crampes and Hollander \(1995\)](#)

extend this result to duopoly. Empirical work on egg demand confirms that consumers substitute among conventional, cage-free, and organic eggs in a manner consistent with quality-based sorting (Bakhtavoryan et al., 2021), which supports the modeling approach taken here. This pattern is consistent with a vertical differentiation setting in which higher-quality buyers have a close fallback option, while lower-quality buyers face weaker incentives to trade up.

There is by now a sizeable empirical literature on California’s egg regulations. On prices, Mullally and Lusk (2018) estimate that earlier regulations (Proposition 2 and AB 1437) raised retail egg prices by roughly 22%, and Malone and Lusk (2016) report comparable magnitudes. Oh and Vukina (2022) put the annual social cost to California consumers in the hundreds of millions of dollars. On the demand side, willingness to pay for cage-free eggs varies substantially, with roughly half of consumers unwilling to pay more than a modest premium (Lusk, 2019), and there is a well-documented gap between voting for animal welfare mandates and actually purchasing the higher-priced products (Paul et al., 2019).

On the supply side, cage-free production costs exceed conventional costs by at least 20% (Sumner et al., 2008), and capital budgeting analyses suggest that cage-free conversion offers modest returns absent regulatory pressure (Trejo-Pech and Thompson, 2021). Supply chains are sensitive to demand shocks, as demonstrated during the COVID-19 pandemic when egg markets experienced significant disruptions (Malone and Lusk, 2021). These studies confirm that the U.S. egg market is highly integrated and that California’s regulations have national reach.

To my knowledge, no existing study simultaneously integrates all three dimensions. The demand-side literature models heterogeneous consumers but abstracts from production costs; the supply-side literature documents conversion challenges but outside a market-clearing framework; and the interstate spillover literature estimates price effects without decomposing welfare across consumer and producer groups in both regions. This paper integrates vertically differentiated demand, heterogeneous producer compliance costs, and interregional trade into a single equilibrium model, allowing welfare decomposition across all affected groups. While the empirical application is to California’s egg market, the framework applies to consumption-based production standards more

generally.

1.4 Economic Framework

To study the economic effects of consumption-based production standards, I develop a partial equilibrium framework with two integrated regions: a regulating jurisdiction and the rest of the market. The regions are connected through trade in the regulated product. Firms in either region can supply the broader market, but regulatory requirements apply only to products sold within the regulating jurisdiction. Markets are perfectly competitive.

Producers can supply either a conventional product or a compliant product that satisfies the regulatory standard. Compliant production is more costly than conventional production because it requires changes in production practices, facilities, or inputs. Producers in both regions differ in their costs of adopting the compliant technology. As discussed in the introduction, the regulation affects domestic and external producers asymmetrically, and the magnitude of this asymmetry depends on how much the jurisdiction produces.

On the demand side, I adopt a vertical differentiation framework following [Shaked and Sutton \(1982\)](#). In this framework, the compliant product is treated as higher quality than the conventional product. Consumers agree on the ranking of product quality but differ in their willingness to pay for the quality difference. Consumers with higher willingness to pay are more likely to purchase the compliant product, while consumers with lower willingness to pay purchase the conventional product or choose the outside option. This structure generates endogenous substitution between product types as prices change.

When the regulation is introduced, consumers within the regulating jurisdiction can no longer purchase the conventional product. They must either switch to the compliant product or reduce their consumption of the regulated good. Consumers outside the jurisdiction continue to have access to both product types, though their purchasing decisions change as prices adjust.

1.4.1 Baseline Model: No Regulation

Supply

On the supply side, I assume each producer i can supply one unit of output, choosing between conventional production and restricted production: $q_i^C = 1$ or $q_i^R = 1$. Each producer incurs a base production cost r . Producers who choose restricted production face an additional incremental cost $\phi_i > 0$ per unit. Throughout the paper, I use the term restricted to denote products that comply with the production standard imposed by the regulation. Producers are heterogeneous in their compliance costs, which are assumed to be uniformly distributed:

$$\phi_i \sim \mathcal{U}(\underline{\phi}, \bar{\phi}), \quad (1.1)$$

where $\bar{\phi}$ and $\underline{\phi}$ are the upper and lower bounds of the incremental cost distribution. The cumulative distribution function is $F(\phi; \bar{\phi}, \underline{\phi}) = \frac{\phi - \underline{\phi}}{\bar{\phi} - \underline{\phi}}$ for $\phi \in [\underline{\phi}, \bar{\phi}]$, with probability density function $f(\phi)$.

The compliance cost ϕ captures both variable costs and the per-unit allocation of incremental fixed costs associated with adopting restricted production. I assume this cost distribution applies to producers both within and outside the regulating jurisdiction.

Let N denote the total mass of producers, with share η located within the jurisdiction I . Via normalization, I impose $0 < N \leq 1$ and $0 < \eta < 1$ to ensure market-clearing conditions hold.

Each producer chooses restricted production if and only if:

$$\begin{aligned} \pi_i = P_0^R - r - \phi_i &> P_0^C - r \\ \iff \phi_i &< P_0^R - P_0^C, \end{aligned} \quad (1.2)$$

where P_0^C and P_0^R denote baseline prices for conventional and restricted products. Define $\phi_0^* = P_0^R - P_0^C$ as the cutoff conversion cost at which a producer is indifferent between the two production modes.

In the baseline, the share of producers choosing conventional or restricted production is identical within and outside the regulating jurisdiction. Using the threshold $\phi_0^* = P_0^R - P_0^C$, the supply

of conventional products within the jurisdiction is:

$$\begin{aligned} Q_{I_0}^{S,C} &= \eta \int_{\phi_0^*}^{\bar{\phi}} N f(\phi) d\phi \\ &= \eta N \left(1 - \frac{P_0^R - P_0^C - \underline{\phi}}{\bar{\phi} - \underline{\phi}} \right). \end{aligned} \quad (1.3)$$

The supply of restricted products within the jurisdiction is:

$$\begin{aligned} Q_{I_0}^{S,R} &= \eta \int_{\underline{\phi}}^{\phi_0^*} N f(\phi) d\phi \\ &= \eta N \frac{P_0^R - P_0^C - \underline{\phi}}{\bar{\phi} - \underline{\phi}}. \end{aligned} \quad (1.4)$$

The supply of conventional products outside the jurisdiction is:

$$\begin{aligned} Q_{O_0}^{S,C} &= (1 - \eta) \int_{\phi_0^*}^{\bar{\phi}} N f(\phi) d\phi \\ &= (1 - \eta) N \left(1 - \frac{P_0^R - P_0^C - \underline{\phi}}{\bar{\phi} - \underline{\phi}} \right). \end{aligned} \quad (1.5)$$

The supply of restricted products outside the jurisdiction is:

$$\begin{aligned} Q_{O_0}^{S,R} &= (1 - \eta) \int_{\underline{\phi}}^{\phi_0^*} N f(\phi) d\phi \\ &= (1 - \eta) N \frac{P_0^R - P_0^C - \underline{\phi}}{\bar{\phi} - \underline{\phi}}. \end{aligned} \quad (1.6)$$

Here, subscripts I_0 and O_0 denote products supplied within and outside the jurisdiction in the baseline period.

Aggregate market supply is given by:

$$\begin{aligned} Q_0^{S,C} &= Q_{I_0}^{S,C} + Q_{O_0}^{S,C}, \\ Q_0^{S,R} &= Q_{I_0}^{S,R} + Q_{O_0}^{S,R}. \end{aligned} \quad (1.7)$$

Demand

I adopt a vertical differentiation framework based on [Shaked and Sutton \(1982\)](#) to model consumer demand. In this framework, products differ by quality, with a clear hierarchy: conventional products rank below restricted products in terms of perceived quality.

Consider consumers indexed by income t , uniformly distributed over the interval $[a, b]$ with density 1. Each consumer makes an indivisible, mutually exclusive purchase: either no purchase, one unit of conventional product, or one unit of restricted product. Consumers who purchase product $k \in \{C, R\}$ at price P^k obtain utility:

$$U(t, k) = u^k(t - P^k), \quad k \in \{C, R\} \quad (1.8)$$

where u^k is the constant marginal utility of income spent on the numeraire good, conditional on consuming product k . A higher value of u^k indicates higher product quality and a greater willingness to pay for product k . I normalize the utility parameter of the numeraire good to one, so that

$$1 < u^C < u^R.$$

Consumers who purchase neither product spend all income on the numeraire and receive reservation utility

$$U^M(t) = t. \quad (1.9)$$

Define:

$$D = \frac{u^R}{u^R - u^C}, \quad E = \frac{u^C}{u^C - 1}, \quad (1.10)$$

with $D > 1$ and $E > 1$, following the notation of [Shaked and Sutton \(1982\)](#). Both D and E are inversely related to the quality gap between adjacent products: D approaches 1 when restricted and conventional products differ substantially in quality, and increases without bound as their quality levels converge. Similarly, E is large when the conventional product is close in quality to the outside

option, implying that consumers are nearly indifferent between purchasing and not purchasing.

Consumer indifference conditions define income thresholds. A consumer with income t_0^{R*} is indifferent between purchasing restricted and conventional goods. A consumer with income t_0^{C*} is indifferent between purchasing conventional goods and the outside option:

$$\begin{aligned} u^R(t_0^{R*} - P_0^R) &= u^C(t_0^{R*} - P_0^C), \\ u^C(t_0^{C*} - P_0^C) &= t_0^{C*}. \end{aligned} \tag{1.11}$$

Solving these conditions yields:

$$\begin{aligned} t_0^{C*} &= EP_0^C, \\ t_0^{R*} &= P_0^C(1 - D) + P_0^R D. \end{aligned} \tag{1.12}$$




Figure 1.1: Utility as a function of income under vertical differentiation.

I normalize the total consumer population to 1 and denote by μ the share of consumers located within the jurisdiction, where $0 < \mu < 1$. Under the uniform income distribution, the demand for

conventional products within the jurisdiction is:

$$\begin{aligned}
Q_{I_0}^{D,C} &= \mu[F(t_0^{R*}) - F(t_0^{C*})] \\
&= \mu \left(\frac{t_0^{R*} - t_0^{C*}}{b - a} \right) \\
&= \mu \frac{D(P_0^R - P_0^C) - (E - 1) P_0^C}{b - a}.
\end{aligned} \tag{1.13}$$

The demand for restricted products within the jurisdiction is:

$$\begin{aligned}
Q_{I_0}^{D,R} &= \mu[1 - F(t_0^{R*})] \\
&= \mu \left(1 - \frac{t_0^{R*} - a}{b - a} \right) \\
&= \mu \frac{b - P_0^C - D(P_0^R - P_0^C)}{b - a}.
\end{aligned} \tag{1.14}$$

Similarly, demand for conventional products outside the jurisdiction is:

$$\begin{aligned}
Q_{O_0}^{D,C} &= (1 - \mu)[F(t_0^{R*}) - F(t_0^{C*})] \\
&= (1 - \mu) \left(\frac{t_0^{R*} - t_0^{C*}}{b - a} \right) \\
&= (1 - \mu) \frac{D(P_0^R - P_0^C) - (E - 1) P_0^C}{b - a}.
\end{aligned} \tag{1.15}$$

Demand for restricted products outside the jurisdiction is:

$$\begin{aligned}
Q_{O_0}^{D,R} &= (1 - \mu)[1 - F(t_0^{R*})] \\
&= (1 - \mu) \left(1 - \frac{t_0^{R*} - a}{b - a} \right) \\
&= (1 - \mu) \frac{b - P_0^C - D(P_0^R - P_0^C)}{b - a}.
\end{aligned} \tag{1.16}$$

Aggregate market demand is given by:

$$\begin{aligned} Q_0^{D,C} &= Q_{I_0}^{D,C} + Q_{O_0}^{D,C}, \\ Q_0^{D,R} &= Q_{I_0}^{D,R} + Q_{O_0}^{D,R}. \end{aligned} \tag{1.17}$$

Equilibrium Conditions before the Regulation

In the baseline setting without any sub-national regulation, equilibrium requires that, for each product type, market demand equals market supply and the common equilibrium price clears both sides of the market. Let asterisks denote equilibrium values. The equilibrium conditions for the conventional and restricted products are:

$$Q_0^{C*} = Q_0^{D,C*} = Q_0^{S,C*}, \quad P_0^{C*} = P_0^{D,C*} = P_0^{S,C*}, \tag{1.18}$$

$$Q_0^{R*} = Q_0^{D,R*} = Q_0^{S,R*}, \quad P_0^{R*} = P_0^{D,R*} = P_0^{S,R*}. \tag{1.19}$$

These conditions ensure market clearing for both conventional and restricted goods. Substituting the explicit demand and supply expressions derived in the previous sections yields the system:

$$\begin{aligned} N \left(1 - \frac{P_0^R - P_0^C - \underline{\phi}}{\bar{\phi} - \underline{\phi}} \right) &= \frac{D(P_0^R - P_0^C) - (E-1)P_0^C}{b-a}, \\ N \frac{P_0^R - P_0^C - \underline{\phi}}{\bar{\phi} - \underline{\phi}} &= \frac{b - P_0^C - D(P_0^R - P_0^C)}{b-a}. \end{aligned} \tag{1.20}$$

Equilibrium prices. Solving this system yields the equilibrium conventional price,

$$P_0^{C*} = \frac{b - (b-a)N}{E}. \tag{1.21}$$

Define the equilibrium price gap as $\Delta_0^* = P_0^{R*} - P_0^{C*}$. Solving for this gap gives

$$\Delta_0^* = \frac{(b-a)N [\bar{\phi} + (E-1)\underline{\phi}] + (\bar{\phi} - \underline{\phi})(E-1)b}{E [D(\bar{\phi} - \underline{\phi}) + N(b-a)]}. \tag{1.22}$$

The equilibrium restricted price then follows as

$$P_0^{R*} = P_0^{C*} + \Delta_0^*. \quad (1.23)$$

Equilibrium quantities. Equilibrium quantities follow from the supply shares implied by the compliance cost distribution:

$$\begin{aligned} Q_0^{R*} &= N \frac{\Delta_0^* - \underline{\phi}}{\bar{\phi} - \underline{\phi}} \\ &= \frac{N [(b-a)N + (E-1)b - DE\underline{\phi}]}{E [D(\bar{\phi} - \underline{\phi}) + N(b-a)]}. \end{aligned} \quad (1.24)$$

$$\begin{aligned} Q_0^{C*} &= N - Q_0^{R*} \\ &= \frac{N [DE\bar{\phi} + (E-1)(N(b-a) - b)]}{E [D(\bar{\phi} - \underline{\phi}) + N(b-a)]}. \end{aligned} \quad (1.25)$$

1.4.2 Post-Regulation Equilibrium

Supply

After the regulation, producers within the jurisdiction can no longer sell conventional products. Those with low enough compliance costs convert to restricted production; those facing higher costs exit. I assume the outside option for producers within the jurisdiction yields zero economic profit (a competitive return). Conventional producers convert to restricted production if and only if conversion generates nonnegative profit. Let P_1^R denote the post-regulation price of restricted products. A producer within the jurisdiction converts if:

$$\pi_{Ri} = P_1^R - r - \phi_i \geq 0 \quad \Longleftrightarrow \quad \phi_i \leq P_1^R - r. \quad (1.26)$$

Define $\phi_{I_1}^* = P_1^R - r$ as the threshold conversion cost at which a producer within the jurisdiction is indifferent between converting to restricted production and exiting the market. Producers with $\phi_i > \phi_{I_1}^*$ exit.

With the new threshold $\phi_{I_1}^* = P_1^R - r$, the supply of conventional products within the jurisdiction becomes:

$$Q_{I_1}^{S,C} = 0, \quad (1.27)$$

reflecting the regulation's prohibition on conventional production for sale within the jurisdiction.

The supply of restricted products within the jurisdiction is:

$$Q_{I_1}^{S,R} = \eta \int_{\underline{\phi}}^{\phi_{I_1}^*} N f(\phi) d\phi = \eta N \frac{P_1^R - r - \underline{\phi}}{\bar{\phi} - \underline{\phi}}. \quad (1.28)$$

The supply of conventional products outside the jurisdiction is:

$$\begin{aligned} Q_{O_1}^{S,C} &= (1 - \eta) \int_{\phi_{O_1}^*}^{\bar{\phi}} N f(\phi) d\phi \\ &= (1 - \eta) N \left(1 - \frac{P_1^R - P_1^C - \underline{\phi}}{\bar{\phi} - \underline{\phi}} \right). \end{aligned} \quad (1.29)$$

The supply of restricted products outside the jurisdiction is:

$$\begin{aligned} Q_{O_1}^{S,R} &= (1 - \eta) \int_{\underline{\phi}}^{\phi_{O_1}^*} N f(\phi) d\phi \\ &= (1 - \eta) N \frac{P_1^R - P_1^C - \underline{\phi}}{\bar{\phi} - \underline{\phi}}. \end{aligned} \quad (1.30)$$

Aggregate market supply is given by:

$$\begin{aligned} Q_1^{S,C} &= Q_{O_1}^{S,C}, \\ Q_1^{S,R} &= Q_{I_1}^{S,R} + Q_{O_1}^{S,R}. \end{aligned} \quad (1.31)$$

Demand

The regulation eliminates consumer access to conventional products within the jurisdiction. Demand for conventional products within the jurisdiction therefore becomes zero, and consumers face a binary choice between restricted products and the outside option.

For consumers within the jurisdiction, a consumer with income $t_{I_1}^{R*}$ is indifferent between purchasing restricted goods at price P_1^R and choosing the outside option:

$$u^R(t_{I_1}^{R*} - P_1^R) = t_{I_1}^{R*}. \quad (1.32)$$

Solving the equation above yields

$$t_{I_1}^{R*} = \frac{DE}{D+E-1} P_1^R. \quad (1.33)$$

For consumers outside the jurisdiction, as in the case before the regulation, consumer indifference conditions define income thresholds. A consumer with income $t_{O_1}^{R*}$ is indifferent between purchasing restricted goods at price P_1^R and purchasing conventional goods at price P_1^C . A consumer with income $t_{O_1}^{C*}$ is indifferent between purchasing conventional goods at price P_1^C and choosing the outside option. These thresholds satisfy

$$\begin{aligned} t_{O_1}^{C*} &= EP_1^C, \\ t_{O_1}^{R*} &= P_1^C(1-D) + P_1^R D. \end{aligned} \quad (1.34)$$

Demand for restricted products within the jurisdiction is:

$$\begin{aligned} Q_{I_1}^{D,R} &= \mu[1 - F(t_{I_1}^{R*})] \\ &= \mu \frac{b - \frac{DE}{D+E-1} P_1^R}{b-a}. \end{aligned} \quad (1.35)$$

Consumers outside the jurisdiction continue to have access to both conventional and restricted products. Demand for conventional products outside the jurisdiction is:

$$\begin{aligned} Q_{O_1}^{D,C} &= (1-\mu)[F(t_{O_1}^{R*}) - F(t_{O_1}^{C*})] \\ &= (1-\mu) \frac{D(P_1^R - P_1^C) - (E-1)P_1^C}{b-a}. \end{aligned} \quad (1.36)$$

Demand for restricted products outside the jurisdiction is:

$$\begin{aligned} Q_{O_1}^{D,R} &= (1 - \mu)[1 - F(t_{O_1}^{R*})] \\ &= (1 - \mu) \frac{b - P_1^C - D(P_1^R - P_1^C)}{b - a}. \end{aligned} \quad (1.37)$$

Aggregate market demand is given by:

$$\begin{aligned} Q_1^{D,C} &= Q_{O_1}^{D,C}, \\ Q_1^{D,R} &= Q_{I_1}^{D,R} + Q_{O_1}^{D,R}. \end{aligned} \quad (1.38)$$

This structure captures three margins of consumer adjustment: remaining with the same product, switching across products, and exiting the market. Because the regulation affects prices and product availability differently across regions, the composition of these responses varies within and outside the jurisdiction.

Equilibrium Conditions after the Regulation

After the regulation, market equilibrium is determined by the clearing conditions in the conventional and restricted markets:

$$Q_1^{C*} = Q_1^{D,C*} = Q_1^{S,C*}, \quad P_1^{C*} = P_1^{D,C*} = P_1^{S,C*}, \quad (1.39)$$

$$Q_1^{R*} = Q_1^{D,R*} = Q_1^{S,R*}, \quad P_1^{R*} = P_1^{D,R*} = P_1^{S,R*}. \quad (1.40)$$

Define the post-regulation equilibrium price gap as $\Delta_1^* = P_1^{R*} - P_1^{C*}$.

Equating supply and demand in the conventional market yields

$$A_1 P_1^C + B_1 \Delta_1 = K_1, \quad (1.41)$$

where

$$\begin{aligned}
A_1 &= -(1 - \mu)(E - 1), \\
B_1 &= (1 - \mu)D + \frac{(b - a)(1 - \eta)N}{\bar{\phi} - \underline{\phi}}, \\
K_1 &= \frac{(b - a)(1 - \eta)N\bar{\phi}}{\bar{\phi} - \underline{\phi}}.
\end{aligned} \tag{1.42}$$

Similarly, equating supply and demand in the restricted market yields

$$A_2 P_1^C + B_2 \Delta_1 = K_2, \tag{1.43}$$

where I define

$$H = \frac{DE}{D + E - 1},$$

and

$$\begin{aligned}
A_2 &= (1 - \mu) + \mu H + \frac{(b - a)\eta N}{\bar{\phi} - \underline{\phi}}, \\
B_2 &= (1 - \mu)D + \mu H + \frac{(b - a)N}{\bar{\phi} - \underline{\phi}}, \\
K_2 &= b + \frac{(b - a)N(\eta r + \underline{\phi})}{\bar{\phi} - \underline{\phi}}.
\end{aligned} \tag{1.44}$$

Equilibrium prices. Equations (1.41) and (1.43) form a linear system in (P_1^C, Δ_1) . Let

$$\Omega = A_1 B_2 - A_2 B_1.$$

Then the equilibrium price gap and conventional price are given by

$$\Delta_1^* = \frac{A_1 K_2 - A_2 K_1}{\Omega}, \quad P_1^{C*} = \frac{K_1 B_2 - K_2 B_1}{\Omega}. \tag{1.45}$$

The restricted price therefore satisfies

$$P_1^{R*} = P_1^{C*} + \Delta_1^*.$$

Equilibrium quantities. Substituting equilibrium prices into the supply equations yields equilibrium quantities:

$$Q_1^{C*} = (1 - \eta)N \frac{\bar{\phi} - \Delta_1^*}{\bar{\phi} - \underline{\phi}}, \quad (1.46)$$

$$Q_1^{R*} = \eta N \frac{P_1^{R*} - r - \underline{\phi}}{\bar{\phi} - \underline{\phi}} + (1 - \eta)N \frac{\Delta_1^* - \underline{\phi}}{\bar{\phi} - \underline{\phi}}. \quad (1.47)$$

1.4.3 Social Welfare

Compensating Variation

I evaluate consumer welfare using compensating variation (CV), defined as the monetary transfer required to make consumers indifferent between the pre- and post-regulation environments (Varian, 1992; Small and Rosen, 1981). In the vertical differentiation model, the regulation changes both available products and prices, inducing consumers to exit the market, switch between product types, or continue consuming at different prices. Compensating variation captures all of these behavioral responses in a single money-metric welfare measure (Trajtenberg, 1989; Berry, Levinsohn and Pakes, 1995).

Consumer income t is uniformly distributed on (a, b) with density $f(t) = \frac{1}{b-a}$. For a consumer who consumed product $j \in \{C, R\}$ before the regulation and chooses option $k \in \{M, C, R\}$ after the regulation, individual compensating variation $CV(t) = \tilde{t} - t$ is the income transfer satisfying

$$U(t - P_0^j, j) = U(\tilde{t} - P_1^k, k).$$

Solving this equation for each behavioral transition yields the individual CV formulas summarized in Table 1.2.

The derivations of these formulas are provided in Appendix 1.9.2.

Within the Jurisdiction. Before the regulation, consumers with $t < t_0^{C*}$ choose the outside option, those with $t \in (t_0^{C*}, t_0^{R*})$ consume conventional goods, and those with $t > t_0^{R*}$ consume

Table 1.2: Individual Compensating Variation by Behavioral Transition

Transition	$CV_i(t) = \tilde{t} - t$
$C \rightarrow M$ (exiter)	$\frac{t - EP_0^C}{E - 1}$
$C \rightarrow R$ (converter)	$-\frac{1}{D}t - \frac{D - 1}{D}P_0^C + P_1^R$
$R \rightarrow R$ (restricted incumbent)	$P_1^R - P_0^R$
$C \rightarrow C$ (conventional incumbent)	$P_1^C - P_0^C$
$R \rightarrow C$ (converter)	$\frac{t - DP_0^R}{D - 1} + P_1^C$
$M \rightarrow R$ (entrant)	$-\frac{D + E - 1}{DE}t + P_1^R$

restricted goods. After the regulation, conventional goods are unavailable within the jurisdiction, so consumers either choose the outside option ($t < t_1^{R*}$) or consume restricted goods ($t > t_1^{R*}$). The relevant thresholds are $t_0^{C*} = EP_0^C$, $t_0^{R*} = P_0^C(1 - D) + P_0^R D$, and $t_1^{R*} = \frac{DE}{D+E-1}P_1^R$.

Depending on the ordering of these thresholds, three cases arise, each generating a different partition of consumers into behavioral groups. For instance, when $t_0^{C*} < t_1^{R*} < t_0^{R*}$ (the empirically relevant case for Proposition 12), former conventional consumers with incomes below t_1^{R*} exit the market, those above t_1^{R*} switch to restricted goods, and restricted incumbents face a price change of $P_1^R - P_0^R$. Aggregate CV within the jurisdiction is obtained by integrating individual CV over each group's income interval and summing. The full case-by-case decomposition and integral calculations are presented in Appendix 1.9.2.

Outside the Jurisdiction. Outside the jurisdiction, both products remain available after the regulation, but at new prices (P_1^C, P_1^R) . The post-regulation thresholds $t_1^{C*} = EP_1^C$ and $t_1^{R*} = P_1^C(1 - D) + P_1^R D$ shift relative to their pre-regulation counterparts, generating exiters, conventional incumbents, converters ($C \rightarrow R$), converters ($R \rightarrow C$), and restricted incumbents. Three cases again arise depending on the ordering of pre- and post-regulation thresholds. Aggregate CV outside the jurisdiction is computed analogously. Details are in Appendix 1.9.2.

To benchmark the magnitude of compensating variation, I define baseline market value as

$$MV_0 = P_0^C Q_0^C + P_0^R Q_0^R. \quad (1.48)$$

Comparing CV_I and CV_O against MV_0 provides a scale-free measure of the welfare impact relative to baseline market size.

Producer Surplus

Producer welfare is measured by changes in producer surplus. Producers are heterogeneous in conversion costs $\phi \sim U[\underline{\phi}, \bar{\phi}]$, and each supplies one unit of output. Before the regulation, a producer chooses restricted production if $\phi \leq \phi_0^* = P_0^R - P_0^C$ and conventional production otherwise.

Within the jurisdiction, the regulation eliminates conventional production. Producers with conversion costs below $\phi_{I_1}^* = P_1^R - r$ remain active as restricted producers; those above this threshold exit. This creates three groups: restricted incumbents ($\phi < \phi_0^*$), who gain or lose depending on the price change net of any compliance cost; converters ($\phi_0^* < \phi < \phi_{I_1}^*$), who switch from conventional to restricted production; and exiters ($\phi > \phi_{I_1}^*$), who lose their pre-regulation conventional profit.

Outside the jurisdiction, both technologies remain available and the post-regulation cutoff shifts to $\phi_{O_1}^* = P_1^R - P_1^C$. Restricted incumbents gain or lose according to the change in the restricted price, conventional incumbents gain or lose according to the change in the conventional price, and converters switch to restricted production at heterogeneous gains depending on their ϕ . The direction and magnitude of price changes depend on the relative shifts of supply and demand in each market.

Aggregate producer surplus changes are obtained by integrating individual surplus changes over the relevant ϕ intervals. The closed-form expressions are presented in Appendix [1.9.3](#).

Total Social Welfare

Total welfare change combines consumer and producer effects:

$$\Delta W = \Delta PS - CV, \tag{1.49}$$

where $CV = CV_I + CV_O$ is total compensating variation (positive values indicate consumer losses) and $\Delta PS = \Delta PS_I + \Delta PS_O$ is the total change in producer surplus. This decomposition allows welfare impacts to be traced to both demand-side adjustments (market exit, product substitution, and price changes for incumbents) and supply-side responses (technology switching, compliance costs, and producer exit). The closed-form expressions are evaluated numerically in Section 3.6 using calibrated parameters.

1.5 Model Calibration

The framework developed in Section 1.4 applies to any consumption-based production standard imposed by a jurisdiction that is both a consumer and a producer of the regulated good. To illustrate the model’s mechanisms in a concrete setting, I calibrate it to California’s Proposition 12 and the U.S. shell egg market. The calibration should be interpreted as an illustrative application of the framework rather than a precise quantitative evaluation of the policy. The calibrated parameters are chosen to match key empirical moments and are broadly consistent with the existing literature where available, with remaining parameters guided by economic intuition and industry evidence.

1.5.1 Application: California’s Egg Market

California’s egg market provides an ideal application of the general framework for several reasons. First, the market exhibits the key structural feature that the model is designed to study: the regulating jurisdiction is a significant consumer but a relatively small producer of the regulated good. California accounts for approximately 12% of national egg consumption but only about 3–4%

of national egg production, making it a large net importer. This asymmetry between consumption and production shares, captured in the model by $\mu = 0.12$ and $\eta = 0.03$, means that the regulation affects both domestic production decisions and interstate trade flows.

Second, the egg market features clear vertical differentiation between conventional and cage-free products, matching the model’s distinction between conventional and restricted goods. Prior to Proposition 12, cage-free eggs sold at a substantial premium over conventional eggs, and the two products coexisted with cage-free eggs accounting for a nontrivial but minority share of the market. This structure maps directly onto the model’s demand system, in which heterogeneous consumers sort across products based on income and willingness to pay for quality.

Third, producers face heterogeneous compliance costs. Some egg producers already operated cage-free systems before the regulation and could comply at little additional cost, while others relied on conventional cage housing and faced substantial capital expenditure to convert (Sumner et al., 2008; Caputo et al., 2023a). This heterogeneity is central to the model’s supply side, where the distribution of conversion costs $[\underline{\phi}, \bar{\phi}]$ determines which producers convert, which exit, and how surplus is redistributed.

Fourth, Proposition 12 is a clean example of a consumption-based standard: beginning January 1, 2022, it prohibited the sale of eggs from caged housing in California, regardless of where the eggs are produced. Combined with Proposition 2, which had already banned caged egg production within California, the regulatory environment prohibits both the production and sale of conventional eggs in the state.

The calibration follows standard quantitative equilibrium approaches used in agricultural policy analysis (e.g., Brester, Marsh and Atwood, 2004; Saitone, Sexton and Sumner, 2015), in which structural parameters are chosen to reproduce observed prices and quantities in the pre-policy equilibrium, and policy effects are evaluated by changing regulatory constraints while holding preferences and production technologies fixed. I take 2018 as the baseline year, when Proposition 12 had been passed but its requirements were not yet in force.

1.5.2 Calibration Strategy

The model has eight free parameters: the income bounds a and b , the quality ratios D and E , the total producer mass N , the conversion cost bounds $\underline{\phi}$ and $\overline{\phi}$, and the base production cost r . I calibrate these parameters jointly by matching the model-implied pre-regulation equilibrium to empirical targets drawn from the 2018 U.S. shell egg market.

Rather than targeting absolute price and quantity levels, which are sensitive to unit definitions and measurement error, I target dimensionless ratios and elasticities. Table 1.3 summarizes the calibration targets and their acceptable ranges.

Table 1.3: Calibration Targets

Target	Range	Description
P_0^R/P_0^C	[1.7, 2.5]	Restricted-to-conventional price ratio
Q_0^C/Q_0^R	[2.5, 4.5]	Conventional-to-restricted quantity ratio
$ \varepsilon_C^D $	[0.1, 1.0]	Conventional demand elasticity
$ \varepsilon_R^D $	[0.5, 5.0]	Restricted demand elasticity
ε_C^S	< 1	Conventional supply elasticity
ε_R^S	< 5	Restricted supply elasticity (soft)

The price and quantity ratio targets are based on USDA Agricultural Marketing Service price data and USDA National Agricultural Statistics Service production data for 2018. The elasticity targets are drawn from the empirical literature on egg demand and supply.

1.5.3 Fixed Parameters and Constraints

Jurisdictional shares. I treat California as part of an integrated national shell egg market. The producer share is set to $\eta = 0.03$, based on California’s share of U.S. egg production from USDA National Agricultural Statistics Service data. The consumer share is set to $\mu = 0.12$, based on California’s share of the U.S. population from Census data. The outside option is normalized so that non-purchasers retain their full income, as described in Section 1.4.

Constraints. I impose the following economically motivated constraints to ensure a well-behaved equilibrium:

- $a > 0$: positive income lower bound;
- $D > 1$: restricted eggs are strictly preferred to conventional ($u^R > u^C$);
- $E > 1$: conventional eggs are strictly preferred to the outside option ($u^C > 1$);
- $r < P_0^C$: producers earn non-negative profit in the baseline;
- Threshold ordering $a < t_C^* < t_R^* < b$, so that both products have positive demand;
- $\underline{\phi} < \Delta < \bar{\phi}$, so that some but not all producers convert in equilibrium;
- Post-regulation feasibility: the policy equilibrium must exist with positive prices and quantities.

1.5.4 Data

I take 2018 as the baseline year, a period in which Proposition 12 had been passed but its production requirements had not yet come into force, and before large-scale anticipatory investments in cage-free capacity were widespread. Data on production quantities are drawn from the USDA National Agricultural Statistics Service, price data from the USDA Agricultural Marketing Service, and population shares from the U.S. Census Bureau.

1.5.5 Estimation Procedure

For any candidate parameter vector $\theta = (a, b, D, E, N, \underline{\phi}, \bar{\phi}, r)$, the model delivers equilibrium prices and quantities

$$(P_0^C(\theta), P_0^R(\theta), Q_0^C(\theta), Q_0^R(\theta)),$$

through the closed-form expressions derived in Section 1.4.1. Because these expressions are analytical, I do not need to solve for equilibrium numerically at each iteration of the calibration.

I calibrate θ by minimizing a weighted loss function that penalizes deviations from the target ranges:

$$\hat{\theta} = \arg \min_{\theta} \sum_k w_k \cdot \left[\max \left(0, \frac{m_k(\theta) - \bar{m}_k}{\bar{m}_k - \underline{m}_k} \right)^2 + \max \left(0, \frac{\underline{m}_k - m_k(\theta)}{\bar{m}_k - \underline{m}_k} \right)^2 \right],$$

where $m_k(\theta)$ is the k -th model-implied moment, $[\underline{m}_k, \bar{m}_k]$ is its target range, and w_k are weights reflecting the relative importance of each target. Model-implied moments that fall within their target ranges incur zero loss. Constraint violations are penalized via large additive terms to ensure feasibility.

I solve this optimization problem using differential evolution (Storn and Price, 1997), a global optimization algorithm well-suited to problems with non-convex objectives and bound constraints. The algorithm is run 12 times in parallel from different random starting values, each with a population size of 50 and a maximum of 3,000 iterations, and the solution with the lowest loss across all runs is selected.

1.5.6 Model Fit

The calibrated parameters are reported in Table 1.4.

Table 1.4: Calibrated Parameters

Parameter	Value	Description
a	5.59	Income lower bound
b	199.77	Income upper bound
D	15.26	Quality ratio $u^R/(u^R - u^C)$
E	1.15	Quality ratio $u^C/(u^C - 1)$
N	0.99	Producer mass
$\underline{\phi}$	5.80	Conversion cost lower bound
$\bar{\phi}$	17.29	Conversion cost upper bound
r	1.96	Base production cost

The large value of $D = 15.26$ implies that the quality gap between restricted and conventional eggs ($u^R - u^C$) is small relative to the level of restricted quality u^R , consistent with vertical differentiation where both products are close substitutes in absolute quality but command different prices due to marginal consumer heterogeneity. The value $E = 1.15$ indicates that conventional eggs

are modestly preferred to the outside option, so that a moderate price increase can shift marginal consumers out of the market.

The calibrated model matches all targeted baseline moments within their acceptable ranges. The price ratio $P_0^R/P_0^C = 2.50$ and the quantity ratio $Q_0^C/Q_0^R = 2.50$. The implied demand elasticities are $\varepsilon_C^D = -0.68$ for conventional eggs and $\varepsilon_R^D = -4.19$ for restricted eggs. Supply elasticities are $\varepsilon_C^S = 0.74$ and $\varepsilon_R^S = 4.61$.

These elasticities are broadly consistent with estimates from the empirical literature on egg demand. [Bakhtavoryan et al. \(2021\)](#) estimate Marshallian own-price elasticities of -0.23 for conventional eggs and -1.00 for cage-free eggs using household scanner data. [Caputo et al. \(2023b\)](#) report own-price elasticities of -0.74 for conventional and -5.84 for cage-free from a choice experiment. Both studies find that cage-free demand is substantially more elastic than conventional demand, the same qualitative pattern as in my calibration. My restricted elasticity of -4.19 lies between these estimates, reflecting a setting in which consumers exhibit both strong quality preferences and meaningful substitution possibilities.

This ranking has a natural interpretation in the vertical differentiation framework. Cage-free buyers are willing to pay a premium for the higher-quality product, but they have a ready fallback: if the premium becomes too large, they can switch to conventional eggs. This close substitute makes their demand relatively elastic. Conventional buyers face a weaker substitution margin. They are unlikely to trade up to cage-free eggs when the premium shrinks, since by revealed preference they have already chosen the lower-quality option. Their primary adjustment margin is instead to reduce consumption, which makes demand less price-responsive.

On the supply side, [Bakhtavoryan et al. \(2021\)](#) use a domestic egg supply elasticity of 0.94 in their tariff simulations. My conventional supply elasticity of 0.74 is in line with this estimate. The restricted supply elasticity of 4.61 is substantially higher, reflecting that cage-free production capacity was expanding and that entry into cage-free production is more responsive to price signals than in the mature conventional segment ([Sumner et al., 2008](#)).

The base production cost $r = 1.96$ is calibrated jointly with the other parameters. Because

r enters both the pre- and post-regulation producer profit expressions symmetrically, it does not affect relative price or quantity ratios in the pre-regulation equilibrium but determines the level of absolute profits and thus the welfare calculations in Section 3.6.

The quantitative results depend on parameter choices, but Section 3.7 shows through a Monte Carlo analysis that the qualitative patterns are robust across a wide range of parameter values.

1.6 Simulation Results

Using the calibrated parameter vector $\hat{\theta}$, I simulate the effects of Proposition 12 by imposing the ban on conventional egg production and consumption in California and solving for the new market equilibrium. I compare outcomes before and after the policy in terms of prices, quantities, trade flows, and welfare for consumers and producers, both within California and in the rest of the United States.

Throughout this section, the before Prop-12 equilibrium corresponds to the 2018 baseline calibration, while the after Prop-12 equilibrium reflects the counterfactual economy in which only cage-free eggs can be produced and consumed in California.

1.6.1 Prices and Quantities

Table 1.5 reports equilibrium prices and quantities before and after Proposition 12 in the calibrated model. Because the calibration targets price and quantity ratios rather than absolute levels, the baseline prices and quantities in the table are expressed in *model units* rather than exactly matching the dollar values and production levels observed in the data.

For reference, in 2018 conventional eggs sold for approximately \$1.28 per dozen and cage-free eggs for about \$2.50 per dozen in the data, with total national production of roughly 89 billion conventional eggs and 20 billion cage-free eggs. My calibration does not directly pin down these absolute levels; instead, it disciplines the relative price P_0^R/P_0^C and the market share $Q_0^R/(Q_0^C+Q_0^R)$, which are the key objects that matter for substitution, welfare, and incidence in the model.

Given this normalization, the baseline values in Table 1.5 should be interpreted as model-consistent magnitudes that preserve the empirical price premium and market shares, rather than literal dollar prices or physical quantities.

After imposing the Proposition 12 constraint, both prices increase in the baseline calibration. The model implies that the price of conventional eggs rises by 5.7%, while the price of restricted eggs increases by 5.0%. As Section 3.7 shows, the restricted price increase is robust across parameter perturbations, while the conventional price effect is more sensitive to parameter values. Quantities move in opposite directions: conventional egg production falls by 7.9%, while restricted production expands by 16.8%. Total market output declines by only 0.9%, indicating that the regulation primarily reshuffles production across technologies rather than reducing total supply.

To understand these results, it is useful to trace the supply and demand shifts within and outside the jurisdiction.

Within California, Proposition 12 eliminates the market for conventional eggs. Conventional demand falls to zero, and in-state producers must either convert to cage-free production or exit. At the same time, California consumers are forced to switch to restricted eggs or leave the egg market entirely. Most consumers switch, so restricted demand increases sharply, although some exit due to higher prices.

Outside California, there is no direct demand shift, since the regulation applies only to sales within the state. However, the increase in the restricted price raises the premium over conventional eggs, creating incentives for producers to convert to cage-free production in order to serve the California market. This voluntary conversion reduces conventional supply and increases restricted supply at the national level.

For conventional eggs, both demand and supply shift to the left. Demand falls because California accounts for roughly 12% of national consumption. Supply also falls, as California producers exit or convert and producers outside the state reallocate capacity toward restricted production. Ex ante, the effect on the conventional price is ambiguous, since both demand and supply contract. In the calibrated equilibrium, however, the supply contraction dominates the demand reduction, so

the conventional price rises.

For restricted eggs, both demand and supply shift to the right. Demand increases sharply due to the California mandate. Supply expands through conversion, but because conversion costs are heterogeneous, not all producers can profitably convert. As a result, the increase in supply falls short of the increase in demand, and the restricted price rises.

The price premium for restricted eggs also widens. In the baseline, restricted eggs are more expensive than conventional eggs by 9.08 model units. After Proposition 12, this gap increases to 9.50 model units, a rise of about 4.6%. The larger premium reflects both higher compliance costs for producers converting to cage-free systems and the increased effective demand for restricted eggs induced by the California regulation.

Table 1.5: Impacts of Proposition 12 on Egg Prices, Quantities, and Price Premium

Variable	Unit	Baseline	Prop 12	% Change
Prices				
Price, conventional eggs (P^C)	model unit	6.05	6.40	+5.7
Price, restricted eggs (P^R)	model unit	15.14	15.90	+5.0
Quantities				
Quantity, conventional eggs (Q^C)	model unit	0.71	0.65	-7.9
Quantity, restricted eggs (Q^R)	model unit	0.28	0.33	+16.8
Total quantity ($Q^C + Q^R$)	model unit	0.99	0.98	-0.9
Price Premium				
Premium ($P^R - P^C$)	model unit	9.08	9.50	+4.6

Notes: “Base” reports the calibrated pre-policy equilibrium and “Prop 12” reports the counterfactual equilibrium under the Proposition 12 constraint. Prices and quantities are reported in model units (as in the equilibrium solution). Percent changes are computed relative to the baseline.

1.6.2 Elasticities

Table 1.6 reports the calibrated own-price demand and supply elasticities before and after Proposition 12. In the baseline, demand for restricted eggs is substantially more elastic (-4.19) than for conventional eggs (-0.68), reflecting its smaller market share and closer substitutability with conventional eggs. Supply is also more elastic for restricted eggs (4.61 vs. 0.74), consistent

with producers facing heterogeneous but adjustable conversion costs.

After Proposition 12, conventional eggs are prohibited in California, so own-price elasticities for that product are not defined within the state. Within California, the restricted demand elasticity is -0.83 , reflecting the binary choice between restricted eggs and the outside option. Outside California, both demand and supply elasticities remain close to their pre-policy values, indicating that the regulation primarily affects behavior within California rather than fundamentally altering national market responses.

Table 1.7 reports cross-price elasticities between conventional and restricted eggs. Positive demand cross-elasticities indicate that the two products are substitutes from the consumer side. Negative supply cross-elasticities reflect reallocation of production across technologies: when the price of one product increases, producers shift production toward that product and away from the other.

Table 1.6: Demand and Supply Elasticities Before and After Proposition 12

	Conventional		Restricted	
	Demand	Supply	Demand	Supply
Base (National Market)				
Own-price elasticity	-0.68	0.74	-4.19	4.61
Prop 12 (California Market)				
Own-price elasticity	N/A	N/A	-0.83	1.95
Prop 12 (Outside California)				
Own-price elasticity	-0.78	0.82	-5.70	4.30

Notes: Demand elasticities ε^D and supply elasticities ε^S are computed from the calibrated model. In California after Proposition 12, conventional eggs are prohibited, so own-price elasticities are not defined for that product.

1.6.3 Trade and Market Integration

California has historically been a net importer of eggs. In 2018, prior to the implementation of Proposition 12, imports already covered approximately 62% of California’s egg consumption. Section 1.6.7 documents how this dependence has strengthened since the regulation took effect.

Table 1.7: Cross-Price Elasticities Before and After Proposition 12

	Conventional	Restricted
Base (National Market)		
Demand (cross)	1.68	1.57
Supply (cross)	−1.84	−1.84
After Prop 12 (Outside California)		
Demand (cross)	1.91	2.15
Supply (cross)	−2.04	−1.73

Notes: Cross-price elasticities measure the percentage change in quantity of one product in response to a 1% change in the price of the other product.

Table 1.8 reports the model-implied trade flows for California before and after Proposition 12. The values in the table are expressed in model units, as they are generated by the calibrated equilibrium rather than taken directly from administrative trade data. They should therefore be interpreted as capturing the direction and relative magnitude of trade adjustments rather than exact quantitative counterparts in the data.

In the baseline equilibrium, the model implies that California is a substantial net importer of both conventional and restricted eggs. Net imports account for approximately three-quarters of consumption in both product categories, consistent with the evidence that California relies heavily on out-of-state supply.

After Proposition 12, the composition of trade changes significantly. Conventional eggs are eliminated from the California market by construction, so conventional imports fall to zero. In contrast, net imports of restricted (cage-free) eggs increase sharply, rising from 0.0255 to 0.0912 in model units. The implied import share for restricted eggs increases from 75% to about 81% of California consumption. This reflects the limited ability of in-state producers to rapidly expand cage-free production in response to the mandate, combined with strong demand within California.

As a result, a substantial portion of the adjustment to Proposition 12 occurs outside California. Producers in other states expand cage-free output and redirect shipments toward California to take advantage of higher prices. This pattern is consistent with observed developments after 2022, when

California became increasingly reliant on out-of-state cage-free production to satisfy Proposition 12 requirements.

Table 1.8: Trade Effects of Proposition 12 on California Egg Market

Product	California Market			
	Demand	Supply	Net Imports	Import Share(%)
Base				
Conventional	0.09	0.02	0.06	75.0
Restricted	0.03	0.01	0.03	75.0
After Prop 12				
Conventional	0.00	0.00	0.00	—
Restricted	0.11	0.02	0.09	81.2

Notes: Net imports are defined as Demand minus Supply in California. Import share is the ratio of net imports to demand. After Proposition 12, conventional eggs are prohibited in California.

1.6.4 Consumer Welfare

I measure changes in consumer welfare using compensating variation (CV), defined as the monetary transfer that would leave consumers at their baseline utility level after the policy. I compute welfare separately for distinct consumer groups based on their consumption choices before and after Proposition 12.

California consumers. Consumer welfare in California declines by a total CV of 0.666, or about 7.8% of baseline market value ($MV_0 = 8.59$). The loss is distributed across three groups. Exiters make up 5.7% of California consumers and account for 3.0% of MV_0 in aggregate CV. Converters are the largest group (65.2%) and contribute 4.5% of MV_0 in CV, as they pay a higher price for cage-free eggs and lose surplus from the product conversion. Restricted incumbents (28.4%) are least affected, with CV of 0.3% of MV_0 from the cage-free price increase alone.

Outside California. Outside California, total CV is 0.403, or 4.7% of MV_0 . Most consumers keep purchasing the same product at a somewhat higher price. Conventional incumbents lose 2.5%

of MV_0 from higher conventional prices, and restricted incumbents lose 1.9% of MV_0 from the cage-free price increase. A small share of consumers switch products or exit the market.

National consumers. At the national level, total consumer CV is 1.068, or 12.4% of MV_0 . Most of this loss is in California, where consumers lose access to conventional eggs entirely. Outside the state, the welfare effects come mainly from moderate price increases. These welfare results reflect different margins of consumer adjustment, including product switching and market exit, which are central to how the model maps price changes into heterogeneous consumer losses.

Figures 1.2 and 1.3 illustrate how consumer segments shift before and after the regulation, within and outside the jurisdiction respectively.

Figure 1.2: Change in income threshold within the jurisdiction (California)




Figure 1.3: Change in income threshold outside the jurisdiction




Table 1.9: Consumer Welfare (CV) Effects of Proposition 12

Region & Group	Consumer Welfare Changes		
	Consumer Share (%)	Group CV	CV / MV_0 (%)
California (CA)			
Exiters (C $\rightarrow$ Exit)	5.7	0.26	3.0
Converters (C $\rightarrow$ R)	65.2	0.38	4.5
Restricted incumbents (R $\rightarrow$ R)	28.4	0.03	0.3
<i>CA Total</i>	100.0	0.67	7.8
Outside California			
Exiters (C $\rightarrow$ Exit)	0.4	0.00	0.0
Conventional incumbents (C $\rightarrow$ C)	65.8	0.21	2.5
Converters (R $\rightarrow$ C)	4.7	0.02	0.2
Restricted incumbents (R $\rightarrow$ R)	28.4	0.17	1.9
<i>Outside-CA Total</i>	100.0	0.40	4.7
National Total			
All consumers	100.0	1.07	12.4

Notes: “Group CV” denotes the integrated compensating variation for each consumer group. CV/MV_0 is expressed as a percentage of baseline market value ($MV_0 = 8.59$).

1.6.5 Producer Surplus

Changes in producer surplus vary substantially across regions and producer types, reflecting differences in exposure to Proposition 12 and in the ability to adjust production technology.

California producers. Total producer surplus in California falls by about 36.8%, despite an increase in total industry revenue in the state. Restricted incumbents, who already operate cage-free systems, see their surplus rise by 13.3% due to higher cage-free prices. Converters, however, experience a 40.7% decline because conversion costs largely offset the price gains. Exiters lose all of their surplus as they leave production entirely. On net, the losses from converters and exiters outweigh the gains to incumbents. California’s conventional egg sector effectively disappears under the regulation, and the expansion of cage-free production does not generate enough additional profit


to compensate.

Outside California. Outside California, total producer surplus rises by about 10.4%. All three groups gain: restricted incumbents benefit from higher national cage-free prices (up 13.3%), converters gain by responding to the increased cage-free demand from California, and conventional incumbents gain 8.4% as reduced conventional supply from California raises prices nationwide.

National producers. At the national level, total producer surplus increases by about 8.9%, from 4.53 to 4.93. The gains, however, go mostly to producers outside California, while California producers as a whole lose surplus.

Figures 1.4 and 1.5 illustrate how the conversion cost threshold shifts for producers within and outside the jurisdiction, respectively.


Figure 1.4: Change in conversion cost threshold within the jurisdiction (California)



1.6.6 Overall Welfare Effects

Combining consumer and producer surplus, total economic welfare declines under Proposition 12. Total consumer compensating variation is 1.068 (12.4% of MV_0), while total producer surplus increases by 0.405 (4.7% of MV_0). The net welfare change is therefore -0.663 , equivalent to a decline of about 7.7% of baseline market value. Although the policy increases industry revenue

Figure 1.5: Change in conversion cost threshold outside the jurisdiction



and benefits producers on aggregate in the baseline calibration, these gains are far from sufficient to offset consumer losses.

The distribution of welfare impacts is highly uneven. Within California, consumers bear the majority of the adjustment costs: CV equals 0.666, or 7.8% of MV_0 (Table 1.9), driven by a combination of exiters who lose access to conventional eggs and converters who must switch to higher-priced cage-free eggs. California producers also lose surplus on net (−36.8%, Table 1.10), despite higher revenues, because conversion costs and exit from conventional production outweigh gains to cage-free incumbents. The state becomes even more reliant on out-of-state supply. Outside California, consumer losses are smaller (CV = 0.403, or 4.7% of MV_0), reflecting price increases in the baseline calibration. Producers outside California, by contrast, gain (+10.4%) as cage-free prices rise and out-of-state cage-free production expands to serve California demand.

These results illustrate a central tradeoff of Proposition 12. While the policy successfully shifts production toward cage-free systems and raises industry revenue, it does so at a substantial cost to consumers, particularly those in California, and at the expense of overall economic efficiency. In effect, the policy induces sizable transfers: from California consumers to producers, and from California to producers in other states, without generating offsetting efficiency gains.

The decline in total surplus operates through two main channels. First, the prohibition of conventional eggs in California removes a lower-cost option that many consumers preferred, leading

Table 1.10: Producer Surplus (PS) Effects of Proposition 12

Region & Group	Producer Surplus Changes			
	Producer Share (%)	PS before	PS after	% change
California (CA)				
Restricted incumbents (R $\rightarrow$ R)	28.6	0.05	0.06	+13.3
Converters (C $\rightarrow$ R)	42.3	0.05	0.03	-40.7
Exiters (C $\rightarrow$ exit)	29.2	0.04	0.00	-100.0
<i>Total PS (CA)</i>	100.0	0.14	0.09	-36.8
Outside California				
Restricted incumbents (R $\rightarrow$ R)	28.6	1.58	1.79	+13.3
Converters (C $\rightarrow$ R)	3.6	0.13	0.16	+19.3
Conventional incumbents (C $\rightarrow$ C)	67.8	2.67	2.90	+8.4
<i>Total PS (outside CA)</i>	100.0	4.39	4.85	+10.4
National Total				
All producers	100.0	4.53	4.93	+8.9

to either higher expenditures or complete exit from the market. Second, substantial producer conversion costs are incurred to comply with the regulation, which represent real resource losses that are not compensated by corresponding gains in consumer utility. Together, these mechanisms explain why total welfare falls even as cage-free production expands.

1.6.7 Comparison with Empirical Evidence

Although the calibration is intended as illustrative rather than predictive, it is useful to compare the model's implications with observed outcomes to assess whether the mechanisms it emphasizes are empirically relevant. The model was calibrated to 2018 data, before Proposition 12's production requirements took effect in January 2022. By now, several years of post-implementation data are available from USDA sources. Two caveats are important. First, outbreaks of highly pathogenic avian influenza (HPAI) in 2022 and 2024 depopulated tens of millions of laying hens nationwide (USDA APHIS), amplifying the supply disruptions beyond what the regulation alone would have caused. Second, several other states adopted cage-free mandates over this period (Massachusetts, Washington, Oregon, Colorado), so the observed national shift toward cage-free produc-

tion reflects multiple regulatory drivers, not Proposition 12 alone. The comparison below therefore focuses on directional consistency and broad magnitudes rather than precise point estimates.

California egg production. Figure 1.6 plots California’s annual egg production from 2013 to 2024 using USDA National Agricultural Statistics Service data. Production peaked at 4,295 million eggs in 2018, the year Proposition 12 was passed. After enforcement began in January 2022, output fell to 3,782 million in 2022, 3,163 million in 2023, and 2,244 million in 2024, a cumulative decline of 48% from the 2018 peak. The late-2024 HPAI outbreak hit California flocks particularly hard. The state’s layer count fell from 9.54 million in December 2023 to 4.73 million in December 2024, a decline of 50%, while the national flock declined by only 3% over the same period ([USDA NASS, 2025](#)). But production was already declining before the outbreak, dropping 17% between 2022 and 2023 when HPAI losses in California were minimal. The model implies that some in-state producers exit rather than bear the cost of converting to cage-free systems. The realized decline is far larger than the model’s predicted reduction in conventional production. This gap reflects two important differences: first, the 48% figure measures total California egg output (conventional plus cage-free) while the model’s 7.9% applies to national conventional production only; second, the model captures a single comparative-static shift and does not account for the HPAI shock or the transition dynamics of a multi-year adjustment. The direction of the effect, contraction of California’s in-state production as producers exit rather than convert, is consistent with the model’s predictions.

Cage-free flock expansion. Figure 1.7 shows the cage-free share of the U.S. laying hen flock from 2016 to 2024, based on USDA Agricultural Marketing Service Cage-Free Shell Egg Reports. The cage-free share rose from roughly 16% in 2018 to 34% in 2022 and approximately 41% by 2024. The model predicts a 16.8% increase in cage-free production at the national level, which is modest relative to the observed expansion. The discrepancy reflects the fact that multiple state-level mandates and corporate pledges drove cage-free conversion simultaneously. Nonetheless, the timing of the acceleration around 2020–2022 is consistent with producers converting capacity in anticipation of Proposition 12 enforcement, and the model correctly predicts the direction of the




Figure 1.6: California egg production, 2013–2024. Data from USDA NASS, California Poultry Reports.

shift.

Prices. The model predicts nationwide price increases of 5.7% for conventional eggs and 5.0% for cage-free eggs, with the largest effects in California. The BLS Average Price series for Grade A large eggs shows that the national average retail price rose from roughly \$1.50 per dozen in 2018 to \$2.86 in 2022 and \$3.17 in 2024. The egg component of the CPI increased 32.2% in 2022 alone, though much of that spike reflects the HPAI supply shock rather than regulation. Separating the regulatory effect from the disease effect is difficult in the aggregate price data. More informative is the California-specific premium: studies of the earlier AB 1437 regulation found that it raised California egg prices by \$0.48–\$0.80 per dozen relative to other states ([Mullally and Lusk, 2018](#)), and [Carter and Scheitrum \(2021\)](#) document significant interstate transmission of California’s regulatory costs. Under Proposition 12, [Oh and Vukina \(2022\)](#) estimate even larger price effects. The pattern of a California-specific price premium coexisting with moderate nationwide price increases is what the model predicts.




Figure 1.7: Cage-free share of U.S. laying hen flock, 2016–2024. Data from USDA AMS, Cage-Free Shell Egg Reports.

Trade patterns and import dependence. The model predicts that California imports approximately 81% of its cage-free egg consumption from other states after the regulation. California was already a net egg importer before Proposition 12, with out-of-state sources covering roughly 62% of consumption in 2018 (the model’s calibrated baseline import share of 75% is higher, reflecting the stylized two-product structure). As of January 2023, the USDA Egg Markets Overview reports a California egg self-sufficiency rate of 30%, implying that approximately 70% of the state’s eggs were sourced from other states (USDA AMS, 2023). By August 2024, in-state layer numbers had declined further to roughly 8.9 million hens, suggesting import dependence of approximately 78–79% (USDA NASS, 2024). The December 2024 flock fell to 4.73 million following the HPAI outbreak (USDA NASS, 2025), implying even greater reliance on out-of-state supply. The direction matches the model’s prediction. The realized import share is broadly consistent with the model’s long-run prediction of 81%.

Distributional effects. The distributional predictions are also broadly consistent with observed industry dynamics. [Trejo-Pech and Thompson \(2021\)](#) document capital costs of cage-free conversion ranging from \$30 to \$40 per bird for housing retrofits, with smaller producers facing disproportionately higher per-unit costs, consistent with the model’s prediction that high-cost producers exit while low-cost producers convert or benefit from higher prices. On the consumer side, [Oh and Vukina \(2022\)](#) estimate the social cost of Proposition 12 to California consumers in the hundreds of millions of dollars annually, and [Lusk \(2019\)](#) finds that roughly half of consumers are unwilling to pay the full cage-free premium voluntarily, implying that the mandate imposes net losses on a substantial share of consumers. These findings support the model’s result that forced converters and market exiters bear the largest welfare costs.

Avian influenza and regulatory vulnerability. The 2024–2025 HPAI outbreak illustrates how Proposition 12 may have amplified California’s exposure to supply shocks. While California’s flock collapsed as documented above, major producing states maintained or expanded their flocks over the same period: Iowa grew by 1%, Ohio by 9%, and Indiana by 6% ([USDA NASS, 2025](#)). California was already losing producers under the cage-free mandate before the outbreak, with layer numbers down 17% between August 2023 and August 2024 ([USDA NASS, 2024](#)). The HPAI wave then reduced the remaining flock by nearly half in four months.

Cage-free housing systems, which Proposition 12 requires, may also increase disease exposure. Approximately 56% of laying hens lost to avian influenza in 2024 were in cage-free systems, even though cage-free hens account for only about 40% of the national flock ([U.S. Department of Agriculture, Animal and Plant Health Inspection Service, 2024](#)). USDA wholesale price data show the consequences: California wholesale egg prices exceeded 800 cents per dozen in late 2024, roughly double the Midwest price over the same period ([U.S. Department of Agriculture, Agricultural Marketing Service, 2024](#)). This California–Midwest price gap has been persistently wider since Proposition 12 took effect in 2022, because the regulation eliminates the conventional-egg buffer that other states can draw on when cage-free supply contracts.

These confounding shocks make it difficult to attribute the observed price spikes and production declines to the regulation alone. The model developed in this paper isolates the regulatory channel by holding disease shocks and concurrent state mandates fixed.

In summary, the post-2022 data confirm the model’s main qualitative predictions: California egg production has declined, the national cage-free flock share has expanded, prices have risen with a California-specific premium, and California’s dependence on out-of-state supply has intensified. The realized magnitudes exceed the model’s comparative-static predictions in several dimensions, reflecting the amplifying effects of the HPAI outbreaks and the concurrent adoption of cage-free mandates by other states. But the structural directions, including producer exit in the regulating jurisdiction, the shift from conventional to cage-free production, and the distributional pattern of welfare effects, are consistent with the model’s predictions.

1.7 Robustness Analysis

The simulation results presented in Section 3.6 rely on a single set of calibrated parameters. To assess whether the qualitative findings are sensitive to the specific parameter values, this section reports a Monte Carlo robustness exercise.

1.7.1 Design

The exercise draws 10,000 parameter vectors from independent normal distributions centered at the calibrated values reported in Table 1.4. Table 1.11 reports the standard deviation for each parameter. The standard deviations are chosen to be small relative to the calibrated values, reflecting uncertainty around the point estimates while maintaining economic plausibility. The quality ratio E , which is close to unity and enters consumer welfare formulas with a large multiplier through $1/(E - 1)$, is perturbed with a particularly small standard deviation to avoid generating implausible threshold values.

For each draw, the base production cost is set to $r = 0.32 \times P_C^0$, where P_C^0 is the pre-regulation conventional price implied by the drawn parameters. The institutional parameters $\mu = 0.12$ and

Table 1.11: Monte Carlo Parameter Perturbation

Parameter	Calibrated value	Std. dev. of draw
a	5.592	0.15
b	199.769	3.0
D	15.265	0.4
E	1.150	0.002
N	0.993	0.015
ϕ	5.799	0.15
$\bar{\phi}$	17.288	0.5

$\eta = 0.03$ are held fixed. Draws that violate feasibility constraints (e.g., $a > 0$, $b > a$, $D > 1$, $E > 1$, interior threshold ordering $a < t_C^0 < t_R^0 < b$, and $\underline{\phi} < \Delta^0 < \bar{\phi}$) are discarded and redrawn. Of all attempted draws, approximately 66% satisfy the constraints.

For each valid parameter vector, the model is solved for both the pre-regulation and post-regulation equilibria, and the full set of price, quantity, trade, and welfare outcomes is computed using the formulas derived in Section 1.4.

1.7.2 Results

The Monte Carlo distributions confirm that the qualitative findings of the baseline simulation are not artifacts of the specific calibrated parameter values. Table 1.12 reports the fraction of draws in which each key outcome has the same sign as the baseline.

Table 1.12: Sign Consistency across 10,000 Monte Carlo Draws

Outcome	Fraction with baseline sign (%)
P_C increases	59.9
P_R increases	88.3
Q_C decreases	100.0
Q_R increases	100.0
CA consumer CV > 0	100.0
Producer $\Delta PS > 0$	64.8
Net welfare < 0	100.0

The central finding, that the regulation reduces total welfare, holds in every single draw.

The quantity reallocation is equally robust: conventional quantity falls and restricted quantity rises in all 10,000 draws, confirming the shift in consumption from conventional to restricted products. California consumers bear a welfare loss in every draw. Restricted prices rise in 88.3% of draws, and conventional prices rise in a majority (59.9%) of draws, though the conventional price effect is more sensitive to parameter values. Producer surplus increases in 64.8% of draws. These results indicate that while the welfare and quantity conclusions are unambiguous, the direction of price-level effects and producer surplus changes is sensitive to the specific parameter configuration.

Table 1.13 reports the distribution of the welfare outcomes, expressed as percentages of baseline market value.

Table 1.13: Welfare Outcomes across Monte Carlo Draws (% of MV_0)

	Median	5th pctile	95th pctile
CA consumer CV / MV_0	6.5	3.9	8.9
Total consumer CV / MV_0	8.4	0.5	16.0
Producer ΔPS / MV_0	1.6	-4.4	7.4
Net welfare / MV_0	-6.8	-8.6	-4.9

The median net welfare loss across draws is 6.8% of baseline market value, close to the baseline estimate of 7.7%, and the direction of the welfare effect is unambiguous: net welfare declines in every draw. The median CA consumer CV is 6.5% of MV_0 , and the 90% confidence interval lies entirely above zero. Producer surplus increases in a majority of draws (median 1.6% of MV_0), though the 90% interval spans both positive and negative values, indicating that the producer surplus effect is more parameter-sensitive than the consumer welfare result.

The range of consumer CV values reflects the sensitivity of welfare losses to model parameters, including the quality ratios D and E and the conversion cost bounds.

Figure 1.8 plots the distributions of pre-regulation equilibrium moments across the Monte Carlo draws. The baseline calibration values fall near the center of each distribution, and the key empirical targets, including the price ratio P_R/P_C and quantity ratio Q_C/Q_R , are well-centered in their respective distributions.

Figure 1.9 shows the distributions of post-regulation percentage changes in prices and quanti-




Figure 1.8: Distribution of Pre-Regulation Equilibrium Moments

ties. The conventional price change distribution has a median of +1.4%, with the baseline value of +5.7% in the upper portion of the distribution. The restricted price change has a median of +3.2%, with most draws showing a positive effect. In contrast, quantity changes are tightly distributed: the conventional quantity decline and restricted quantity increase are consistent across draws.

Figure 1.10 displays the distribution of California's restricted egg import share after the regulation. The median import share is approximately 80%, confirming that California's reliance on out-of-state cage-free supply is a robust implication of the model.

Figure 1.11 plots the welfare outcome distributions. The net welfare panel confirms the universal negative sign, while the consumer CV and producer surplus panels show that the baseline estimates are well-situated within their Monte Carlo distributions.

In sum, the Monte Carlo exercise demonstrates that the principal conclusions of the paper, namely that the regulation shifts consumption toward restricted products, generates consumer welfare losses concentrated in California, and reduces total welfare, are qualitatively robust to per-




Figure 1.9: Distribution of Post-Regulation Price and Quantity Changes




Figure 1.10: Distribution of California Import Share for Restricted Eggs




Figure 1.11: Distribution of Welfare Outcomes

turbations of the calibrated parameters. The quantity reallocation and welfare-sign results hold in every draw. Restricted prices rise in the vast majority of draws (88.3%), and conventional prices rise in a majority (59.9%). The direction of aggregate producer surplus changes is positive in most draws but more parameter-sensitive than the other results.

1.8 Conclusion

This paper proposes a partial equilibrium framework for analyzing consumption-based production standards in integrated markets. The model applies to any setting in which a jurisdiction requires that goods sold within its borders satisfy a production standard, while the jurisdiction itself participates in both the consumption and production of the regulated commodity. The framework incorporates vertically differentiated products, income-based consumer sorting, heterogeneous producer compliance costs, and explicit trade between the regulating jurisdiction and the rest of the market.

I calibrate the model to the U.S. shell egg market and use Proposition 12 as an illustrative application. In the baseline calibration, the regulation raises cage-free egg prices by 5.0% and conventional prices by 5.7% nationwide; Monte Carlo analysis confirms that the cage-free price increase is robust across parameter perturbations, while the conventional price effect is more sensitive to parameter values. Conventional egg consumption falls by 7.9%, while cage-free consumption rises by 16.8%. California becomes heavily dependent on out-of-state cage-free imports, sourcing approximately 81% of its cage-free consumption from other states. Consumer welfare losses are concentrated in the regulating jurisdiction: compensating variation for California consumers amounts to 7.8% of baseline market value, compared with 4.7% for consumers elsewhere. The producer-side effects are similarly asymmetric. California producer surplus falls by 36.8%, as conversion costs and exit among conventional producers more than offset gains to cage-free incumbents. The decline is large enough that California producers lose on net even though the regulation raises egg prices and expands cage-free demand. Outside California, producer surplus rises by 10.4%. At the national level, aggregate producer surplus increases by 8.9% in the baseline, but these gains accrue almost

entirely to out-of-state producers. Monte Carlo analysis confirms that the aggregate producer surplus gain is present in most parameter draws but is sensitive to parameter values. The analysis highlights that the effects on domestic producers in the regulating jurisdiction are a first-order welfare consideration that has received limited quantitative attention in the existing literature.

These implications are consistent with the available empirical evidence. Studies of California’s earlier housing regulations find significant California-specific price increases ([Mullally and Lusk, 2018](#); [Malone and Lusk, 2016](#); [Carter and Scheitrum, 2021](#)), and [Oh and Vukina \(2022\)](#) estimate substantial consumer costs under the stricter Proposition 12 mandate. USDA data confirm the expansion of cage-free production and the contraction of California’s in-state flock, matching the model’s predicted reallocation. The capital costs of conversion documented by [Trejo-Pech and Thompson \(2021\)](#) are consistent with the heterogeneous compliance cost structure in the model, and survey evidence on willingness to pay ([Lusk, 2019](#)) supports the finding that consumers forced to switch or exit bear the largest losses. The existing empirical work documents individual dimensions of these effects, including price changes, consumer costs, and production shifts, but to my knowledge, no prior study fully integrates the production and consumption sides within an equilibrium framework or quantifies how much producers in the regulating jurisdiction lose. This paper does both, and shows how distinguishing between switching and market exit helps clarify the distribution of welfare losses across consumers.

A Monte Carlo exercise over 10,000 perturbed parameter vectors shows that the quantity reallocation, consumer welfare losses, and net welfare decline are robust to perturbations of the calibrated parameters, holding across all draws. Restricted prices rise in 88% of draws and conventional prices rise in 60%, while the direction of aggregate producer surplus gains holds in about 65% of draws.

The framework, however, is not specific to eggs or to California. Consumption-based production standards have become increasingly common across policy domains and levels of governance. At the state level, Massachusetts’s Question 3 and similar statutes in Washington, Colorado, and other states impose animal housing requirements on products sold within their borders. At the

supranational level, the European Union’s Deforestation Regulation prohibits both the sale and the export of deforestation-linked soy, beef, palm oil, coffee, and cacao. Proposition 12 and the EUDR differ in legal structure: Proposition 12 bans both noncompliant production within California and sales of noncompliant products in the state, while the EUDR is silent on production practices but prohibits domestic sale and export alike. In both cases, however, domestic producers who cannot meet the standard lose their markets and must convert or exit. The model developed here applies to either type of regulation, because the central mechanism is the same: domestic producers face a harder constraint than external suppliers, who can redirect noncompliant output to other markets. Applying the framework to a different setting requires only the relevant market parameters.

The welfare consequences depend critically on market structure. The jurisdiction’s share of consumption and production determines whether the regulation operates mainly through trade (as with California and pork, or the EU and coffee) or also reshapes the domestic industry (as with California and eggs, or the EU and beef). Heterogeneity in compliance costs governs how many producers convert versus exit and how large the resulting efficiency losses are. The extent of vertical differentiation between compliant and non-compliant products shapes consumer substitution patterns and the magnitude of welfare losses for those forced to switch products or leave the market.

Several limitations deserve mention. The model assumes perfect competition and abstracts from market power among processors or retailers. Compliance costs are static, setting aside dynamic investment and learning that may reduce costs over time. The calibration relies on aggregate market moments rather than farm-level micro-data. The model also omits non-market benefits such as animal welfare improvements, which proponents regard as the primary justification for these regulations and which could offset some of the measured efficiency losses.

Future work could incorporate dynamic adjustment, allowing producers to invest in compliant technology over time in response to anticipated regulation. Estimating the compliance cost distribution from farm-level data would improve the precision of the welfare calculations. Applying the framework to other consumption-based standards, including the EU Deforestation Regulation, state-level pork mandates, and emerging sustainability requirements for textiles and electronics,

would test whether the mechanisms identified here generalize beyond the egg market and provide policymakers with a systematic basis for comparing the distributional consequences of this growing class of regulations.

These regulations are becoming more common. Several U.S. states have adopted animal housing mandates similar to Proposition 12, and the European Union is considering extending its animal welfare standards to imported products, which would apply the same consumption-based logic to a much larger market. The analysis in this paper shows that the effects on domestic producers in the regulating jurisdiction are a first-order welfare consideration that existing work has not quantified. Policymakers evaluating these regulations should account for what their own producers stand to lose, not only the consequences for consumer prices.

References

- Bakhtavoryan, Rafael, Vardges Hovhannisyan, Stephen Devadoss, and Jose Lopez.** 2021. “An empirical evaluation of egg demand in the United States.” *Journal of Agricultural and Applied Economics*, 53(2): 280–300.
- Beghin, John C, Miet Maertens, and Johan Swinnen.** 2015. “Nontariff measures and standards in trade and global value chains.” *Annual Review of Resource Economics*, 7: 425–450.
- Berry, Steven, James Levinsohn, and Ariel Pakes.** 1995. “Automobile prices in market equilibrium.” *Econometrica*, 63(4): 841–890.
- Böhringer, Christoph, Jared C Carbone, and Thomas F Rutherford.** 2012. “Unilateral climate policy design: Efficiency and equity implications of alternative instruments to reduce carbon leakage.” *Energy Economics*, 34: S208–S217.
- Brester, G. W., J. M. Marsh, and J. A. Atwood.** 2004. “Distributional Impacts of Country-of-Origin Labeling in the U.S. Meat Industry.” *Journal of Agricultural and Resource Economics*, 29(2): 206–227.
- Caputo, Vincenzina, Aaron J Staples, Glynn T Tonsor, and Jayson L Lusk.** 2023a. “Egg Producer Attitudes and Expectations Regarding the Transition to Cage-Free Production: A Mixed-Methods Approach.” *Poultry Science*, 102(11): 103058.

- Caputo, Vincenzina, Jayson L Lusk, Glynn T Tonsor, and Aaron J Staples.** 2023b. “The Transition to Cage-Free Eggs.” Michigan State University, Purdue University, and Kansas State University. Report prepared for the United Egg Producers, FMI, and the United Egg Association.
- Carter, Colin A, Schaefer K. Aleks, and Daniel Scheitrum.** 2021. “Piecemeal farm regulation and the US Commerce Clause.” *American Journal of Agricultural Economics*, 103(3): 1141–1163.
- Crampes, Claude, and Abraham Hollander.** 1995. “Duopoly and quality standards.” *European Economic Review*, 39(1): 71–82.
- Disdier, Anne-Célia, Lionel Fontagné, and Mondher Mimouni.** 2008. “The impact of regulations on agricultural trade: Evidence from the SPS and TBT agreements.” *American Journal of Agricultural Economics*, 90(2): 336–350.
- Fischer, Carolyn, and Alan K Fox.** 2012. “Comparing policies to combat emissions leakage: Border carbon adjustments versus rebates.” *Journal of Environmental Economics and Management*, 64(2): 199–216.
- Fowlie, Meredith, Stephen P Holland, and Erin T Mansur.** 2012. “What do emissions markets deliver and to whom? Evidence from Southern California’s NOx trading program.” *American Economic Review*, 102(2): 965–993.
- Greenstone, Michael.** 2002. “The impacts of environmental regulations on industrial activity: Evidence from the 1970 and 1977 Clean Air Act amendments and the Census of Manufactures.” *Journal of Political Economy*, 110(6): 1175–1219.
- Lee, Hanbin, Richard J Sexton, and Daniel A Sumner.** 2023. “National and subnational regulation of farm practices for consumer products sold within a jurisdiction: California’s Proposition 12.” *Agricultural Economics*, 54(6): 838–853.
- Lichtenberg, Erik, and David Zilberman.** 1988. “Efficient regulation of environmental health risks.” *Quarterly Journal of Economics*, 103(1): 167–178.

- Lusk, Jayson L.** 2019. “Consumer Preferences for Cage-Free Eggs and Impacts of Retailer Pledges.” *Agribusiness*, 35(2): 129–148.
- Malone, Trey, and Jayson L Lusk.** 2016. “Putting the Chicken Before the Egg Price: An Ex Post Analysis of California’s Battery Cage Ban.” *Journal of Agricultural and Resource Economics*, 41(3): 518–532.
- Malone, Trey, Schaefer, K. Aleks, and Jayson L Lusk.** 2021. “Unscrambling U.S. Egg Supply Chains amid COVID-19.” *Food Policy*, 101: 102046.
- Marette, Stephan, and John C Beghin.** 2010. “Are standards always protectionist?” *Review of International Economics*, 18(1): 179–192.
- Mullally, Conner, and Jayson L Lusk.** 2018. “The impact of farm animal housing restrictions on egg prices, consumer welfare, and production in California.” *American Journal of Agricultural Economics*, 100(3): 649–669.
- Mussa, Michael, and Sherwin Rosen.** 1978. “Monopoly and product quality.” *Journal of Economic theory*, 18(2): 301–317.
- Oh, Sohae Eve, and Tomislav Vukina.** 2022. “The price of cage-free eggs: Social cost of Proposition 12 in California.” *American Journal of Agricultural Economics*, 104(4): 1293–1326.
- Paul, Allen S, Jayson L Lusk, F Bailey Norwood, and Glynn T Tonsor.** 2019. “An Experiment on the Vote-Buy Gap with Application to Cage-Free Eggs.” *Journal of Behavioral and Experimental Economics*, 79: 102–109.
- Ronnen, Uri.** 1991. “Minimum quality standards, fixed costs, and competition.” *RAND Journal of Economics*, 22(4): 490–504.
- Saitone, Tina L., Richard J. Sexton, and Daniel A. Sumner.** 2015. “What Happens When Food Marketers Require Restrictive Farming Practices?” *American Journal of Agricultural Economics*, 97(4): 1021–1043.

- Shaked, Avner, and John Sutton.** 1982. “Relaxing price competition through product differentiation.” *The review of economic studies*, 49(1): 3–13.
- Small, Kenneth A., and Harvey S. Rosen.** 1981. “Applied welfare economics with discrete choice models.” *Econometrica*, 49(1): 105–130.
- Storn, Rainer, and Kenneth Price.** 1997. “Differential evolution – a simple and efficient heuristic for global optimization over continuous spaces.” *Journal of Global Optimization*, 11(4): 341–359.
- Sumner, Daniel A.** 2017. “Economics of US state and local regulation of farm practices, with emphasis on restrictions of interstate trade.” *Annual Review of Resource Economics*, 9: 13–31.
- Sumner, Daniel A., J. Thomas Rosen-Molina, William A. Matthews, Joy A. Mench, and Kurt R. Richter.** 2008. *Economic effects of proposed restrictions on egg-laying hen housing in California*. Davis, CA:University of California Agricultural Issues Center.
- Supreme Court of the United States.** 2023. “National Pork Producers Council v. Ross.” 598 U.S. 356 (2023).
- Swinnen, Johan, and Thijs Vandemoortele.** 2011. “Trade and the political economy of food standards.” *Journal of Agricultural Economics*, 62(2): 259–280.
- Trajtenberg, Manuel.** 1989. “The welfare analysis of product innovations, with an application to computed tomography scanners.” *Journal of Political Economy*, 97(2): 444–479.
- Trejo-Pech, Carlos J O, and Jada M Thompson.** 2021. “Discounted Cash Flow Valuation of Conventional and Cage-Free Production Investments.” *International Food and Agribusiness Management Review*, 24(2): 197–214.
- USDA AMS.** 2023. “Egg Markets Overview.” U.S. Department of Agriculture, Agricultural Marketing Service, Livestock, Poultry, and Grain Market News. Data derived from USDA NASS and ERS; U.S. Census Bureau.

- USDA NASS.** 2024. “Pacific Region Poultry Report.” U.S. Department of Agriculture, National Agricultural Statistics Service.
- USDA NASS.** 2025. “Pacific Region Poultry Report.” U.S. Department of Agriculture, National Agricultural Statistics Service.
- U.S. Department of Agriculture, Agricultural Marketing Service.** 2024. “Shell Egg Market Reports.” Livestock, Poultry, and Grain Market News. Weekly wholesale egg price data.
- U.S. Department of Agriculture, Animal and Plant Health Inspection Service.** 2024. “2022–2025 Confirmations of Highly Pathogenic Avian Influenza in Commercial and Backyard Flocks.” Data accessed January 2025.
- U.S. Department of Agriculture, National Agricultural Statistics Service.** 2019. “Chickens and Eggs 2018 Summary.” USDA NASS.
- Varian, Hal R.** 1992. *Microeconomic Analysis*. . 3 ed., New York:W. W. Norton & Company.
- Zhang, Wei.** 2018. “Costs of a practice-based air quality regulation: Dairy farms in the San Joaquin Valley.” *American Journal of Agricultural Economics*, 100(3): 762–785.

1.9 Appendix

1.9.1 Comparative Statics

Before the regulation, the equilibrium conventional price is

$$P_0^{C*} = \frac{b - (b - a)N}{E}. \quad (\text{A1})$$

This expression implies that the conventional price is increasing in the upper bound of income, b , increasing in the lower bound of income, a , and decreasing in total supply, N :

$$\frac{\partial P_0^{C*}}{\partial b} = \frac{1 - N}{E} > 0, \quad \frac{\partial P_0^{C*}}{\partial a} = \frac{N}{E} > 0, \quad \frac{\partial P_0^{C*}}{\partial N} = -\frac{b - a}{E} < 0.$$

The conventional price is independent of the conversion-cost parameters $\underline{\phi}$ and $\bar{\phi}$, and it is also independent of D . Economically, this is because the conventional price is pinned down by the market segment comparing conventional eggs to the outside option, so it depends on E but not on the degree of differentiation between restricted and conventional eggs.

The restricted price can be written as

$$P_0^{R*} = P_0^{C*} + \Delta_0^*, \quad (\text{A2})$$

where the equilibrium price premium is

$$\Delta_0^* = \frac{N(b - a) [\bar{\phi} + (E - 1)\underline{\phi}] + (\bar{\phi} - \underline{\phi})(E - 1)b}{E [D(\bar{\phi} - \underline{\phi}) + N(b - a)]}. \quad (\text{A3})$$

This decomposition is useful because the effects of parameters on the restricted price can be separated into their effects on the conventional price and on the price premium.

The role of total supply. An increase in N lowers the conventional price directly. Its effect on the premium is generally ambiguous:

$$\frac{\partial \Delta_0^*}{\partial N} = \frac{(b-a)(\bar{\phi} - \underline{\phi}) [D(\bar{\phi} + (E-1)\underline{\phi}) - (E-1)b]}{E [D(\bar{\phi} - \underline{\phi}) + N(b-a)]^2}. \quad (\text{A4})$$

Accordingly, the effect of N on P_0^{R*} is also ambiguous in general. The economic intuition is straightforward: more producers expand supply and push prices down, but they may also change the composition of production across conventional and restricted eggs, which affects the equilibrium premium.

The role of product differentiation. The premium is decreasing in D :

$$\frac{\partial \Delta_0^*}{\partial D} = -\frac{(\bar{\phi} - \underline{\phi}) [N(b-a)(\bar{\phi} + (E-1)\underline{\phi}) + (\bar{\phi} - \underline{\phi})(E-1)b]}{E [D(\bar{\phi} - \underline{\phi}) + N(b-a)]^2} < 0. \quad (\text{A5})$$

Since P_0^{C*} does not depend on D , it follows that

$$\frac{\partial P_0^{R*}}{\partial D} = \frac{\partial \Delta_0^*}{\partial D} < 0. \quad (\text{A6})$$

Thus, a larger D lowers the restricted price. In this model, an increase in D corresponds to a smaller utility gap between restricted and conventional eggs. As the two products become closer substitutes, consumers are less willing to pay a large premium for the restricted product, which reduces its equilibrium price.

The role of conversion costs. Higher conversion costs raise the restricted price through the premium. Intuitively, when conversion becomes more expensive, fewer producers find it profitable to supply the restricted product, shifting restricted supply inward and increasing the equilibrium price. This is the main economic mechanism, and the formal derivatives simply confirm it. In particular, the premium is increasing in both $\underline{\phi}$ and, under standard parameter restrictions, also

increasing in $\bar{\phi}$. Therefore, higher conversion costs tend to raise P_0^{R*} and reduce the equilibrium quantity of restricted eggs.

Equilibrium quantities. The equilibrium quantities are

$$Q_0^{R*} = \frac{N [(b-a)N + (E-1)b - DE\phi]}{E [D(\bar{\phi} - \phi) + N(b-a)]}, \quad (\text{A7})$$

and

$$Q_0^{C*} = \frac{N [DE\bar{\phi} + (E-1)(N(b-a) - b)]}{E [D(\bar{\phi} - \phi) + N(b-a)]}. \quad (\text{A8})$$

As total supply rises, both equilibrium quantities tend to increase. By contrast, higher conversion costs shift production away from restricted eggs and toward conventional eggs. Likewise, as D rises and the price premium shrinks, production reallocates away from restricted eggs and toward conventional eggs.

Overall, the main comparative-statics results before the regulation are intuitive. More supply lowers prices; higher conversion costs make restricted production less attractive and therefore raise the restricted premium; and a larger D compresses the premium because the two products become closer substitutes. Additional derivative-based results can be reported in an appendix.

1.9.2 Compensating Variation

This appendix derives the individual compensating variation formulas presented in Table 1.2 and provides the full case-by-case aggregate expressions for both within-jurisdiction and outside-jurisdiction consumers.

Individual CV

Recall from the model that $D = u^R/(u^R - u^C)$ and $E = u^C/(u^C - 1)$, so the quality ratios entering the CV calculations are $u^C = E/(E - 1)$, $u^R/u^C = D/(D - 1)$, and $u^R = DE/((D -$

1)($E - 1$)). For each behavioral transition, I solve for the post-regulation income $\tilde{t}$ that equates pre- and post-regulation utility.

Exiters ($C \rightarrow M$). Pre-regulation utility is $\frac{E}{E-1}(t - P_0^C)$; post-regulation utility is $\tilde{t}$. Equating:

$$\frac{E}{E-1}(t - P_0^C) = \tilde{t} \implies \tilde{t} = \frac{E}{E-1}(t - P_0^C).$$

Therefore,

$$CV^{C \rightarrow M}(t) = \tilde{t} - t = \frac{t - EP_0^C}{E-1}. \quad (\text{A9})$$

Converters ($C \rightarrow R$). Pre-regulation utility is $\frac{E}{E-1}(t - P_0^C)$; post-regulation utility is $\frac{DE}{(D-1)(E-1)}(\tilde{t} - P_1^R)$. Setting these equal and simplifying by the common factor $\frac{E}{E-1}$:

$$t - P_0^C = \frac{D}{D-1}(\tilde{t} - P_1^R) \implies \tilde{t} = \frac{D-1}{D}(t - P_0^C) + P_1^R.$$

Therefore,

$$CV^{C \rightarrow R}(t) = -\frac{1}{D}t - \frac{D-1}{D}P_0^C + P_1^R. \quad (\text{A10})$$

Restricted incumbents ($R \rightarrow R$). Pre- and post-regulation utilities share the same quality multiplier, so the quality ratio cancels:

$$t - P_0^R = \tilde{t} - P_1^R \implies \tilde{t} = t + P_1^R - P_0^R.$$

Therefore,

$$CV^{R \rightarrow R}(t) = P_1^R - P_0^R. \quad (\text{A11})$$

Conventional incumbents ($C \rightarrow C$). By analogous reasoning,

$$CV^{C \rightarrow C}(t) = P_1^C - P_0^C. \quad (\text{A12})$$

Converters ($R \rightarrow C$). Pre-regulation utility is $\frac{DE}{(D-1)(E-1)}(t - P_0^R)$; post-regulation utility is $\frac{E}{E-1}(\tilde{t} - P_1^C)$. Setting these equal and simplifying:

$$\frac{D}{D-1}(t - P_0^R) = \tilde{t} - P_1^C \implies \tilde{t} = \frac{D}{D-1}(t - P_0^R) + P_1^C.$$

Therefore,

$$CV^{R \rightarrow C}(t) = \frac{1}{D-1}t - \frac{D}{D-1}P_0^R + P_1^C. \quad (\text{A13})$$

Entrants ($M \rightarrow R$). Pre-regulation utility is t ; post-regulation utility is $\frac{DE}{(D-1)(E-1)}(\tilde{t} - P_1^R)$. Setting these equal:

$$t = \frac{DE}{(D-1)(E-1)}(\tilde{t} - P_1^R) \implies \tilde{t} = \frac{(D-1)(E-1)}{DE}t + P_1^R.$$

Therefore,

$$CV^{M \rightarrow R}(t) = -\frac{D+E-1}{DE}t + P_1^R. \quad (\text{A14})$$

Exiters ($R \rightarrow M$). Pre-regulation utility is $\frac{DE}{(D-1)(E-1)}(t - P_0^R)$; post-regulation utility is $\tilde{t}$. Setting these equal:

$$\frac{DE}{(D-1)(E-1)}(t - P_0^R) = \tilde{t} \implies \tilde{t} = \frac{DE}{(D-1)(E-1)}(t - P_0^R).$$

Therefore,

$$CV^{R \rightarrow M}(t) = \frac{D+E-1}{(D-1)(E-1)}t - \frac{DE}{(D-1)(E-1)}P_0^R. \quad (\text{A15})$$

This case arises only when the post-regulation thresholds shift sufficiently that some former restricted consumers can no longer afford any product.

Aggregate CV: Within the Jurisdiction

Aggregate CV for each behavioral group is obtained by integrating individual CV over the relevant income interval with density $f(t) = 1/(b-a)$. Since individual CV is either constant or linear in t , all integrals evaluate in closed form.

For constant CV (restricted and conventional incumbents):

$$\int_{\underline{t}}^{\bar{t}} c \cdot f(t) dt = c \cdot \frac{\bar{t} - \underline{t}}{b - a}.$$

For linear CV of the form $CV(t) = \alpha t + \beta$:

$$\int_{\underline{t}}^{\bar{t}} (\alpha t + \beta) f(t) dt = \frac{1}{b - a} \left[\frac{\alpha}{2} (\bar{t}^2 - \underline{t}^2) + \beta (\bar{t} - \underline{t}) \right].$$

The three cases for within-jurisdiction consumers are:

Case 1: $t_1^{R*} \leq t_0^{C*} < t_0^{R*}$. The post-regulation restricted threshold falls below the pre-regulation conventional threshold. Some former non-consumers now purchase restricted goods. Aggregate CV is:

$$CV_I^{(1)} = \underbrace{\int_{t_1^{R*}}^{t_0^{C*}} CV^{M \rightarrow R}(t) f(t) dt}_{\text{entrants}} + \underbrace{\int_{t_0^{C*}}^{t_0^{R*}} CV^{C \rightarrow R}(t) f(t) dt}_{\text{converters}} + \underbrace{\int_{t_0^{R*}}^b CV^{R \rightarrow R}(t) f(t) dt}_{\text{incumbents}}. \quad (\text{A16})$$

Case 2: $t_0^{C*} < t_1^{R*} < t_0^{R*}$. The empirically relevant case for Proposition 12. Some former conventional consumers exit the market, others convert to restricted goods. Aggregate CV is:

$$CV_I^{(2)} = \underbrace{\int_{t_0^{C*}}^{t_1^{R*}} CV^{C \rightarrow M}(t) f(t) dt}_{\text{exiters}} + \underbrace{\int_{t_1^{R*}}^{t_0^{R*}} CV^{C \rightarrow R}(t) f(t) dt}_{\text{converters}} + \underbrace{\int_{t_0^{R*}}^b CV^{R \rightarrow R}(t) f(t) dt}_{\text{incumbents}}. \quad (\text{A17})$$

Evaluating the integrals for Case 2 explicitly:

$$\int_{t_0^{C^*}}^{t_1^{R^*}} CV^{C \rightarrow M}(t) f(t) dt = \frac{1}{b-a} \left[\frac{1}{2(E-1)} \left((t_1^{R^*})^2 - (t_0^{C^*})^2 \right) - \frac{EP_0^C}{E-1} (t_1^{R^*} - t_0^{C^*}) \right], \quad (\text{A18})$$

$$\int_{t_1^{R^*}}^{t_0^{R^*}} CV^{C \rightarrow R}(t) f(t) dt = \frac{1}{b-a} \left[-\frac{1}{2D} \left((t_0^{R^*})^2 - (t_1^{R^*})^2 \right) + \left(P_1^R - \frac{D-1}{D} P_0^C \right) (t_0^{R^*} - t_1^{R^*}) \right], \quad (\text{A19})$$

$$\int_{t_0^{R^*}}^b CV^{R \rightarrow R} f(t) dt = (P_1^R - P_0^R) \cdot \frac{b - t_0^{R^*}}{b-a}. \quad (\text{A20})$$

Case 3: $t_0^{C^*} < t_0^{R^*} \leq t_1^{R^*}$. All former conventional consumers exit the market. Some former restricted consumers may also exit. Aggregate CV is:

$$CV_I^{(3)} = \underbrace{\int_{t_0^{C^*}}^{t_0^{R^*}} CV^{C \rightarrow M}(t) f(t) dt}_{\text{exiters from C}} + \underbrace{\int_{\max(t_0^{R^*}, t_1^{R^*})}^b CV^{R \rightarrow R} f(t) dt}_{\text{incumbents}}. \quad (\text{A21})$$

Aggregate CV: Outside the Jurisdiction

Outside the jurisdiction, both products remain available after the regulation. The post-regulation thresholds are $t_1^{C^*} = EP_1^C$ and $t_1^{R^*} = P_1^C(1-D) + P_1^R D$. If $P_1^C > P_0^C$, then $t_1^{C^*} > t_0^{C^*}$. Three cases arise depending on the relative magnitudes of the threshold shifts:

Case 1: $t_0^{C^*} < t_1^{C^*} < t_1^{R^*} < t_0^{R^*}$.

$$CV_O^{(1)} = \int_{t_0^{C^*}}^{t_1^{C^*}} CV^{C \rightarrow M} f dt + \int_{t_1^{C^*}}^{t_1^{R^*}} CV^{C \rightarrow C} f dt + \int_{t_1^{R^*}}^{t_0^{R^*}} CV^{C \rightarrow R} f dt + \int_{t_0^{R^*}}^b CV^{R \rightarrow R} f dt. \quad (\text{A22})$$

Case 2: $t_0^{C^*} < t_1^{C^*} < t_0^{R^*} < t_1^{R^*}$.

$$CV_O^{(2)} = \int_{t_0^{C^*}}^{t_1^{C^*}} CV^{C \rightarrow M} f dt + \int_{t_1^{C^*}}^{t_0^{R^*}} CV^{C \rightarrow C} f dt + \int_{t_0^{R^*}}^{t_1^{R^*}} CV^{R \rightarrow C} f dt + \int_{t_1^{R^*}}^b CV^{R \rightarrow R} f dt. \quad (\text{A23})$$

Case 3: $t_0^{C^*} < t_0^{R^*} < t_1^{C^*} < t_1^{R^*}$.

$$CV_O^{(3)} = \int_{t_0^{C^*}}^{t_0^{R^*}} CV^{C \rightarrow M} f dt + \int_{t_0^{R^*}}^{t_1^{C^*}} CV^{R \rightarrow M} f dt + \int_{t_1^{C^*}}^{t_1^{R^*}} CV^{R \rightarrow C} f dt + \int_{t_1^{R^*}}^b CV^{R \rightarrow R} f dt. \quad (\text{A24})$$

All integrals evaluate in closed form using the linear integration formula above.

1.9.3 Producer Surplus

This appendix provides the closed-form expressions for aggregate producer surplus changes.

Within the Jurisdiction

Before the regulation, producers choose restricted production if $\phi \leq \phi_0^* = P_0^R - P_0^C$ and conventional production otherwise. After the regulation, only restricted production is permitted, and a producer with cost ϕ remains active only if $\phi \leq \phi_{I_1}^* = P_1^R - r$. Individual surplus changes are:

$$\begin{aligned}\Delta PS^{R \rightarrow R}(\phi) &= (P_1^R - r - \phi) - (P_0^R - r - \phi) = P_1^R - P_0^R, \\ \Delta PS^{C \rightarrow R}(\phi) &= (P_1^R - r - \phi) - (P_0^C - r) = P_1^R - P_0^C - \phi, \\ \Delta PS^{exit}(\phi) &= -(P_0^C - r).\end{aligned}$$

Integrating over the uniform distribution $g(\phi) = 1/(\bar{\phi} - \underline{\phi})$:

$$\Delta PS_I^{R \rightarrow R} = \frac{\phi_0^* - \underline{\phi}}{\bar{\phi} - \underline{\phi}} (P_1^R - P_0^R), \quad (\text{A25})$$

$$\Delta PS_I^{C \rightarrow R} = \frac{1}{\bar{\phi} - \underline{\phi}} \left[(P_1^R - P_0^C)(\phi_{I_1}^* - \phi_0^*) - \frac{1}{2} ((\phi_{I_1}^*)^2 - (\phi_0^*)^2) \right], \quad (\text{A26})$$

$$\Delta PS_I^{exit} = -\frac{\bar{\phi} - \phi_{I_1}^*}{\bar{\phi} - \underline{\phi}} (P_0^C - r). \quad (\text{A27})$$

Total within-jurisdiction producer surplus change is $\Delta PS_I = \Delta PS_I^{R \rightarrow R} + \Delta PS_I^{C \rightarrow R} + \Delta PS_I^{exit}$.

Outside the Jurisdiction

Outside the jurisdiction, both technologies remain available. The post-regulation cutoff is $\phi_{O_1}^* = P_1^R - P_1^C$. Individual surplus changes are:

$$\Delta PS^{R \rightarrow R}(\phi) = (P_1^R - r - \phi) - (P_0^R - r - \phi) = P_1^R - P_0^R,$$

$$\Delta PS^{C \rightarrow R}(\phi) = (P_1^R - r - \phi) - (P_0^C - r) = P_1^R - P_0^C - \phi,$$

$$\Delta PS^{C \rightarrow C}(\phi) = (P_1^C - r) - (P_0^C - r) = P_1^C - P_0^C.$$

Integrating:

$$\Delta PS_O^{R \rightarrow R} = \frac{\phi_0^* - \phi}{\bar{\phi} - \underline{\phi}} (P_1^R - P_0^R), \quad (\text{A28})$$

$$\Delta PS_O^{C \rightarrow R} = \frac{1}{\bar{\phi} - \underline{\phi}} \left[(P_1^R - P_0^C)(\phi_{O_1}^* - \phi_0^*) - \frac{1}{2} ((\phi_{O_1}^*)^2 - (\phi_0^*)^2) \right], \quad (\text{A29})$$

$$\Delta PS_O^{C \rightarrow C} = \frac{\bar{\phi} - \phi_{O_1}^*}{\bar{\phi} - \underline{\phi}} (P_1^C - P_0^C). \quad (\text{A30})$$

Total outside-jurisdiction producer surplus change is $\Delta PS_O = \Delta PS_O^{R \rightarrow R} + \Delta PS_O^{C \rightarrow R} + \Delta PS_O^{C \rightarrow C}$.

Total economy-wide producer surplus change is $\Delta PS = \Delta PS_I + \Delta PS_O$.

Chapter 2

Beyond Access: How Does Electricity Quality Affect the Adoption of Energy-Consuming Appliances?

Abstract

Most research on electrification in developing countries focuses on whether households have access to the grid. This paper asks a different question: conditional on being connected, does the quality of electricity service affect which appliances households adopt? Using household-level data from the World Bank Multi-Tier Framework survey in Zambia, I estimate the effects of electricity quality, defined as achieving the highest reliability and availability tiers, on both aggregate appliance capacity and ownership of 30 individual appliances spanning 3 to 2,000 watts. For urban households, I estimate OLS and Probit models controlling for income and demographics. For rural households, I instrument for electricity quality using the distance to the nearest distribution substation, exploiting the fact that voltage declines with distance along distribution lines. Improved quality increases aggregate appliance capacity by 308 watts in urban areas (OLS) and 798 to 1,329 watts in rural areas (OLS and IV). Appliance-level estimates reveal that these effects concentrate on appliances whose value depends on continuous power: refrigerator ownership increases by 16.7 percentage points in urban areas and 22.6 percentage points in rural areas, and electric iron ownership increases by 20.3 and 30.9 percentage points, respectively. Entertainment

appliances used for sustained daily activities (fans, color televisions, DVD players) also respond significantly. By contrast, appliances that can be used in brief flexible episodes (kettles, rice cookers) and very high-wattage appliances constrained by purchase cost (air conditioners, space heaters) show little response. The quality effect does not vary significantly by household income or gender of the household head, consistent with reliability operating as a supply-side constraint that affects all connected households. Comparing quality to access effects, grid connection increases appliance capacity by roughly 1,500 to 1,800 watts, approximately six times the quality effect in urban areas. For the approximately 1.2 million Zambian households already on the grid, investments in distribution infrastructure that improve service reliability can yield meaningful welfare gains by enabling adoption of appliances that households currently forgo due to poor power quality.

2.1 Introduction

A large literature in development economics studies the effects of expanding electricity access on household welfare. Electrification programs in South Africa and Brazil have been shown to increase employment, raise incomes, and improve educational outcomes (Dinkelman, 2011; Lipscomb, Mobarak and Barham, 2013). At the same time, experimental and quasi-experimental evidence from Kenya and India suggests that grid connections alone may not generate large short-run welfare gains if households cannot afford appliances or if the supply is unreliable (Lee, Miguel and Wolfram, 2020; Burlig and Preonas, 2024). As countries in Sub-Saharan Africa continue to expand grid coverage, off-grid and mini-grid solutions have emerged as alternatives for remote areas, but the majority of newly connected households still rely on centralized grid infrastructure (Blimpo and Cosgrove-Davies, 2019). For these households, a growing body of evidence indicates that the quality of service, not merely access itself, is the binding constraint on their ability to benefit from electrification (Chakravorty, Pelli and Marchand, 2014; Allcott, Collard-Wexler and O’Connell, 2016; Meeks et al., 2023).

This paper asks whether the quality of electricity service affects which appliances households

adopt. I use household-level data from the World Bank Multi-Tier Framework (MTF) survey in Zambia ([Bhatia and Angelou, 2015](#)), which measures electricity access along multiple dimensions, including supply reliability, availability, and voltage stability, rather than relying on the traditional binary metric of grid connection. I estimate the effects of electricity quality on both aggregate appliance capacity and ownership of 30 individual appliances spanning power loads from 3 watts (LED lights) to 2,000 watts (air conditioners). For urban households, I estimate OLS and Probit models controlling for income and demographics. For rural households, where electricity quality may be endogenous to unobserved household characteristics, I instrument for quality using the distance from each household to the nearest distribution substation.

I find that improved electricity quality significantly increases household appliance capacity in both urban and rural areas. The effects are not uniform across appliance types: ownership increases concentrate among appliances that require continuous power to deliver value, such as refrigerators and electric irons, while appliances that can be used in brief episodes or that are constrained by purchase cost show little response. These patterns suggest that the relevant distinction for quality-driven adoption is not wattage per se, but whether an appliance’s functionality depends on uninterrupted supply.

This paper contributes to the literature on electricity and economic development in three ways. First, most research on electrification in developing countries examines the extensive margin of grid access ([Dinkelman, 2011](#); [Lipscomb, Mobarak and Barham, 2013](#); [Lee, Miguel and Wolfram, 2020](#); [Burlig and Preonas, 2024](#)). A smaller but growing literature studies the intensive margin by examining the role of supply reliability for firms ([Allcott, Collard-Wexler and O’Connell, 2016](#)), agricultural wages ([Nag, 2024](#)), and household willingness to pay ([Deutschmann, Postepska and Sarr, 2021](#)). I contribute to this latter strand by estimating how electricity quality affects household appliance portfolios, a channel through which reliability shapes welfare but which has received limited empirical attention. Second, I provide systematic evidence on the heterogeneous effects of electricity quality across the full spectrum of household appliances, showing that the pattern of adoption responses is driven by appliance use characteristics rather than by wattage alone. Third,

by comparing quality effects to access effects within the same country and dataset, I quantify the welfare gains available to the approximately 1.2 million Zambian households already on the grid (Luzi et al., 2019) from reliability improvements relative to the gains from initial electrification.

The remainder of this paper is organized as follows. Section 2 describes the background on electricity supply and energy use in Zambia. Section 3 reviews the related literature. Section 4 presents the data. Section 5 introduces the empirical strategy. Section 6 reports results on electricity quality effects on appliance capacity, appliance-level ownership, heterogeneity by income and gender, and a comparison to electricity access. Section 7 concludes.

2.2 Background

2.2.1 Electricity Quality and Appliance Usage

The quality of an electricity supply is defined by its voltage stability, frequency stability, and continuity of service (Bollen, 2000). In Zambia, the nominal standard is 230 volts at 50 hertz. When voltage drops below acceptable levels (brownouts) or spikes above them (surges), appliances may shut down, operate inefficiently, or sustain permanent damage to electrical components. Frequency deviations can cause motors in refrigerators, fans, and air conditioners to run at incorrect speeds, accelerating wear and reducing efficiency.

These quality dimensions are particularly relevant for appliance adoption decisions. Households that experience frequent outages or voltage fluctuations face a higher effective cost of owning sensitive appliances: food spoils in unreliable refrigerators, electronic components degrade faster, and the expected service life of appliances decreases. Rational households will therefore be less willing to invest in appliances whose value depends on stable, continuous power when they anticipate poor supply quality.

2.2.2 Energy Use in Zambia

Zambia’s electricity sector is dominated by hydropower, which accounts for about 83% of total installed capacity of 3,356.6 MW as of 2022.¹ Three major stations, Kariba Dam, Kafue Gorge, and Itzhi-Tezhi, have historically generated enough power to meet domestic demand and export the surplus: peak national demand is approximately 2,300 MW against installed capacity that exceeds it by roughly 1,100 MW. The mining sector consumes 51% of generated electricity, with the domestic sector accounting for 33%.

ZESCO, the state-owned vertically integrated utility, owns and operates most of the generation, transmission, and distribution infrastructure. The transmission network carries high-voltage power from remote hydropower stations to load centers, where substations step the voltage down for distribution to end users. Because the major generating stations are concentrated along the Zambezi and Kafue rivers in the south, households in the northern and eastern provinces are served by long transmission and distribution lines. The length of these lines matters for service quality: electrical resistance causes voltage to drop over distance, so households far from substations tend to receive lower and less stable voltage. This geographic structure generates the spatial variation in electricity quality that I exploit in my empirical strategy.

Despite aggregate generating capacity that exceeds peak demand, 66% of Zambians lack electricity access, a figure that rises to 96% in rural areas.² The low access rate reflects the cost of extending distribution lines to dispersed rural settlements rather than a shortage of generation. Among households that are connected, service quality varies considerably. In my survey data, approximately 60% of urban electrified households achieve the highest reliability and availability tiers, while the corresponding share among rural electrified households is much lower (Section 5.1).

Dependence on rainfall also makes supply vulnerable to drought. In 2024, an El Niño-driven drought produced a generation deficit of 700 MW, forcing ZESCO to import electricity to sustain

¹https://www.moe.gov.zm/?page_id=2198

²<https://unfccc.int/climate-action/momentum-for-change/financing-for-climate-friendly-investment/beyond-the-grid-fund-for-Zambia>

copper mining operations.³ Such generation shortfalls are transmitted to connected households through load shedding and increased outage frequency, compounding the voltage quality problems created by long distribution lines.

Zambia has begun diversifying its generation mix, though progress remains limited. With approximately 2,000 to 3,000 hours of sunshine per year,⁴ the country has substantial solar potential, and the 54 MW Bangweulu Solar Power Plant is among the largest on the continent. Solar generation complements hydropower because it peaks during dry seasons when water levels are low. However, solar still accounts for only about 3% of installed capacity, with coal and heavy fuel oil contributing another 11%.

2.3 Literature Review

A large body of research examines the economic effects of electricity access in developing countries. [Dinkelman \(2011\)](#) exploits the gradual rollout of electrification in rural South Africa and finds that access increased female employment by shifting time from home production to market work. [Lipscomb, Mobarak and Barham \(2013\)](#) use the topographic suitability for hydropower placement as an instrument in Brazil and document positive effects on housing values, urbanization, and labor productivity. However, recent experimental evidence has tempered some of these findings. [Lee, Miguel and Wolfram \(2020\)](#) conduct a randomized experiment in rural Kenya and find that new grid connections had limited effects on income, employment, and educational outcomes in the short run, suggesting that access alone is not sufficient to generate large welfare gains. [Burlig and Preonas \(2024\)](#) reach a similar conclusion using a regression discontinuity design around India's national rural electrification program.

The effect of electricity service quality on users in developing countries has received increasing attention. At the firm and aggregate level, [Chakravorty, Pelli and Marchand \(2014\)](#) find that grid access combined with good electricity quality increases non-agricultural incomes in rural India by

³<https://www.reuters.com/business/energy/zambia-seeks-power-imports-key-mining-sector-2024-04-22/>

⁴<http://www.smartsolar-zambia.com/solar-sector-information/solar-power-potential-zambia/>

28.6 percent between 1994 and 2005, compared with roughly 9 percent for grid access alone, indicating that quality is a quantitatively important margin. [Allcott, Collard-Wexler and O’Connell \(2016\)](#) document how electricity shortages reduce manufacturing output in India, with firms partially adapting through costly self-generation; average reported shortages reduce plant revenues and producer surplus by 5 to 10 percent. [Nag \(2024\)](#) show that reliable electricity supply reduces agricultural labor wages in Indian villages by enabling mechanization and shifting labor demand. [Burgess et al. \(2020\)](#) examine the political economy of electricity provision in India and argue that treating electricity as a right leads to underpricing, fiscal losses for utilities, and poor service quality, illustrating how the institutional structure of supply can undermine the welfare gains from access.

At the household level, [McRae \(2010\)](#) develops a model of household electricity demand that incorporates both short-run consumption decisions and long-run appliance investment, and uses data from Colombia to show how improvements in supply reliability shift the portfolio of household appliances. [Meeks et al. \(2023\)](#) provide experimental evidence from the Kyrgyz Republic, where randomized improvements in grid infrastructure reduced voltage fluctuations and increased household electricity consumption, particularly during peak months. [Deutschmann, Postepska and Sarr \(2021\)](#) estimate household willingness to pay for improved reliability in Senegal and find that consumers value fewer outages even when current tariffs are already subsidized. [Jimenez Mori \(2021\)](#) use MTF survey data from multiple countries and find that household satisfaction with electricity service is driven primarily by supply quality rather than price, suggesting that reliability improvements may matter more for welfare than tariff reductions. Supply reliability clearly matters for economic outcomes, but evidence on how quality affects household appliance portfolios in Sub-Saharan Africa remains limited.

A related strand of literature examines the determinants of appliance adoption and household energy demand in developing countries. [Wolfram, Shelef and Gertler \(2012\)](#) argue that rising incomes in developing countries will dramatically increase energy consumption, and that understanding appliance adoption patterns is central to forecasting future demand. [Gertler et al. \(2016\)](#) show that appliance ownership follows a nonlinear S-curve with income in Mexico, with rapid adop-

tion of energy-intensive durables such as air conditioners and refrigerators once income thresholds are reached. [Lee, Miguel and Wolfram \(2016\)](#) find that even after gaining electricity access in rural Kenya, households adopt very few appliances beyond basic lighting, with income constraints serving as the primary barrier to productive use. [Li et al. \(2019\)](#) develop projections of appliance adoption under different population, urbanization, and income scenarios.

My paper bridges the literatures on electricity quality and appliance adoption. While the existing quality literature focuses primarily on aggregate outcomes such as income, employment, or total electricity consumption, I provide disaggregated evidence on how quality affects the ownership of individual appliances across the full range of household power loads. While the appliance adoption literature emphasizes income as the key determinant, I show that supply reliability is an additional and distinct constraint, particularly for appliances whose value depends on continuous power. My identification strategy exploits geographic variation in the distance to distribution substations in Zambia, and my use of the MTF survey data allows me to measure electricity quality along multiple dimensions rather than relying on a binary access indicator. To my knowledge, this is the first paper to estimate appliance-level marginal effects of electricity quality across an entire Sub-Saharan African country.

2.4 Data

2.4.1 Household Survey Data

My household-level data come from the Zambia Multi-Tier Framework (MTF) Survey, conducted by the World Bank Group in collaboration with the Energy Sector Management Assistance Program (ESMAP). Rather than classifying households as simply connected or unconnected, the MTF methodology measures electricity access along a continuum of tiers that capture capacity, availability, reliability, quality, and affordability.

The survey employed a two-stage stratified sampling design, drawing on the 2010 Census of Population and Housing at the enumeration area (EA) level. Data were collected between September

Table 2.1: Distribution of sampled households

Province	Urban			Rural			Nationwide		
	Grid	Non	Total	Grid	Non	Total	Grid	Non	Total
Central	69	42	111	18	192	210	87	234	321
Copperbelt	361	165	526	15	81	96	376	246	622
Eastern	37	32	69	7	301	308	44	333	377
Luapula	24	42	66	18	190	208	42	232	274
Lusaka	415	249	664	40	58	98	455	307	762
Muchinga	23	17	40	6	131	137	29	148	177
North Western	15	39	54	0	100	100	15	139	154
Northern	26	38	64	14	172	186	40	210	250
Southern	85	55	140	30	206	236	115	261	376
Western	21	21	42	0	182	182	21	203	224
Total	1,076	700	1,776	148	1,613	1,761	1,224	2,313	3,537

2017 and March 2018. A total of 3,537 households (1,776 urban and 1,761 rural) were interviewed across 260 EAs in 849 villages spanning all 10 Zambian provinces, with equal allocation between urban and rural strata. In each EA, 14 households were sampled. Table 2.1 presents the distribution of grid-connected and non-grid households. In urban areas, approximately 60% of sampled households are connected to the grid, compared with fewer than 10% in rural areas.

Table 2.2 describes the key variables used in the analysis, including the weekly frequency and duration of unexpected disruptions, daily hours of electricity availability, and appliance ownership. The MTF methodology aggregates these raw measures into reliability and availability tiers on a scale of 1 to 5. Table 2.3 reports sample-weighted summary statistics for the 1,224 grid-connected households. Disruption frequency, duration, and availability are specific to grid service and are therefore reported only for this subsample. Among these households, the average weekly frequency

of unexpected disruptions is 1.06 and the average weekly duration is 1.68 hours, corresponding to a mean reliability tier of 4.7. Average daily electricity availability is 21 hours, corresponding to a mean availability tier of 4.4.

The survey records household ownership of 30 appliances with power consumption ranging from a few watts (LED lights, mobile phone chargers) to several thousand watts (space heaters, air conditioners). I classify these appliances into five load categories based on their rated wattage (Table 2.4) and construct a measure of household aggregated appliance capacity by summing the average wattage of each owned appliance. The mean aggregated capacity is 545.4 watts nationally, with a pronounced urban-rural gap between urban and rural areas.

2.4.2 Grid Data

I obtain electricity transmission and distribution grid data from the World Bank Group. The dataset covers planned and existing grid lines across Africa, tagged by country, operational status, and voltage level (from sub-transmission to 700 kV extra-high voltage). I restrict the sample to existing grid lines in Zambia as of 2018. I merge the grid line data with the World Bank’s point-level database of electrical substations and transformers in Zambia.

Figure 2.1 maps the locations of surveyed households alongside major power plants, substations, and high-voltage grid lines. Households without grid access tend to be located far from grid infrastructure, while those reporting good electricity quality are generally near substations. Households with poor quality, however, are found at varying distances from infrastructure, suggesting that proximity alone does not guarantee reliable service. On average, surveyed households are 11.7 kilometers from the nearest substation and 20.9 kilometers from the nearest high-voltage grid line.

2.4.3 Climate Data

Zambia, located in Central and Southern Africa, experiences a predominantly tropical climate influenced by its elevation. According to the Köppen climate classification, most of the country falls under the humid subtropical or tropical wet-and-dry categories, with small areas in the southwest

Table 2.2: Data Description for Selected Variables

Variables	Descriptions
<i>Raw Variable</i>	
dis_frq	The weekly frequency of unexpected electricity disruptions
dis_dur (hours)	The weekly duration of unexpected electricity disruptions
dur_day (hours)	The hours of electricity are available each day and night
Reliability_tier	Tier 2: >14 disruptions per week Tier 3: ≤ 14 disruptions per week or ≤ 3 disruptions per week with total duration of >2 hours Tier 4: >3 to 14 disruptions per week or ≤ 3 disruptions per week with >2 hours of outage Tier 5: ≤ 3 disruptions per week with total duration of <2 hours
Availability_tier	Tier 0: Less than 4 hours Tier 2: At least 4 hours Tier 3: At least 8 hours Tier 4: At least 16 hours Tier 5: At least 23 hours
appl_use_i	Appliances ownership of different load levels
ZambiaMinDist2Sub (meters)	The Distance from household to its nearest substation
ZambiaMinDist2Grid (meters)	The shortest distance from household to high voltage transmission and distribution grid line
<i>Outcome Variables</i>	
appl_capacity (W)	Household Aggregated Appliance Capacity
sh_energy	Share of household monthly expenditure on energy
<i>Control Variables</i>	
Avg_consumption (local)	Household monthly average consumption
hhsiz	Household size
gender	Household Head Gender

Source: Multi-Tier Framework for Energy Access (MTF), The World Bank Group

Table 2.3: Descriptive statistics

Variable	N	Mean	Std. Dev	Min	Max
dis_frq	1,224	1.06	2.41	0	30
dis_dur (hours)	1,224	1.68	5.16	0	61
dur_day (hours)	1,224	21.05	5.97	2	24
Reliability_tier	1,224	4.74	0.44	2	5
Availability_tier	1,224	4.43	2.05	0	5
appl_capacity (W)	3,342	545.38	1157.42	0	1,2759
ZambiaMinDist2Sub (meters)	3,364	11,674	17,546	35	135,816
ZambiaMinDist2Grid (meters)	3,364	20,929	38,930	1.71	190,546
Avg_consumption (local)	3,537	2,068	3,357	72	75,913
sh_energy	3,537	0.088	0.08	0	0.89

Table 2.4: Appliances by load level

Load level	Indicative electric appliances
Very low load (3–25 W)	Torch/flashlight/lantern, regular mobile phone charger, LED light bulb, incandescent light bulb, smartphone (internet phone) charger
Low load (75–150 W)	Compact fluorescent light (CFL) bulb, radio/CD players/sound system, fluorescent tube, television (b&w), electric sewing machine, fan, rice cooker, VCD/DVD, flat color TV, computer, regular color TV
Medium load (200–750 W)	Freezer, refrigerator, electric food processor/blender, hair dryer, washing machine, electric water pump
High load (1,000–1,500 W)	Electric iron, electric hot water pot/kettle, microwave oven, electric water heater, indoor air cooler
Very high load (2,000 W or more)	Air conditioner, solar based water heater, space heater

Note: Wattages reflect typical appliance ratings used in our analysis.




Figure 2.1: Zambia: Electricity Access and Quality Distribution

characterized by a semi-arid steppe climate (Beck et al., 2018).


Figure 2.2 illustrates four distinct climate zones in Zambia. Based on the Beck et al. (2018) Köppen-Geiger classification at 1-km resolution, 19.40% of Zambia’s land area falls within the Tropical Savanna (Aw) climate zone, characterized by pronounced wet and dry seasons. 2.62% is classified as Arid Steppe (BSh), a semi-arid region with a dry climate. 70.91% of the country is situated in the humid subtropical (Cwa) climate zone, experiencing dry winters and hot, humid summers. The remaining 7.07% falls within the subtropical highland (Cwb) climate zone, which features a temperate oceanic climate with dry winters and warm summers.

I classify each surveyed household into its corresponding climate zone using its GPS coordinates and the Beck et al. (2018) raster at 1-km resolution. Panel (b) of Figure 2.2 maps the resulting classification. Among my surveyed households, 10.72% are located in Aw, 1.64% in BSh, 82.57% in Cwa, and 5.06% in Cwb. The sample distribution broadly mirrors the national land-area distribution, with all four climate zones represented, confirming that my data provide geographic coverage across Zambia’s diverse climatic regions. The moderate overrepresentation of the Cwa zone in the sample reflects the concentration of Zambia’s population in the central and southern provinces, where the majority of urban centers and grid infrastructure are located.

2.5 Estimation Strategy


2.5.1 Model

I restrict the sample to all electrified households, including both grid-connected and off-grid users. Off-grid households, who lack grid service, are classified as having poor quality ($Q_i = 0$) by construction. This is why the regression samples (1,066 urban, 272 rural) are larger than the grid-connected counts in Table 2.1. Among grid-connected households, a household is classified as having good quality ($Q_i = 1$) if it achieves the highest tier (tier 5) on both the MTF availability and reliability dimensions. This requires fewer than 3 disruptions per week with total duration under 2 hours and at least 23 hours of daily electricity availability. All other electrified households are



(a) Köppen Climate Classification

Surveyed Households by Köppen Climate Zone in Zambia



(b) Surveyed Households by Climate Zone

Figure 2.2: Zambia: Köppen Climate Classification and Sample Distribution

Source: Beck et al. (2018) Köppen-Geiger climate classification.

classified as having poor quality ($Q_i = 0$).

My baseline specification regresses outcome Y_i for household i on the quality indicator Q_i and a vector of household characteristics X_i :

$$Y_i = \beta Q_i + \Theta X_i + \epsilon_i \quad (2.1)$$

where:

$$Q_i = \begin{cases} 0 & \text{if Availability Tier} \neq 5 \text{ or Reliability Tier} \neq 5 \\ 1 & \text{if Availability Tier} = 5 \text{ and Reliability Tier} = 5 \end{cases} \quad (2.2)$$

The coefficient β captures the effect of good electricity quality on outcome Y_i . The vector X_i includes household size, gender of household head, number of rooms, dwelling type, and province fixed effects. Standard errors are clustered at the enumeration area level. I estimate the model separately for urban and rural households.

Table 2.5 compares the characteristics of households with good electricity quality ($Q = 1$) and poor quality ($Q = 0$). In urban areas, the two groups are well balanced: there are no statistically significant differences in age, gender of household head, monthly expenditure, or appliance capacity. Only household size differs (4.4 vs. 4.1, $p < 0.05$). This balance supports the credibility of the urban OLS estimates. In rural areas, several imbalances emerge. Poor-quality households are larger, older-headed, more likely to be male-headed, and located 12.7 km farther from the nearest substation ($p < 0.01$), consistent with the first-stage relationship exploited in the IV strategy. Good-quality households have higher appliance capacity (3,120 vs. 2,167 watts, $p < 0.01$) and higher monthly expenditure (2,507 vs. 1,934 kwacha), though the expenditure difference is not statistically significant. These rural imbalances reinforce the importance of the IV approach for identification.

Table 2.5: Balance Table: Good vs. Poor Electricity Quality

	Urban			Rural		
	Poor	Good	Diff	Poor	Good	Diff
Household Size	4.4 (2.4)	4.1 (2.5)	−0.4**	5.4 (2.9)	3.6 (2.7)	−1.9***
Age of Head	42.4 (12.7)	42.5 (13.7)	0.1	44.5 (15.4)	38.8 (12.7)	−5.7***
Female Head (%)	25.2 (43.5)	25.4 (43.6)	0.2	16.3 (37.0)	27.7 (45.0)	11.4**
Monthly Expenditure	4,189 (5,755)	3,867 (4,492)	−322.2	1,934 (3,809)	2,507 (2,143)	573.3
Appliance Capacity (W)	3,935 (2,301)	4,139 (1,945)	203.7	2,167 (2,145)	3,120 (2,040)	952.7***
Dist. to Substation (km)	2.2 (2.6)	2.0 (2.5)	−0.2	19.7 (16.0)	7.0 (7.5)	−12.7***
Observations	401	665	1,066	178	94	272

Notes: Sample is electrified households. Poor and Good refer to electricity quality.

Means with standard deviations in parentheses. Female Head reported in percent.

Stars on differences from Welch t-tests. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

2.5.2 Urban

For urban households, the primary endogeneity concern is that higher-income households may simultaneously enjoy better electricity quality and own more appliances. Urban households generally have similar physical access to retail markets, so geographic sorting into areas with different appliance availability is less of a concern. I address the income channel by controlling for household monthly expenditure as a proxy for income:

$$Y_i = \beta Q_i + \Theta X_i + \phi_i + \epsilon_i \quad (2.3)$$

where ϕ_i denotes the household income proxy. I present results across three specifications with progressively richer controls: (1) province fixed effects only, (2) adding demographic and income controls, and (3) adding climate zone indicators.

2.5.3 Rural

Rural households pose additional identification challenges. Unlike in urban areas, where access to appliance markets is relatively uniform, rural households are geographically dispersed and vary considerably in their proximity to retail centers. A household located near a city may find it easier to purchase appliances regardless of electricity quality. I therefore add the distance from each household to the nearest of 20 major Zambian cities (selected by population and GDP) as a spatial control, denoted λ_i .

Even after controlling for income and distance to the nearest city, electricity quality may remain endogenous if unobserved household characteristics are correlated with both quality and appliance ownership. I address this concern using an instrumental variables (IV) strategy. My instrument is the Euclidean distance from each household to the nearest distribution substation, denoted D_i . The physical basis for this instrument is that electrical resistance in distribution lines causes voltage to decline with distance, so that households farther from substations receive lower-quality power ([Bollen, 2000](#)).

For the instrument to be valid, two conditions must hold. First, distance must predict electricity quality (relevance): $\text{Cov}(D_i, Q_i) \neq 0$. Figure 2.3 provides visual evidence for this condition, showing that the probability of good electricity quality declines with distance to the nearest substation. I formally test relevance using the first-stage F-statistic in Section 6. Second, distance must affect appliance ownership only through its effect on quality (exclusion): $\text{Cov}(D_i, \epsilon_i) = 0$. The inclusion of distance to the nearest city as a separate control helps to absorb any direct effect of remoteness on appliance access, strengthening the case that substation distance operates through the quality channel.




Figure 2.3: Distribution of Distance to the Nearest Substation, Electricity Access and Quality

The second-stage equation for rural households is:

$$Y_i = \beta Q_i + \Theta X_i + \phi_i + \lambda_i + \epsilon_i \quad (2.4)$$

where Q_i is instrumented by D_i in the first stage.

2.5.4 Appliance-Level Models

To examine heterogeneous effects across appliance types, I estimate Probit models for each of the 30 appliances in the survey. Let A_{ij} indicate whether household i owns appliance j :

$$P(A_{ij} = 1 \mid Q_i, X_i) = \Phi(\beta_j Q_i + \Theta_j X_i) \quad (2.5)$$

where $\Phi(\cdot)$ is the standard normal CDF. I report average marginal effects (AME) of the quality indicator, computed as the sample mean of the individual marginal effects $\phi(\beta_j Q_i + \Theta_j X_i) \cdot \beta_j$. The AME gives the average percentage-point change in the probability of owning appliance j when electricity quality shifts from poor to good, holding other covariates at their observed values. Estimating separate models for each appliance allows the effect of quality to vary freely across the appliance portfolio, rather than imposing a common coefficient.

2.6 Empirical Results

2.6.1 Impacts of Electricity Quality

I begin by estimating the effect of electricity quality on household aggregated appliance capacity, measured in watts. For each subsample I present three specifications with progressively richer controls, all including province fixed effects and clustering standard errors at the enumeration area level.

Urban Households. Table 2.6 reports OLS estimates for urban electrified households. The point estimate is 278.69 watts without controls, 287.13 watts with demographic and income controls, and 308.04 watts with the full set of controls including climate zone indicators. The coefficient is statistically significant at the 5% level in the preferred specification. The stability of the estimate across specifications indicates that the quality-capacity relationship is not driven by observable differences between households with good and poor electricity quality.

Table 2.6: Urban: Impacts of Electricity Quality

	Aggregated Appliance Capacity (W)		
	(1)	(2)	(3)
Electricity Quality	278.69 * (143.57)	287.13 ** (143.01)	308.04 ** (143.61)
Observations	1,066	1,066	1,066
R^2	0.058	0.332	0.332
Clusters	277	277	277
Demographic & Income	No	Yes	Yes
Climate Controls	No	No	Yes

Notes: This table presents OLS estimates for urban electrified households.

Robust standard errors clustered by EAs in parentheses.

*** Significant at the 1 percent level.

** Significant at the 5 percent level.

* Significant at the 10 percent level.

Rural Households. Table 2.7 reports estimates for rural electrified households. The OLS results in Panel A show a coefficient of 854.85 watts without controls, declining to 764.73 watts with demographic and income controls and adjusting to 797.78 watts with the full set of controls including climate zone and distance-to-city indicators. All three specifications are significant at the 5% level or better. The larger magnitude relative to urban areas is consistent with rural households facing more binding supply constraints: when quality improves, the marginal value of adopting additional appliances is higher for households that previously made limited use of their electricity connection.

As discussed in Section 5.3, electricity quality may be endogenous in rural areas if households that sort closer to grid infrastructure differ systematically in their demand for appliances. Panel B reports IV estimates that instrument for quality using distance to the nearest distribution substation. Table 2.8 presents the first-stage results. Distance to the nearest substation is a strong predictor of quality: each additional 10 kilometers reduces the probability of good quality by 10 to 12 percentage points across specifications. Distance to the nearest high-voltage grid line, by contrast, is not statistically significant, consistent with the fact that high-voltage power must be stepped down through local substations before reaching households. I therefore use only substation distance as my instrument.

The IV estimates are larger than the OLS estimates, ranging from 2,110.35 watts without controls to 1,329.28 watts with the full set of controls. The upward shift from OLS to IV could reflect attenuation bias in the OLS estimates due to measurement error in the binary quality variable, or it could indicate that households induced to receive good quality by the instrument (compliers) have above-average returns to quality. The Kleibergen-Paap F-statistics range from 14 to 36 across specifications, above the Stock and Yogo (2005) critical value for 15% maximum bias. The consistency in sign and significance between OLS and IV supports a causal interpretation of the quality effect on appliance adoption in rural Zambia.

Table 2.7: Rural: Impacts of Electricity Quality

	Aggregated Appliance Capacity (W)		
	(1)	(2)	(3)
<i>Panel A. OLS</i>			
Electricity Quality	854.85 *** (307.17)	764.73 ** (376.89)	797.78 ** (392.59)
<i>Panel B. IV</i>			
Electricity Quality	2,110.35 *** (711.84)	1,228.63 * (737.18)	1,329.28 * (748.90)
First stage F-statistic	36.52	26.81	14.68
Observations	272	272	272
R^2 (Panel A)	0.134	0.301	0.302
Clusters	135	135	135
Demographic & Income	No	Yes	Yes
Climate & Spatial Controls	No	No	Yes

Notes: This table presents estimates for rural electrified households. Panel B instruments for Electricity Quality using Distance to the Nearest Substation. F-statistic is for the heteroskedasticity and cluster-robust Kleibergen-Paap weak instrument test. Robust standard errors clustered by EAs in parentheses.

*** Significant at the 1 percent level.

** Significant at the 5 percent level.

* Significant at the 10 percent level.

Table 2.8: Rural: First Stage Estimation

	Electricity Quality		
	(1)	(2)	(3)
<i>Panel A. Distance to the Nearest Substation (10km)</i>			
Distance	-0.12 *** (0.02)	-0.1 *** (0.02)	-0.11 *** (0.02)
<i>Panel B. Shortest Distance to the High Voltage Grid Lines (10km)</i>			
Distance	-0.02 (0.02)	-0.015 (0.02)	-0.01 (0.02)
Observations	272	272	272
Clusters	135	135	135
Demographic & Income	No	Yes	Yes
Climate & Spatial Controls	No	No	Yes

2.6.2 Household Appliance Ownership

The aggregate results in Section 6.1 may mask heterogeneous responses across appliance types. To examine which appliances drive the overall effect, I estimate appliance-level Probit models for both urban and rural areas. Table 2.9 reports the baseline ownership rates among households with poor electricity quality alongside the estimated average marginal effects, organized by appliance load category. Figures 2.4 and 2.5 present these results graphically.

Very Low Load Appliances (3–25 W). These include torches, mobile phone chargers, LED bulbs, incandescent bulbs, and smartphone chargers. In urban areas, only LED bulb ownership shows a statistically significant response, increasing by 6.4 percentage points from a baseline of 25.3%. The remaining very low load appliances show small and insignificant effects: smartphone charger ownership increases by 4.5 percentage points, while torches and mobile phone chargers show changes of less than 2 percentage points. This pattern is consistent with the nature of these devices. They draw minimal power and function adequately even under unstable supply, so electricity quality is not the binding constraint on their adoption. In rural areas, LED bulb and smartphone charger ownership increase significantly (19.4 and 7.7 percentage points, respectively), and incandescent bulb ownership rises by 11.9 percentage points. Regular mobile phone charger ownership declines by 11.0 percentage points, likely reflecting substitution toward smartphones: as quality improves, households upgrade from basic phones to smartphones, consistent with the simultaneous increase in smartphone charger ownership. The larger rural effects likely reflect lower baseline ownership rates, leaving more room for quality-driven adoption.

Low Load Appliances (75–150 W). This category spans lighting (compact fluorescent bulbs, fluorescent tubes), entertainment (radios, TVs, DVD players), and small household devices (fans, rice cookers, computers). The results reveal that quality effects concentrate on appliances used for sustained daily activities. Fan ownership increases by 10.6 percentage points in urban areas and 9.5 percentage points in rural areas. VCD/DVD player ownership rises by 10.1 and 22.2 percentage

Table 2.9: Appliance Baseline Ownership and Average Marginal Effects of Electricity Quality

Appliance	Load	Watts	Urban		Rural	
			Baseline (%)	ME (pp)	Baseline (%)	ME (pp)
Torch/flashlight/lantern	Very Low	3	19.5	-1.4	32.4	-9.5
Regular mobile phone charger	Very Low	5	65.9	0.7	66.1	-11.0*
LED Light Bulb	Very Low	10	25.3	6.4**	10.2	19.4***
Incandescent Light Bulb	Very Low	25	37.8	-3.3	20.3	11.9**
Smartphone charger	Very Low	25	45.1	4.5	12.4	7.7**
Compact Fluorescent Light Bulb	Low	75	28.3	3.5	14.1	2.5
Radio/CD Players/sound system	Low	75	62.9	3.6	41.5	12.2*
Fluorescent Tube	Low	80	24.1	3.0	10.2	4.5
Black & White TV	Low	100	2.8	-2.2***	1.1	1.0
Electric sewing machine	Low	100	2.8	-1.2	—	—
Fan	Low	100	40.4	10.6***	14.7	9.5**
Rice cooker	Low	100	6.5	0.1	2.3	0.3
VCD/DVD	Low	100	46.6	10.1***	18.2	22.2***
Flat color TV	Low	120	35.6	2.8	9.6	8.9**
Computer	Low	150	16.8	1.0	7.9	-1.3
Regular Color TV	Low	150	46.4	7.4**	22.6	26.1***
Freezer	Medium	200	32.8	5.2*	15.3	0.3
Refrigerator	Medium	350	52.4	16.7***	14.7	22.6***
Electric food processor/blender	Medium	400	5.8	-1.6	—	—
Hair dryer	Medium	500	4.8	2.2	4.0	-2.6
Washing machine	Medium	500	4.0	1.4	—	—
Electric Water Pump	Medium	750	2.0	-0.1	4.0	-2.6
Electric Iron	High	1,000	61.7	20.3***	18.6	30.9***
Electric hot water pot/kettle	High	1,000	21.3	2.2	6.2	4.8
Microwave oven	High	1,000	18.8	0.9	7.3	-0.8
Electric water heater	High	1,500	7.3	-2.0	1.7	5.9**
Indoor Air cooler	High	1,500	1.5	2.8**	—	—
Air Conditioner (AC)	Very High	2,000	3.8	0.3	—	—
Solar based water heater	Very High	2,000	0.3	0.2	—	—
Space Heater	Very High	2,000	2.5	0.2	1.7	1.5

Notes: Baseline is ownership rate (%) among electrified households with poor electricity quality.

ME is average marginal effect in percentage points from Probit.

— indicates insufficient rural observations for estimation.

*** Significant at the 1 percent level. ** Significant at the 5 percent level. * Significant at the 10 percent level.

points, respectively. Regular color TV ownership increases by 7.4 percentage points in urban areas and 26.1 percentage points in rural areas. These entertainment and comfort appliances benefit from stable electricity because their value to the household depends on regular, uninterrupted use. By contrast, black and white TV ownership declines by 2.2 percentage points in urban areas, consistent with households upgrading to color televisions when quality improves. Rice cooker and computer ownership show no significant response in either setting, suggesting that cooking and computing decisions are driven by factors other than supply reliability.

Medium Load Appliances (200–750 W). Refrigerators stand out in this tier. Urban refrigerator ownership increases by 16.7 percentage points from a baseline of 52.4%, and rural refrigerator ownership rises by 22.6 percentage points from a baseline of 14.7%. These are among the largest effects in the entire table. Refrigerators require continuous, stable power to maintain temperature, and even brief outages can spoil stored food. Households are therefore unlikely to invest in a refrigerator unless they expect reliable supply. Freezer ownership shows a modest increase of 5.2 percentage points in urban areas, though the rural estimate is not significant. The remaining medium-load appliances (blenders, hair dryers, washing machines, water pumps) show small and mostly insignificant effects, with limited rural observations precluding estimation for several of them.

High Load Appliances (1,000–1,500 W). The electric iron is the most responsive appliance in the entire sample. Urban iron ownership increases by 20.3 percentage points from a baseline of 61.7%, and rural iron ownership increases by 30.9 percentage points from a baseline of 18.6%. Unlike a kettle that boils water in a few minutes, ironing a load of clothes typically requires 30 to 60 minutes of continuous use, making it difficult to complete during short windows between outages. Irons are also sensitive to voltage fluctuations, which can prevent the heating element from reaching the required temperature. The indoor air cooler shows a small but significant urban increase of 2.8 percentage points. Electric water heater ownership increases by 5.9 percentage points in rural areas, consistent with the logic that appliances requiring sustained heating benefit from reliable supply. Electric kettles and microwave ovens, despite being common appliances, show no significant

response to quality in either setting (2.2 and 0.9 percentage points in urban areas, respectively). This is consistent with their use pattern: kettles and microwaves operate for brief heating cycles of a few minutes that households can schedule around outage periods.

Very High Load Appliances (2,000 W). Air conditioners, solar water heaters, and space heaters constitute this tier. The effects are uniformly small and statistically insignificant. Air conditioner ownership increases by only 0.3 percentage points in urban areas, and the rural estimate cannot be obtained due to insufficient observations. Solar water heater ownership shows no meaningful response (0.2 pp urban). Space heater ownership shows similarly negligible changes. For these appliances, the binding constraint is purchase cost and affordability rather than electricity quality. Even with perfectly reliable power, few households in this income distribution can afford the upfront investment and operating costs of 2,000-watt appliances.

In summary, three patterns emerge across load tiers. First, the largest effects concentrate on appliances that require continuous or stable power for their core function, particularly refrigerators (+16.7 pp urban, +22.6 pp rural) and electric irons (+20.3 pp urban, +30.9 pp rural), rather than on the highest-wattage appliances. Second, entertainment and comfort devices that depend on sustained daily use (fans, TVs, DVD players) show moderate and significant effects, while appliances used in brief, flexible episodes (kettles, rice cookers, microwaves) do not respond to quality. Third, very high-wattage appliances show near-zero effects because their adoption is constrained by household income, not supply reliability. These findings are consistent with a framework in which electricity quality reduces the effective operating cost of appliances whose value depends on uninterrupted power, while leaving unaffected those appliances whose use can be scheduled around outages or whose purchase price is the binding constraint.

2.6.3 Comparison to Electricity Access

To benchmark the magnitude of quality effects, I estimate the same specifications with electricity access (grid connection versus no connection) as the treatment variable, using the full sample




Figure 2.4: Average Marginal Effects on Household Appliance Ownership: Percentage Points Change




Figure 2.5: Average Marginal Effects on Household Appliance Ownership: Percentage Change

of urban and rural households.

Urban Households. Table 2.10 reports OLS estimates for all urban households. Gaining electricity access increases aggregated appliance capacity by 1,728.32 watts without controls and 1,731.47 watts with the full set of controls. The access effect is roughly six times larger than the quality effect (308.04 watts). This ordering is expected: access represents a shift from zero electricity use to a connected household, while quality captures an improvement within the population of already-connected households.

Rural Households. Table 2.11 presents OLS (Panel A) and IV (Panel B) estimates for rural households. The OLS access effect ranges from 1,766.27 watts without controls to 1,784.27 watts with the full set of controls, all significant at the 1% level. The IV estimates yield coefficients of 1,898.35 watts without controls and 1,515.26 watts with the full set of controls. First-stage F-statistics range from 12 to 16, above conventional weak-instrument thresholds. The exclusion restriction requires that substation distance affects appliance capacity only through its effect on access; as in the quality regressions, I control for distance to the nearest city to absorb any direct effect of remoteness on appliance availability. The consistency between OLS and IV confirms a large, robust access effect in rural areas.

The comparison between access and quality effects illustrates how the two operate at different margins. For the majority of Zambian households that remain unconnected, gaining grid access is the first-order priority, yielding roughly 1,500 to 1,800 watts of appliance capacity. For the approximately 1.2 million households already on the grid (Luzi et al., 2019), improving service quality can unlock an additional 300 to 800 watts of capacity, equivalent to adding a refrigerator or several smaller appliances to the household portfolio. The quality margin is smaller in absolute terms, but it applies to a population that has already made the initial investment in grid connection and is positioned to benefit from incremental improvements in reliability.

Table 2.10: Urban: Impacts of Electricity Access

	Aggregated Appliance Capacity (W)		
	(1)	(2)	(3)
Electricity Access	1,728.32 *** (117.77)	1,731.07 *** (116.66)	1,731.47 *** (115.97)
Observations	1,706	1,706	1,706
R^2	0.278	0.484	0.484
Clusters	346	346	346
Demographic & Income	No	Yes	Yes
Climate Controls	No	No	Yes

Notes: This table presents OLS estimates for all urban households.

Robust standard errors clustered by EAs in parentheses.

*** Significant at the 1 percent level.

** Significant at the 5 percent level.

* Significant at the 10 percent level.

Table 2.11: Rural: Impacts of Electricity Access

	Aggregated Appliance Capacity (W)		
	(1)	(2)	(3)
<i>Panel A. OLS</i>			
Electricity Access	1,766.27 *** (206.54)	1,769.47 *** (207.78)	1,784.27 *** (210.46)
<i>Panel B. IV</i>			
Electricity Access	1,898.35 *** (306.96)	1,497.72 *** (335.08)	1,515.26 *** (348.96)
First stage F-statistic	15.61	13.82	12.45
Observations	1,658	1,658	1,658
R^2 (Panel A)	0.350	0.421	0.422
Clusters	538	538	538
Demographic & Income	No	Yes	Yes
Climate & Spatial Controls	No	No	Yes

Notes: This table presents estimates for all rural households. Panel B instruments for Electricity Access using Distance to the Nearest Substation. F-statistic is for the heteroskedasticity and cluster-robust Kleibergen-Paap weak instrument test. Robust standard errors clustered by EAs in parentheses.

*** Significant at the 1 percent level.

** Significant at the 5 percent level.

* Significant at the 10 percent level.

2.6.4 Heterogeneity by Income and Gender

The aggregate quality effect estimated in Section 6.1 may differ across household types. Wealthier households can afford a broader range of appliances and may respond more strongly to quality improvements, or alternatively, lower-income households may face more binding quality constraints that, once relaxed, allow them to adopt their first major appliance. Similarly, the gender of the household head may shape the response if male- and female-headed households differ in their appliance priorities or purchasing power.

I test for heterogeneity along two dimensions: household income (proxied by monthly expenditure, split at the within-subsample median) and gender of the household head. Table 2.12 reports results from interaction models estimated on the full urban electrified sample. I focus the interaction analysis on urban households because the rural sample ($N = 272$) lacks sufficient power for interaction models. Column (1) interacts the quality indicator with an above-median income dummy. The quality effect for below-median households is 246.6 watts ($p = 0.02$), while the interaction term is small and not statistically significant (-24.2 , $p = 0.92$), indicating that the quality effect for above-median households is similar in magnitude. Column (2) interacts quality with a female-headed household indicator. The quality effect for male-headed households is 231.6 watts, and the interaction term is small and insignificant (-97.7 , $p = 0.66$), suggesting no meaningful differential effect by gender.

Table 2.13 presents split-sample estimates as a complementary check. The quality coefficient is similar across subgroups: 163.9 watts for below-median income households and 238.0 watts for above-median households, 222.4 watts for male-headed households and 159.3 watts for female-headed households. None of the split-sample estimates are individually significant, which is expected given the halved sample sizes. The split-sample point estimates differ from the interaction estimates because the split regressions allow all coefficients, not just the quality effect, to vary across subgroups. The rural split-sample estimates are imprecise due to small cell sizes and are reported for completeness.

The absence of significant heterogeneity suggests that electricity quality operates as a supply-

Table 2.12: Urban: Heterogeneity in Quality Effects (Interaction Models)

	(1)	(2)
	Income Interaction	Gender Interaction
Electricity Quality	246.6** (106.0)	
Above-Median Income	1380.6*** (278.7)	
Quality $\times$ Above-Median Income	-24.2 (232.5)	
Electricity Quality		231.6 (161.1)
Female-Headed Household		-21.3 (207.2)
Quality $\times$ Female Head		-97.7 (219.5)
Observations	1066	1066
R^2	0.397	0.334

Notes: Sample is urban electrified households. Dependent variable is aggregated appliance capacity (watts). All specifications include province fixed effects, demographic controls, income, and climate zone indicators. Standard errors clustered at the enumeration area level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 2.13: Heterogeneity in Quality Effects by Income and Gender (Split Sample)

	By Income		By Gender of HH Head	
	Below Median	Above Median	Male Head	Female Head
<i>Panel A. Urban</i>				
Electricity Quality	163.9 (101.0)	238.0 (231.1)	222.4 (164.1)	159.3 (180.9)
Observations	533	533	796	270
R^2	0.175	0.254	0.324	0.466
<i>Panel B. Rural</i>				
Electricity Quality	267.1 (179.0)	304.8 (547.1)	611.9 (442.2)	293.2 (612.7)
Observations	136	136	217	55
R^2	0.241	0.306	0.352	0.264

Notes: Dependent variable is aggregated appliance capacity (watts). All specifications include province fixed effects, demographic controls (household size, gender of head), income (monthly expenditure), and climate zone indicators.

Standard errors clustered at the enumeration area level in parentheses. Income split at the within-subsample median.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

side constraint that affects connected households broadly, rather than interacting with purchasing power or gender-specific preferences. Quality improvements appear to expand the feasible set of appliances for all electrified households, not just those at a particular point in the income distribution.

2.7 Conclusion

This paper estimates the effects of electricity quality on household appliance adoption in Zambia, using household-level data from the World Bank Multi-Tier Framework survey. I distinguish between electricity access, whether a household has any source of electricity, and electricity quality, whether electrified households receive reliable, continuous service. For rural households, I instrument for electricity quality using the distance to the nearest distribution substation to address potential endogeneity.

Four findings emerge. First, improved electricity quality increases aggregated appliance capacity by 308 watts in urban areas and 798 watts in rural areas under OLS, with IV estimates for rural households yielding 1,329 watts. The larger rural coefficients are consistent with more binding supply constraints in areas where grid infrastructure is less developed. Second, appliance-level Probit estimates reveal that the aggregate effect is driven by a specific subset of appliances. The largest ownership increases occur for appliances that require continuous or stable power, particularly refrigerators (+16.7 pp urban, +22.6 pp rural) and electric irons (+20.3 pp urban, +30.9 pp rural). Entertainment and comfort devices used in sustained daily activities, such as fans, color televisions, and DVD players, also respond significantly. Appliances used in brief, flexible episodes (kettles, rice cookers, microwaves) and very high-wattage appliances constrained by purchase cost (air conditioners, space heaters) show little response. This pattern indicates that the relevant distinction is not wattage per se, but whether an appliance’s value to the household depends on uninterrupted power. Third, the quality effect does not differ significantly by household income or by the gender of the household head, suggesting that poor service reliability constrains appliance adoption broadly across the population of connected households. Fourth, electricity access increases appliance capacity by roughly 1,500 to 1,800 watts, approximately six times the quality effect in urban areas, confirming that access remains the first-order margin for unconnected households.

These results carry direct policy implications. For the majority of Zambian households that remain unconnected, expanding grid access is the first-order priority. However, for the approximately 1.2 million households already on the grid ([Luzi et al., 2019](#)), investments in distribution infrastructure that improve service reliability can yield meaningful welfare gains by enabling the adoption of appliances that households currently forgo due to poor power quality. The appliance-level results suggest that these gains will concentrate among refrigerators, irons, fans, and televisions, appliances that directly affect food preservation, household productivity, comfort, and access to information. As grid coverage expands across Sub-Saharan Africa, the quality of service delivered through existing connections will become an increasingly relevant policy margin.

Several limitations of this study point to directions for future work. My analysis relies on

cross-sectional data, which limits my ability to trace the dynamic process of appliance adoption following quality improvements. Panel data covering households before and after infrastructure upgrades would allow for sharper causal identification. The binary quality measure, while transparent, collapses multiple dimensions of service reliability into a single indicator; future work could exploit continuous variation in outage frequency, duration, and voltage stability separately. Finally, my estimates capture the effect of quality on ownership but not on utilization. A household that owns a refrigerator but runs it intermittently due to unreliable power derives less welfare than one that operates it continuously. Linking quality improvements to appliance usage intensity and household expenditure patterns would provide a more complete picture of the welfare returns to infrastructure investment.

References

- Allcott, Hunt, Allan Collard-Wexler, and Stephen D O’Connell.** 2016. “How do electricity shortages affect industry? Evidence from India.” *American Economic Review*, 106(3): 587–624.
- Beck, Hylke E, Niklaus E Zimmermann, Tim R McVicar, Noemi Vergopolan, Alexis Berg, and Eric F Wood.** 2018. “Present and future Köppen-Geiger climate classification maps at 1-km resolution.” *Scientific data*, 5(1): 180214.
- Bhatia, Mikul, and Niki Angelou.** 2015. “Beyond Connections: Energy Access Redefined.” World Bank, ESMAP Technical Report 008/15.
- Blimpo, Moussa P, and Malcolm Cosgrove-Davies.** 2019. *Electricity access in Sub-Saharan Africa: Uptake, reliability, and complementary factors for economic impact*. World Bank Publications.
- Bollen, Math HJ.** 2000. *Understanding Power Quality Problems: Voltage Sags and Interruptions*. Vol. 3, IEEE press New York.
- Burgess, Robin, Michael Greenstone, Nicholas Ryan, and Anant Sudarshan.** 2020. “The consequences of treating electricity as a right.” *Journal of Economic Perspectives*, 34(1): 145–169.
- Burlig, Fiona, and Louis Preonas.** 2024. “Out of the darkness and into the light? Development effects of rural electrification.” *Journal of Political Economy*, 132(9): 2937–2971.

- Chakravorty, Ujjayant, Martino Pelli, and Beyza Ural Marchand.** 2014. “Does the quality of electricity matter? Evidence from rural India.” *Journal of Economic Behavior & Organization*, 107: 228–247.
- Deutschmann, Joshua W., Agnieszka Postepska, and Leopold Sarr.** 2021. “Measuring willingness to pay for reliable electricity: Evidence from Senegal.” *World Development*, 138: 105209.
- Dinkelman, Taryn.** 2011. “The effects of rural electrification on employment: New evidence from South Africa.” *American Economic Review*, 101(7): 3078–3108.
- Gertler, Paul J, Ori Shelef, Catherine D Wolfram, and Alan Fuchs.** 2016. “The demand for energy-using assets among the world’s rising middle classes.” *American Economic Review*, 106(6): 1366–1401.
- Jimenez Mori, Raul.** 2021. “It’s not price; it’s quality. Satisfaction and price fairness perception.” *World Development*, 139: 105302.
- Lee, Kenneth, Edward Miguel, and Catherine Wolfram.** 2016. “Appliance ownership and aspirations among electric grid and home solar households in rural Kenya.” *American Economic Review: Papers & Proceedings*, 106(5): 89–94.
- Lee, Kenneth., Edward. Miguel, and Catherine. Wolfram.** 2020. “Experimental evidence on the economics of rural electrification.” *Journal of Political Economy*, 128(4): 1523–1565.
- Li, Mingquan, Rui Shan, Mauricio Hernandez, Varun Mallampalli, and Dalia Patiño-Echeverri.** 2019. “Effects of population, urbanization, household size, and income on electric appliance adoption in the Chinese residential sector towards 2050.” *Applied Energy*, 236: 293–306.
- Lipscomb, Molly, A Mushfiq Mobarak, and Tania Barham.** 2013. “Development effects of electrification: Evidence from the topographic placement of hydropower plants in Brazil.” *American Economic Journal: Applied Economics*, 5(2): 200–231.

- Luzi, Lucia, Yunhui Lin, Brian Bonsuk Koo, Dana Rysankova, and Elisa Portale.** 2019. “Zambia – Beyond Connections: Energy Access Diagnostic Report Based on the Multi-Tier Framework.” World Bank, ESMAF.
- McRae, Shaun.** 2010. “Reliability, appliance choice, and electricity demand.” *Department of Economics, University of Michigan, Working Paper*.
- Meeks, Robyn C., Arstan Omuraliev, Ruslan Isaev, and Zhenxuan Wang.** 2023. “Impacts of electricity quality improvements: Experimental evidence on infrastructure investments.” *Journal of Environmental Economics and Management*, 120: 102838.
- Nag, Suryadeepo.** 2024. “Does reliable electricity mean lesser agricultural labor wages? Evidence from Indian villages.” *Energy Policy*, 190: 114151.
- Stock, James H., and Motohiro Yogo.** 2005. “Testing for Weak Instruments in Linear IV Regression.” In *Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg*, ed. Donald W. K. Andrews and James H. Stock, 80–108. Cambridge University Press.
- Wolfram, Catherine, Orie Shelef, and Paul Gertler.** 2012. “How will energy demand develop in the developing world?” *Journal of Economic Perspectives*, 26(1): 119–138.

Chapter 3

Weather Shocks, Supply Dynamics, and Price Persistence: Evidence from U.S. Lettuce Markets

Abstract

This paper constructs a weekly dataset that integrates USDA shipment data, FOB (free on board) prices, and weather data to study weather-driven supply and price dynamics in U.S. lettuce markets from 2010 to 2026. Using an empirical strategy that combines market-level aggregation for prediction with district-level panel variation for identification, I document three results. First, simple linear models outperform tree-based machine learning methods. OLS achieves substantially lower prediction error than XGBoost and Random Forest (RMSE \$4.36 versus \$6.67 and \$6.81), reflecting strong price persistence and a data-generating process that is largely linear conditional on lagged prices. Second, weather effects on prices are present but attenuated in aggregate data. In district-level panels, each additional freezing degree day increases weekly prices by \$0.85 ($p < 0.001$), while these effects are diluted in market-level regressions due to price persistence and spatial aggregation. Third, these patterns are robust across lettuce varieties, including Romaine, and across alternative validation strategies. These results show that downstream market outcomes, shipment volumes and FOB prices, can be used to identify weather-driven supply responses even in the absence of direct yield measures.

3.1 Introduction

Fresh produce supply chains are exposed to weather risk, but translating weather shocks into quantitative supply and price signals remains difficult. For most fresh vegetables, the data infrastructure that enables high-frequency crop monitoring in grain markets simply does not exist. There are no futures contracts, no weekly crop progress reports, and no county-level yield estimates. Buyers, distributors, and supply chain managers rely instead on trade reports, informal market intelligence, and experience. However, it remains unclear whether downstream market data can be used to recover supply responses to weather shocks.

This paper studies how weather shocks affect supply and prices in the U.S. iceberg lettuce market. Lettuce is the largest fresh vegetable crop in the United States by volume and is produced year-round through a seasonal rotation between California (spring through fall) and Arizona (winter). I construct a weekly dataset that integrates USDA truck shipment volumes, FOB shipping point prices, and gridded weather data from PRISM, covering six production districts over the period 2010–2026.

My empirical strategy combines market-level aggregation for prediction with district-level panel variation for identification. At the market level, I aggregate all districts into a single weekly time series and estimate price and volume prediction models using OLS, XGBoost, and Random Forest. At the district level, I retain a three-district panel and use fixed effects to identify the causal effect of local weather shocks on prices. Both layers are necessary because aggregation improves data coverage (from 72% to 98% of weeks with observed prices) but dilutes the weather signal.

First, simple linear models outperform tree-based machine learning methods in level prediction tasks and perform comparably in change models. The best market-level price model achieves an out-of-sample RMSE of \$4.36, a 14.7% improvement over a naive autoregressive baseline. XGBoost and Random Forest perform substantially worse than the baseline in level models, but in change models all methods converge to similar performance. This reflects a data-generating process dominated by strong price persistence and approximate linearity, under which additional model complexity does

not improve predictive performance and reduces robustness in extreme periods.

Second, weather effects on prices are present but obscured by price persistence in aggregate data. In level models, the three price lags together account for nearly half of total feature importance, leaving weather variables with little apparent explanatory power. When I model week-to-week price changes instead, removing the autoregressive component by construction, weather variables become substantially more important. At the district level, each additional freezing degree day increases the local weekly price by \$0.85 ($p < 0.001$), while this effect is attenuated in market-level aggregates. The mechanism runs from weather to local supply disruption to local price increases, which then persist through the autoregressive structure of the market.

Third, all results are robust across lettuce varieties. Replicating the analysis with Romaine lettuce produces consistent coefficient signs and comparable, and in fact lower, prediction errors (RMSE of \$2.69 vs. \$4.36 for Iceberg).

The rest of the paper is organized as follows. Section 3.2 reviews the related literature. Section 3.3 describes the institutional structure of the U.S. lettuce market. Section 3.4 presents the data sources and variable construction. Section 3.5 lays out the empirical strategy. Section 3.6 reports the main results, including the underlying mechanism linking weather, supply, and prices. Section 3.7 presents robustness checks, and Section 3.8 concludes.

3.2 Literature Review

This paper contributes to four related strands of literature: the econometrics of weather and agricultural production, the transmission of weather shocks to prices, climate impacts on specialty crops, and the use of machine learning in applied economics.

A large literature studies how weather affects agricultural production, primarily through crop yields. Schlenker and Roberts (2009) show that the relationship between temperature and yields of major U.S. crops is highly nonlinear, with extreme heat driving large output losses. Subsequent work extends this framework using improved weather measures and broader datasets (e.g., Roberts, Schlenker and Eyer, 2013), and finds that weather variation explains a substantial share of yield

fluctuations. More recent studies document that climate change has already reduced global agricultural productivity (e.g., [Lobell, Schlenker and Costa-Roberts, 2011](#); [Ortiz-Bobea et al., 2021](#)). However, this approach relies critically on high-frequency yield data, which are typically available only for major row crops at annual frequency and are not available for most fresh vegetables such as lettuce.

A related literature examines how weather shocks translate into market outcomes such as prices. Early evidence is provided by [Roll \(1984\)](#), who show that futures prices respond to weather shocks. [Roberts and Schlenker \(2009\)](#) and [Roberts and Schlenker \(2013\)](#) estimate supply and demand elasticities using weather-driven yield shocks, finding that demand is highly inelastic and that small supply disruptions generate large price movements. More recent work shows that prices also respond to expectations about weather before harvest ([Letta, Montalbano and Pierre, 2022](#)). Most of this literature relies on yield shocks to identify price effects. In contrast, I show that weather effects on prices can be identified directly from market outcomes, and that these effects are attenuated in aggregate data due to price persistence.

Despite their economic importance and greater sensitivity to short-run weather variation, perishable specialty crops have received relatively little attention in this literature. [Lobell, Cahill and Field \(2007\)](#) estimate climate effects on yields of several California crops, including lettuce, using annual data, while [Kistner et al. \(2018\)](#) document that specialty crops are particularly sensitive to extreme weather events such as freezes and excess rainfall. These studies focus on annual yield responses. I extend this work by using weekly data to study short-run dynamics and by focusing on market outcomes, prices and shipment volumes, which provide a high-frequency proxy for supply in settings where yield data are unavailable.

Finally, a growing literature examines when machine learning methods improve upon traditional econometric approaches. [Mullainathan and Spiess \(2017\)](#) emphasize that machine learning excels in prediction tasks with complex, nonlinear relationships, while economic structure remains essential for interpretation. [Storm, Baylis and Heckelei \(2020\)](#) show that machine learning methods perform well in high-dimensional agricultural prediction settings, and [Crane-Droesch \(2018\)](#) find

gains in crop yield prediction when nonlinearities are important. My results provide a counterpoint. In the lettuce market, tree-based methods (XGBoost and Random Forest) underperform both OLS and a naive baseline. This reflects the fact that price dynamics are dominated by a strong autoregressive component that is effectively captured by linear models. Once persistence is removed by differencing, all methods perform similarly, suggesting that nonlinearities do not play a large role in predictive performance in this setting. More broadly, my findings contribute to understanding when additional model flexibility improves predictive performance and when simple linear models are sufficient.

3.3 Market Background

The U.S. lettuce market is highly geographically concentrated, with California and Arizona accounting for over 95% of domestic production ([USDA National Agricultural Statistics Service, 2023](#)). Production is organized around a seasonal geographic rotation between two primary regions.

Salinas-Watsonville, California is the dominant production district. Located in the Salinas Valley on the central California coast, it benefits from a mild maritime climate well suited to lettuce during the spring, summer, and fall months (roughly April through November, corresponding to weeks 15–50 of the year). It supplies the majority of national shipments during this period.

Western Arizona (Yuma) takes over during the winter months (roughly November through March, weeks 45–15). The region’s warm desert climate allows lettuce production when Salinas is too cold, and shipment volumes are comparable to Salinas during peak winter weeks.

Imperial Valley, California serves as a supplementary desert production district, with a seasonal pattern similar to Yuma but smaller volumes. Three other California districts, Santa Maria, Oxnard, and the San Joaquin Valley, contribute smaller and more intermittent shipments.

The annual rotation between California and Arizona is largely calendar-driven and follows a stable pattern over time. Production transitions occur around weeks 15–20 (as Salinas ramps up and Arizona winds down) and weeks 45–50 (as Salinas winds down and Arizona ramps up). These transition periods are associated with temporary supply contractions and elevated price volatility.

They are also particularly sensitive to weather shocks, as supply is already constrained during the shift between regions. Figure 3.1 illustrates this seasonal rotation. Salinas dominates shipments from roughly weeks 15–50, while Western Arizona dominates the winter weeks. Price spikes are concentrated around the transition periods, when aggregate supply is temporarily reduced.




Figure 3.1: Seasonal patterns in shipment volume and prices by production district.

Lettuce has a short growing cycle of approximately 60 to 90 days from planting to harvest. Unlike annual grain crops, production occurs on a continuous, rolling basis throughout the season, with fields replanted multiple times per year. As a result, there is no single harvest event that determines total output. Instead, supply in any given week reflects planting decisions made two to three months earlier, realized growing conditions, and the pace of harvest operations. This structure generates a natural lag between weather shocks and realized supply.

This production environment has an important implication for empirical analysis. High-frequency yield data, which are available for many grain crops at the county level, do not exist for lettuce. There is no weekly or even monthly measure of production by district. Shipment data, which record the volume of lettuce loaded onto trucks at the point of origin, provide the closest available proxy for realized supply. Because lettuce is highly perishable, with a shelf life of roughly one to two weeks, storage is limited and most production is shipped shortly after harvest. Shipment volumes therefore provide a high-frequency measure of supply that is central to my empirical strategy.

3.4 Data

3.4.1 Sources

I combine three publicly available data sources covering January 2010 through April 2026.

Shipment volume. Weekly truck shipment reports from the USDA Agricultural Marketing Service ([USDA Agricultural Marketing Service, 2026b](#)) record the volume of each commodity shipped by truck from each production district, measured in units of 10,000 pounds. These reports provide the highest-frequency measure of supply available for lettuce. Unlike grain crops, which have annual county-level yield data, lettuce is harvested continuously within each growing season and no comparable high-frequency production data exist. Shipment volume therefore serves as a proxy for realized supply: because lettuce is highly perishable and cannot be stored for more than a few days, most production is shipped within the week it is harvested.

Prices. Daily FOB (free on board) shipping point prices from USDA Market News ([USDA Agricultural Marketing Service, 2026a](#)) report the price at which lettuce is sold at the point of origin, measured in dollars per carton and before transportation costs. I focus on FOB prices rather than retail prices because they more directly reflect production-side conditions. Prices are reported by district, commodity, and package type; I aggregate daily observations to weekly averages.

Weather. Daily temperature and precipitation data come from the PRISM Climate Group at Oregon State University ([Daly et al., 2008](#)), which provides gridded climate data at 4 km resolution for the contiguous United States. For each production district, I extract daily maximum temperature, minimum temperature (both in degrees Celsius), and precipitation (in millimeters) at the district centroid and aggregate these variables to weekly measures. This construction aligns weather exposure with the timing of shipment and pricing outcomes.

3.4.2 Production Districts

The data cover six lettuce production districts in California and Arizona: Salinas-Watsonville, Western Arizona (Yuma), Imperial Valley, Santa Maria, Oxnard, and the San Joaquin Valley.

Together, these districts account for the vast majority of U.S. lettuce production.

District activity and reporting vary substantially over time. Salinas-Watsonville, the largest district, ships in 560 of the 849 sample weeks and reports prices in 91% of those weeks. Western Arizona ships in 429 weeks, primarily during the winter months, with 82% price coverage. The remaining districts have more limited and irregular reporting: Oxnard reports prices in only 2% of its active weeks, Santa Maria in 69%, and the San Joaquin Valley in 76%.

This uneven coverage has two implications. First, price data are missing non-randomly across districts and time. Second, relying on any single district would substantially reduce the usable sample. These features motivate the dual-layer empirical strategy described in Section 3.5.1, which combines market-level aggregation to maximize coverage with district-level variation for identification.

3.4.3 Variable Construction

Market-level price. For each week, I construct a shipment-weighted average FOB price across all districts that report a price (Equation 3.1). Weighting by shipment volume ensures that districts with larger production have greater influence on the market price.

Market-level weather. Weather variables are aggregated to the market level using shipment weights. Market-level weather reflects conditions in the districts that are actively supplying the market in a given week, rather than a fixed geographic average. When production shifts seasonally from Salinas to Yuma, the market-level weather measures shift accordingly.

Extreme weather. Following the degree-day approach standard in the agricultural economics literature (Schlenker and Roberts, 2009), I construct two continuous measures of extreme weather exposure from the daily data. The variable dd_freeze is the weekly sum of degrees below 0°C (i.e., $\sum_d \max(0 - T_d^{\min}, 0)$ over days d in the week), capturing both the occurrence and severity of freezing conditions. The variable dd_heat is the analogous measure above 40°C. The variable $heavy_rain$ is an indicator equal to one if total weekly precipitation exceeds 25 mm; I retain an indicator here because the continuous precipitation level (ppt_total) is already included as a sep-

arate variable. At the market level, degree-day measures are aggregated using shipment weights, consistent with the other weather variables.

Lagged variables. For prices, shipment volume, maximum temperature, and precipitation, I include lags at 1, 2, and 4 weeks. I also include 4-week rolling averages of temperature and precipitation. These lagged and smoothed weather measures allow the model to capture delayed effects of weather on supply, consistent with the 60–90 day lettuce growth cycle.

Rolling averages. Four-week rolling means of volume and weather variables capture the recent background state of the market.

Market structure. I include the coverage ratio C_t (Equation 3.2) and the number of active districts n_t as controls to capture variation in data quality and the geographic concentration of supply.

Seasonality. Month and week-of-year enter as continuous variables to capture the strong seasonal pattern in production and prices.

3.4.4 Summary Statistics

Table 3.1 reports summary statistics for the market-level dataset. The mean weekly iceberg lettuce FOB price is \$16.27 per carton, with a standard deviation of \$11.80. The distribution is right-skewed, with a maximum of \$94.33 reached during the 2022 Impatiens Necrotic Spot Virus (INSV) outbreak. Mean weekly shipment volume is 4,422 (10,000 lb units). Extreme weather events are relatively infrequent at the market level: freezing conditions (degree days below $0^\circ\text{C} > 0$) occur in 14 of 849 weeks (1.6%), extreme heat (degree days above $40^\circ\text{C} > 0$) in 6 weeks (0.7%), and heavy rain in 68 weeks (8.0%). The average coverage ratio is 0.93, indicating that most shipment volume is associated with observed prices. In addition, prices are observed in 98.2% of weeks, ensuring a nearly complete time series. On average, 2.9 production districts are active in a given week, where an active district is defined as one with positive shipment volume. This reflects the seasonal concentration of supply across regions.

Table 3.1: Summary statistics for market-level data, 2010–2026

Variable	Mean	SD	Min	Max
Price (\$/carton)	16.27	11.80	5.99	94.33
Volume (10,000 lbs)	4,422	749	2,313	6,993
Coverage ratio	0.93	0.19	0.00	1.00
Tmax (°C)	21.9	3.2	14.8	35.7
Tmin (°C)	10.4	2.5	1.9	17.7
Precipitation (mm)	1.95	5.93	0.00	80.52
Active districts (count)	2.9	1.1	2.0	6.0

3.5 Methodology and Empirical Strategy

This section describes the empirical framework. Section 3.5.1 outlines the dual-layer design that motivates the two datasets. Sections 3.5.2–3.5.4 present the model specifications, and Section 3.5.5 details the out-of-sample validation strategy.

3.5.1 Dual-Layer Design

The U.S. lettuce market spans six production districts across California and Arizona, but not all districts report FOB prices every week. In particular, the three smaller districts (Santa Maria, Oxnard, and San Joaquin Valley) exhibit substantial gaps in price reporting. A district-level panel including all six districts would have a price coverage rate of only 72%, making it unsuitable for time-series prediction.

I address this by constructing two datasets from the same underlying weekly panel. The first aggregates all six districts into a single market-level time series. The second retains the three core districts (Salinas-Watsonville, Western Arizona, and Imperial Valley) as a panel.

Dataset A: Market-level aggregate. Each week, I compute a shipment-weighted average price across all districts that report a price:

$$P_t = \frac{\sum_i V_{it} \cdot P_{it}}{\sum_i V_{it}}, \quad (3.1)$$

where V_{it} is iceberg lettuce shipment volume from district i in week t , and the sum is taken over districts with nonmissing prices. Weather variables are aggregated in the same way using shipment weights, so that conditions in higher-volume districts receive greater weight. This yields 849 weekly observations, with 98.2% of weeks having at least one observed price. The coverage ratio,

$$C_t = \frac{\sum_{i \in \text{priced}} V_{it}}{\sum_i V_{it}}, \quad (3.2)$$

measures the share of total shipments for which a price is observed and serves as a data quality indicator.

Dataset B: District-level panel. I retain the three districts with both substantial shipment volume and relatively consistent price reporting: Salinas-Watsonville, Western Arizona, and Imperial Valley. This panel contains 1,466 observations with 78% price coverage, higher than the 72% in the full six-district panel.

The two datasets serve distinct roles. Dataset A is used for prediction, where complete coverage and a continuous time series are essential. Dataset B is used for identification, where cross-district variation in weather provides the identifying variation.

3.5.2 Market-Level Prediction Model

The market-level price model takes the form

$$P_t = \alpha + \beta'_1 \mathbf{W}_t + \beta'_2 \mathbf{L}_t + \beta'_3 \mathbf{S}_t + \beta'_4 \mathbf{X}_t + \varepsilon_t, \quad (3.3)$$

where P_t is the shipment-weighted FOB price in week t , $\mathbf{W}_t$ is a vector of contemporaneous weather variables, $\mathbf{L}_t$ contains lagged price and volume terms, $\mathbf{S}_t$ includes seasonal indicators, and $\mathbf{X}_t$ contains controls such as the coverage ratio and the number of active districts.

I estimate Equation (3.3) by OLS with heteroskedasticity-robust standard errors (White, 1980). I also estimate the same specification using XGBoost (Chen and Guestrin, 2016) and Random Forest (Breiman, 2001) as nonlinear benchmarks.

For the tree-based models, I use conservative hyperparameters to limit overfitting. The XGBoost model uses 200 trees, maximum depth 3, learning rate 0.03, subsample ratio 0.7, and L1/L2 regularization ($\alpha = 1.0$, $\lambda = 5.0$). The Random Forest uses 500 trees, maximum depth 6, and minimum leaf size 10. These settings are intentionally restrictive, as the goal is to evaluate out-of-sample performance rather than maximize in-sample fit.

3.5.3 Change Model

Lettuce prices are highly persistent: the first-order autocorrelation exceeds 0.9, and in level regressions the lagged price explains most of the variation. This persistence can mask the effects of other variables, including weather. To isolate the role of short-run shocks, I estimate models in first differences:

$$\Delta P_t = P_t - P_{t-1} = \alpha + \boldsymbol{\gamma}'\mathbf{W}_t + \boldsymbol{\delta}'\mathbf{M}_t + u_t, \quad (3.4)$$

where $\mathbf{W}_t$ contains weather variables and $\mathbf{M}_t$ includes volume, coverage, and seasonal controls. Lagged prices are excluded because differencing removes the autoregressive component.

Comparing the level and change specifications allows me to assess whether weather effects are attenuated by price persistence.

3.5.4 District-Level Panel and Identification

The district panel exploits cross-sectional variation in weather exposure that is averaged out at the market level. A freezing week in Yuma, for example, affects local prices directly but is diluted when aggregated with regions that do not experience the same shock.

Identification relies on the assumption that short-run weather variation affects supply but not demand. Weather shocks influence prices primarily through crop damage and harvest disruptions rather than contemporaneous shifts in consumer demand. This assumption is standard in the agricultural economics literature and is particularly plausible at the weekly frequency considered here, where consumption patterns are relatively stable and unlikely to respond immediately to

short-run weather fluctuations.

The district panel specification is

$$\Delta P_{it} = \phi' \mathbf{W}_{it} + \alpha_i + \mu_q + \eta_{iq} + \nu_{it}, \quad (3.5)$$

where ΔP_{it} is the week-over-week price change for district i in week t , $\mathbf{W}_{it}$ is the local weather vector, α_i are district fixed effects, μ_q are quarter fixed effects, and η_{iq} captures district-by-season interactions. Standard errors are clustered at the district level. Identification comes from cross-district variation in weather within the same week.

I also estimate a level specification that includes lagged prices:

$$P_{it} = \phi' \mathbf{W}_{it} + \psi' \mathbf{L}_{it} + \alpha_i + \mu_q + \eta_{iq} + \nu_{it}. \quad (3.6)$$

Comparing the level and change specifications at the district level provides an additional check on whether persistence attenuates estimated weather effects.

3.5.5 Validation Strategy

Out-of-sample performance is evaluated using an expanding-window approach. For each test year τ from 2014 to 2026, models are trained on all prior years and evaluated on year τ , mimicking a real-time forecasting setting.

I require a minimum of four years of training data, so the first test year is 2014 (trained on 2010–2013). The expanding window is preferred over a rolling window because the sample is relatively short and additional training data improves stability; I verify this in Section 3.7.

Model performance is measured using root mean squared error (RMSE) in price levels. For change models, predictions are converted back to price levels before evaluation. The naive baseline predicts $\hat{P}_t = P_{t-1}$, corresponding to zero predicted price change.

3.6 Results

3.6.1 Market-Level Price Prediction

Table 3.2 reports average out-of-sample prediction errors for the market-level price models over the 2014–2026 test period. The naive baseline, which predicts $\hat{P}_t = P_{t-1}$, has an average RMSE of \$5.11. The OLS level model with the full feature set achieves an RMSE of \$4.36, a 14.7% improvement over the naive baseline and the best performance among the models considered.

Both tree-based methods perform worse than OLS in level models. XGBoost averages \$6.67 and Random Forest \$6.81, each substantially worse than the naive baseline. This pattern is not driven by hyperparameter choice: the conservative settings I use are designed to limit overfitting, and the models capture meaningful in-sample variation. Rather, the underperformance reflects two features of this setting. First, lettuce prices are driven primarily by a strong autoregressive component that is well captured by linear models, leaving limited scope for nonlinear methods to add value. Second, the sample of 849 weekly observations provides relatively limited training data for flexible nonparametric methods, which typically require larger samples to learn stable patterns.

In contrast, the change models yield similar predictive performance across all methods. OLS, XGBoost, and Random Forest produce RMSE values within a narrow range of \$5.00 to \$5.12, with no method clearly outperforming the others. This convergence reflects the removal of the dominant autoregressive component present in price levels. Once price persistence is differenced out, the remaining variation is weaker and harder to predict.

The year-by-year results reveal the source of the gap most clearly in 2022. That year, an INSV outbreak in the Salinas Valley disrupted production severely, and the weekly FOB price for iceberg lettuce reached \$94.33, roughly five times the sample mean. OLS tracked the spike with an RMSE of \$7.77 for that year, while XGBoost and Random Forest produced much larger errors, with RMSEs of \$24.37 and \$23.83, respectively. Because the training data contain no comparable price spike, the tree-based models substantially underpredict the 2022 peak. This illustrates a limitation of tree-based methods in small-sample time series settings: they are generally poor at extrapolating

beyond the range of the training data, while linear autoregressive models can extend trends through the lagged price coefficient.

Table 3.2: Out-of-sample performance for price models

Model	Avg RMSE (\$)	Avg MAE (\$)	Improvement vs. Naive
Naive ($P_t = P_{t-1}$)	5.11	3.35	baseline
<i>Level models</i>			
OLS (full)	4.36	3.14	+14.7%
XGBoost	6.67	4.26	−30.5%
Random Forest	6.81	4.32	−33.3%
<i>Change models</i>			
OLS ΔP	5.11	3.60	+0.1%
XGBoost ΔP	5.12	3.56	−0.2%
Random Forest ΔP	5.00	3.40	+2.2%
OLS log ΔP	5.00	3.41	+2.2%

Figure 3.2 plots actual and predicted market-level prices from the OLS model. The model tracks normal years closely, while most large errors occur during the 2022 price spike driven by the INSV outbreak. Panel (b) shows that these errors are predominantly positive, indicating systematic underprediction during the spike.

3.6.2 Price Persistence and Hidden Weather Effects

In the full level regression, the lagged price coefficient is 1.369 ($p < 0.001$), and the in-sample R^2 is 0.891. The one-week lagged price alone accounts for 30% of total feature importance in the XGBoost level model, and the three price lags together account for 45%, indicating that price persistence dominates the predictive structure. Precipitation is statistically significant ($\hat{\beta} = 0.123$, $p = 0.010$), but its economic magnitude is small relative to the lagged price term.

The change models tell a different story. When I difference out price persistence by modeling ΔP_t instead of P_t , two things happen. First, all methods converge: OLS, XGBoost, Random Forest, and the log-differenced OLS all produce average RMSEs within a narrow range of \$5.00 to \$5.12. The gap between OLS and tree-based methods disappears because the dominant linear feature, lagged




Figure 3.2: Market-level prices: actual and OLS predictions.

price, has been removed. The residual relationship between weather and price changes appears to be relatively simple, and nonlinear methods do not deliver additional predictive gains in this setting.

Second, the relative importance of weather variables increases substantially. Figure 3.3 illustrates this shift using XGBoost feature importance, which provides a consistent measure of the relative contribution of each variable to predictive performance. In the level model, temperature ($tmax$) contributes only 2.8% of total importance. In the change model, where price lags are excluded by construction, temperature importance rises to 6.0%, and other weather variables (precipitation, diurnal range) increase as well. This pattern indicates that the weather signal is present in the data but is largely obscured in level specifications by the dominant role of price persistence.




Figure 3.3: Feature importance in price models.

3.6.3 Mechanism: Weather, Supply, and Price

This subsection provides evidence on the underlying economic mechanisms using the district-level panel specification described in Section 3.5.4. The prediction results above show that weather affects lettuce prices, but that these effects are masked by price persistence in level models. I use this

panel to identify how weather affects prices and to explain why market-level aggregates understate these effects.

Table 3.3 reports the district-level estimates. Freezing degree days are the strongest and most robust weather predictor: each additional degree day below 0°C is associated with a price increase of approximately \$0.85 per carton ($p < 0.001$), and this estimate is stable across specifications (columns M1–M3). Minimum temperature has a significant negative coefficient ($-\$0.14$ to $-\$0.17$ per degree Celsius, $p < 0.01$), consistent with the degree-day measure.

Heat degree days are significant in the level specification ($\$1.21/\text{degree day}$, $p = 0.015$) but not in the change models, suggesting that heat effects may operate over a longer horizon than week-to-week changes capture. Precipitation is positive but not statistically significant at the district level, possibly because its effects operate primarily through volume adjustments.

Table 3.3: District-level effects of weather on prices

Variable	ΔP_{it} (M1)	$\Delta P_{it} + \text{vol}$ (M2)	+ weather lags (M3)	Level (M4)
dd_freeze	0.850***	0.838***	0.817***	0.872***
tmin_avg	-0.163**	-0.168***	-0.142***	-0.115***
dd_heat	0.334	0.309	0.472	1.209**
ppt_total	0.103	0.103	0.102	0.088
heavy_rain	-0.859	-0.863	-0.768	-2.094

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. All models include district FE, quarter FE, and district $\times$ season interactions. SE clustered by district.

These effects reflect short-run supply disruptions. Freezes damage crops and delay harvests, reducing available supply and raising local prices. Because identification comes from within-week cross-district variation, these estimates isolate localized supply shocks rather than broader demand shifts.

At the market level, these effects are attenuated for two reasons. First, spatial aggregation dilutes local shocks. A freeze in Western Arizona has a direct and immediate effect on the local price, but when averaged with Salinas (which may not experience a freeze in the same week), the signal is diluted. This helps explain why freezing degree days are strongly identified in the district panel

but appear weaker in market-level regressions. Second, price persistence obscures contemporaneous weather effects. Freezing degree days are large and significant in all four specifications, but the coefficient is somewhat attenuated in the level model (column M4) as lagged prices absorb some of the variation.

Comparing the shipment and price responses also provides a rough indication of short-run demand conditions. Weather shocks such as freezes generate moderate reductions in shipments but disproportionately large increases in prices, consistent with relatively inelastic short-run demand in fresh lettuce markets and limited short-run substitution possibilities.

Taken together, the evidence supports the following mechanism. Weather shocks, particularly freezes, reduce local supply and raise prices in affected districts. These effects are diluted through spatial aggregation and further obscured by the persistence of prices over time. The dual-layer design is not merely a convenience for handling missing data; it is essential for identifying localized weather effects that aggregation would otherwise obscure.

3.6.4 Volume Prediction

Table 3.4 reports results for market-level shipment volume prediction. Consistent with the price results, OLS substantially outperforms XGBoost. The OLS level model achieves an average RMSE of 288, a 37.6% improvement over the naive baseline of 463, while XGBoost achieves 354, improving on the baseline but remaining well behind OLS.

Both precipitation and freezing conditions are significant predictors of volume. In the OLS level model, the coefficient on weekly precipitation is -8.7 (in 10,000 lb units) per millimeter ($p = 0.019$), reflecting disruptions to field operations and harvest logistics. Freezing degree days have a coefficient of -110 per degree day ($p < 0.001$), indicating that more severe freezes cause larger reductions in shipments. The degree-day specification captures freeze severity more precisely than a binary indicator, and the effect is highly significant. These patterns are consistent with the mechanism identified above, where weather affects prices indirectly through its impact on shipment volumes.

Table 3.4: Out-of-sample performance for shipment volume models

Model	Avg RMSE (10,000 lbs)	Avg MAE (10,000 lbs)	Improvement vs. Naive
Naive ($V_t = V_{t-1}$)	463	347	baseline
OLS Level	288	225	+37.6%
XGBoost Level	354	265	+23.6%
OLS ΔV	459	345	+0.8%
XGBoost ΔV	453	345	+2.1%

3.7 Robustness

This section examines the sensitivity of the main results along several dimensions: commodity choice, validation design, data quality, extreme events, coefficient stability, and missing data treatment.

3.7.1 Romaine Lettuce

I replicate the core analysis using Romaine lettuce prices instead of Iceberg. Table 3.5 compares the two varieties. The market-level OLS model achieves an out-of-sample RMSE of \$2.69 for Romaine, compared to \$4.37 for Iceberg. The lower RMSE reflects lower price volatility rather than improved model fit; the in-sample R^2 is actually higher for Romaine (0.952 vs. 0.892).

Table 3.5: Robustness: iceberg vs. romaine lettuce

	Iceberg	Romaine
<i>Model performance</i>		
In-sample R^2 (level)	0.892	0.952
In-sample RMSE	\$3.95	\$2.47
OOS avg RMSE	\$4.37	\$2.69
Price mean	\$16.27	\$17.80
Price coverage	98.2%	98.4%
<i>Key coefficients</i>		
price_lag1 (level)	1.373***	1.637***
ppt_total (ΔP)	0.190***	0.129***
dd_freeze (ΔP)	1.202***	0.807**
heavy_rain (ΔP)	-1.982**	-1.530***

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

The key finding is that the weather coefficients have consistent signs across the two varieties. Precipitation is positive and significant for both ($p < 0.01$), freezing degree days are positive and significant for both, and heavy rain is negative and significant in both. The core findings of this paper, that weather affects prices through supply disruptions and that these effects are masked by price persistence, are not specific to Iceberg lettuce. They reflect the structure of the lettuce market more broadly.

3.7.2 Expanding vs. Rolling Window

Table 3.6 compares expanding-window and five-year rolling-window validation. The expanding window produces a lower average RMSE (\$4.36 vs. \$4.57). The two approaches yield nearly identical results in early test years, when they share the same training set. The gap widens in later years: in 2021, the expanding window achieves \$3.41 vs. \$4.61 for rolling, and in 2022 the gap is \$7.77 vs. \$9.22. The expanding window benefits from a larger training sample, and the data-generating process appears sufficiently stable that older observations remain informative.

Table 3.6: Robustness to expanding and rolling window validation

Year	Expanding		Rolling (5yr)	
	RMSE	MAE	RMSE	MAE
2014	2.47	1.86	2.47	1.86
2015	3.60	2.83	3.60	2.83
2016	2.66	1.84	2.71	1.85
2017	4.92	3.40	5.00	3.48
2018	5.12	2.94	5.00	2.88
2019	4.61	3.71	4.69	3.78
2020	4.17	2.95	4.27	2.97
2021	3.41	2.71	4.61	3.92
2022	7.77	4.93	9.22	5.81
2023	4.62	3.16	4.67	3.15
2024	4.56	3.53	4.71	3.54
2025	4.16	2.97	4.12	2.88
2026	4.59	4.04	4.28	3.77
Average	4.36	3.14	4.57	3.29

3.7.3 Coverage and Prediction Error

Table 3.7 shows that weeks with low price coverage are harder to predict. When the coverage ratio falls below 0.5, the average MAE is \$4.98, nearly double the \$2.83 for weeks with coverage above 0.95. This pattern is consistent with measurement error in the shipment-weighted price: when fewer districts report prices, the weighted average is noisier, and the model has less reliable training signal. Including the coverage ratio as a feature helps the model partially adjust for this, but it cannot fully compensate for missing price information.

The low-coverage bin (0.5–0.7) contains only one observation and should be interpreted with caution. Figure 3.4 shows the underlying scatter, confirming the generally declining relationship between coverage and prediction error.

One possible concern is that low-coverage periods may coincide with seasonal transition windows (approximately weeks 15–20 and 45–50), when production shifts between California and Arizona and prices tend to be more volatile. Coverage is somewhat lower during these periods (0.927 vs. 0.935), and the standard deviation of district-level prices is substantially higher than during non-transition weeks (\$18.42 vs. \$9.07). However, the relationship between low coverage and larger prediction errors remains present outside transition weeks as well, suggesting that incomplete price reporting contributes additional noise beyond the volatility associated with seasonal transitions.

Table 3.7: Prediction error by coverage ratio

Coverage bin	N	MAE (\$)	Median error (\$)	MAPE (%)
Very low (< 0.5)	17	4.98	4.43	29.6
Low (0.5–0.7)	1	3.97	3.97	16.8
Medium (0.7–0.85)	69	3.55	2.56	18.2
High (0.85–0.95)	69	3.99	2.87	19.5
Very high (> 0.95)	480	2.83	1.84	18.1




Figure 3.4: Price coverage ratio and absolute prediction error.

3.7.4 The 2022 Extreme Year

The 2022 crop year was an outlier by any measure. Table 3.8 summarizes the key statistics. Weekly iceberg lettuce prices averaged \$31.89 with a standard deviation of \$27.09, compared to typical years where the mean is around \$15 and the standard deviation is around \$5. The maximum weekly price reached \$94.33, driven by an INSV outbreak in the Salinas Valley that severely disrupted production.

The disparity between OLS and tree-based methods in 2022 highlights an extrapolation issue: tree-based models are generally unable to extrapolate beyond the range of their training data, while OLS can follow the upward movement through the autoregressive component. I do not interpret 2022 as a failure of the model. Extreme phytosanitary shocks are inherently difficult to forecast using weather and market data alone. The relevant finding is that OLS performs better under extreme conditions, rather than producing highly unstable predictions.

Table 3.8: 2022 extreme year: model performance comparison

	2022	Other years
<i>Price statistics</i>		
Mean price	\$31.89	~\$15
Max price	\$94.33	~\$50
Price std	\$27.09	~\$5
<i>Model RMSE</i>		
OLS Level	\$7.77	\$3.95
XGBoost Level	\$24.37	\$4.75
Random Forest Level	\$23.83	\$4.89
<i>Prediction error (OLS)</i>		
MAE	\$4.93	\$2.93

3.7.5 Feature Stability

I split the sample into three sub-periods and re-estimate the market-level OLS model on each. Table 3.9 reports the key coefficients. The lagged price coefficient ranges from 1.153 to 1.519 across periods, and the precipitation coefficient from 0.105 to 0.183. There is no clear evidence of a structural break in the core relationships.

The temperature variables ($tmax$, $tmin$) show sign instability, which likely reflects collinearity between them rather than instability in the underlying relationship. The degree-day measures vary across periods, consistent with their low frequency making them sensitive to individual events within each sub-sample.

Table 3.9: Coefficient stability across sub-periods (market-level OLS)

Variable	2010–2015	2016–2020	2021–2026
price_lag1	1.245	1.153	1.519
ppt_total	0.105	0.118	0.183
dd_freeze	1.613	−11.493	−50.089
heavy_rain	−2.056	−0.670	−2.220
dd_heat	0.000	0.000	−26.138

3.7.6 Missing Price Prediction

Not all district-week observations have reported prices. Out of 2,441 district-week observations with positive shipment volume, 682 (28%) are missing a price. I train a model on observed prices using weather, volume, district indicators, and seasonal controls, then predict prices for the missing observations. OLS achieves an average out-of-sample RMSE of \$11.08 (expanding-window validation), compared to \$13.27 for XGBoost.

Table 3.10 compares predicted and observed mean prices by district. The predicted prices for missing observations are systematically higher than the observed means. For example, the predicted mean for Oxnard’s missing weeks is \$17.72 vs. an observed mean of \$15.30. This pattern is consistent with non-random missingness in price reporting: districts may be less likely to report prices during periods of high prices or supply stress.

Table 3.10: Missing price prediction by district

District	Missing obs	Predicted mean (\$)	Observed mean (\$)
Oxnard	238	17.72	15.30
Santa Maria	199	19.57	13.86
Western Arizona	76	18.93	17.26
Imperial Valley	71	20.93	15.49
Salinas-Watsonville	49	23.72	16.14
San Joaquin Valley	49	22.18	20.40
Total	682		

OLS RMSE = \$11.08, XGBoost RMSE = \$13.27 (expanding window).

3.8 Conclusion

This paper studies weather-driven supply responses in a setting where high-frequency yield data are not available. In contrast to major row crops such as corn and soybeans, lettuce production is geographically concentrated and occurs as a continuous flow throughout the season rather than as a single annual harvest. As a result, standard approaches that rely on county-level yield data are not feasible. Instead, I use weekly shipment volumes and FOB prices as downstream market

outcomes to infer supply responses. Identification relies on the assumption that short-run weather variation affects supply but not demand. Under this assumption, variation in prices and shipments can be used to recover the impact of weather on supply.

I document three main results. First, simple linear models outperform tree-based machine learning methods for price prediction. OLS achieves an out-of-sample RMSE of \$4.36, compared to \$6.67 for XGBoost and \$6.81 for Random Forest. Lettuce prices exhibit strong persistence, and most of the predictable variation is captured by lagged prices. Conditional on this persistence, the remaining relationship between weather and prices appears to be relatively simple, as more flexible methods do not improve predictive performance and perform poorly during extreme events.

Second, weather effects on prices are real but masked by price persistence. In level models, the three price lags together account for nearly half of total feature importance, leaving limited scope for weather variables to contribute. When I model price changes instead of price levels, weather variables become substantially more important. At the district level, each additional freezing degree day increases weekly prices by \$0.85 ($p < 0.001$), but this effect is attenuated when aggregated to the market level. The dual-layer approach, combining market-level aggregation for prediction with district-level panels for identification, is therefore necessary to uncover localized weather effects.

Third, these findings are robust. The results replicate for Romaine lettuce, coefficients are stable across sub-periods, and alternative validation strategies yield similar conclusions. Data quality, as measured by price coverage, is also an important determinant of prediction accuracy.

Several limitations remain. The identification assumption that weather affects supply but not demand is plausible at the weekly frequency considered here, but longer-run climate trends may affect both sides of the market. My estimates are based on short-run weather variation and therefore do not capture long-run climate adaptation. They provide a useful benchmark for understanding climate-related supply risks, but extrapolating from weather responses to climate responses remains an open question. The 2022 virus outbreak highlights that non-weather shocks can dominate market outcomes, and such events are inherently difficult to predict using historical price and weather data. In addition, the analysis focuses on six production districts in California and Arizona; extending the

framework to other crops or regions would require additional data and may yield different dynamics.

Despite these limitations, the results show that downstream market outcomes, shipment volumes and FOB prices, can be used to recover weather-driven supply responses even in the absence of high-frequency yield data. More broadly, the paper illustrates how market data can serve as a proxy for production conditions in perishable specialty crops. For participants in fresh produce markets, publicly available USDA data combined with linear models provide a transparent and practical approach to monitoring weather-related price risk.

References

- Breiman, Leo.** 2001. “Random Forests.” *Machine Learning*, 45(1): 5–32.
- Chen, Tianqi, and Carlos Guestrin.** 2016. “XGBoost: A Scalable Tree Boosting System.” 785–794.
- Crane-Droesch, Andrew.** 2018. “Machine Learning Methods for Crop Yield Prediction and Climate Change Impact Assessment in Agriculture.” *Environmental Research Letters*, 13(11): 114003.
- Daly, Christopher, Michael Halbleib, Joseph I. Smith, Wayne P. Gibson, Matthew K. Doggett, George H. Taylor, Jan Curtis, and Phillip P. Pasteris.** 2008. “Physiographically Sensitive Mapping of Climatological Temperature and Precipitation across the Conterminous United States.” *International Journal of Climatology*, 28(15): 2031–2064.
- Kistner, E., O. Kellner, J. Andresen, D. Todey, and L. W. Morton.** 2018. “Vulnerability of Specialty Crops to Short-Term Climatic Variability and Adaptation Strategies in the Midwestern USA.” *Climatic Change*, 146(1): 145–158.
- Letta, Marco, Pierluigi Montalbano, and Guillaume Pierre.** 2022. “Weather Shocks, Traders’ Expectations, and Food Prices.” *American Journal of Agricultural Economics*, 104(3): 1100–1119.
- Lobell, David B., Kimberly Nicholas Cahill, and Christopher B. Field.** 2007. “Historical Effects of Temperature and Precipitation on California Crop Yields.” *Climatic Change*, 81(2): 187–203.

- Lobell, David B., Wolfram Schlenker, and Justin Costa-Roberts.** 2011. "Climate Trends and Global Crop Production since 1980." *Science*, 333(6042): 616–620.
- Mullainathan, Sendhil, and Jann Spiess.** 2017. "Machine Learning: An Applied Econometric Approach." *Journal of Economic Perspectives*, 31(2): 87–106.
- Ortiz-Bobea, Ariel, Toby R. Ault, Carlos M. Carrillo, Robert G. Chambers, and David B. Lobell.** 2021. "Anthropogenic Climate Change Has Slowed Global Agricultural Productivity Growth." *Nature Climate Change*, 11: 306–312.
- Roberts, Michael J., and Wolfram Schlenker.** 2009. "World Supply and Demand of Food Commodity Calories." *American Journal of Agricultural Economics*, 91(5): 1235–1242.
- Roberts, Michael J., and Wolfram Schlenker.** 2013. "Identifying Supply and Demand Elasticities of Agricultural Commodities: Implications for the US Ethanol Mandate." *American Economic Review*, 103(6): 2265–2295.
- Roberts, Michael J., Wolfram Schlenker, and Jonathan Eyer.** 2013. "Agronomic Weather Measures in Econometric Models of Crop Yield with Implications for Climate Change." *American Journal of Agricultural Economics*, 95(2): 236–243.
- Roll, Richard.** 1984. "Orange Juice and Weather." *American Economic Review*, 74(5): 861–880.
- Schlenker, Wolfram, and Michael J. Roberts.** 2009. "Nonlinear Temperature Effects Indicate Severe Damages to U.S. Crop Yields under Climate Change." *Proceedings of the National Academy of Sciences*, 106(37): 15594–15598.
- Storm, Hugo, Kathy Baylis, and Thomas Heckelei.** 2020. "Machine Learning in Agricultural and Applied Economics." *European Review of Agricultural Economics*, 47(3): 849–892.
- USDA Agricultural Marketing Service.** 2026a. "Market News: Fruit and Vegetable Shipping Point Prices." <https://marketnews.usda.gov/mnp/fv-home>, Accessed April 2026.

USDA Agricultural Marketing Service. 2026b. “Truck Rate Reports: Fruit and Vegetable Shipments.” <https://usda.library.cornell.edu/concern/publications/mg74qm08h>, Accessed April 2026.

USDA National Agricultural Statistics Service. 2023. “Vegetables Summary.” https://www.nass.usda.gov/Surveys/Guide_to_NASS_Surveys/Vegetables/, Accessed April 2026.

White, Halbert. 1980. “A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity.” *Econometrica*, 48(4): 817–838.

3.9 Appendix

This appendix provides supplementary materials. Appendix 3.9.1 defines all variables used in the analysis. Appendix 3.9.2 reports the complete OLS coefficient estimates for the market-level price model. Appendix 3.9.3 presents out-of-sample RMSE by year for the six main models.

3.9.1 Variable Definitions

Table A1: Variable definitions and data sources

Variable	Definition	Unit	Source
<i>Outcome variables</i>			
price	Shipment-weighted FOB price	\$/carton	USDA Market News
dprice	$P_t - P_{t-1}$	\$/carton	Constructed
volume	Total weekly truck shipments	10,000 lbs	USDA AMS
<i>Weather variables</i>			
tmax_avg	Weekly mean of daily max temperature	°C	PRISM
tmin_avg	Weekly mean of daily min temperature	°C	PRISM
ppt_total	Total weekly precipitation	mm	PRISM
diurnal_range	tmax_avg - tmin_avg	°C	Constructed
dd_freeze	$\sum_d \max(0 - T_d^{\min}, 0)$	°C-days	Constructed
dd_heat	$\sum_d \max(T_d^{\max} - 40, 0)$	°C-days	Constructed
heavy_rain	$\mathbf{1}[\text{weekly precip} > 25 \text{ mm}]$	Binary	Constructed
<i>Lagged variables</i>			
price_lag k	$P_{t-k}, k \in \{1, 2, 4\}$	\$/carton	Constructed
volume_lag k	$V_{t-k}, k \in \{1, 2\}$	10,000 lbs	Constructed
tmax_lag k	tmax_avg $_{t-k}, k \in \{1, 2\}$	°C	Constructed
ppt_lag1	ppt_total $_{t-1}$	mm	Constructed
<i>Rolling averages</i>			
volume_roll4_mean	4-week rolling mean of volume	10,000 lbs	Constructed
tmax_avg_roll4_mean	4-week rolling mean of tmax	°C	Constructed
ppt_total_roll4_mean	4-week rolling mean of precip	mm	Constructed
<i>Market and seasonal</i>			
coverage	Share of shipments with observed price	[0, 1]	Constructed
n_districts	Number of active shipping districts	Count	Constructed
month	Calendar month	1–12	Constructed
week_of_year	Calendar week	1–52	Constructed

3.9.2 Full OLS Regression Output

Table A2: Full OLS regression: market-level price model

Variable	Coeff.	Robust SE	<i>z</i>	<i>p</i> -value
<i>Price and volume lags</i>				
price_lag1	1.369	0.042	32.31	<0.001
price_lag2	−0.575	0.049	−11.77	<0.001
price_lag4	0.070	0.024	2.91	0.004
volume_lag1	−0.001	0.001	−1.32	0.188
volume_lag2	0.002	0.000	3.74	<0.001
<i>Contemporaneous weather</i>				
tmax_avg	24.829	27.888	0.89	0.373
tmin_avg	−25.010	27.899	−0.90	0.370
ppt_total	0.121	0.048	2.54	0.011
diurnal_range	−24.982	27.899	−0.90	0.371
dd_heat	−31.832	22.796	−1.40	0.163
dd_freeze	1.715	0.245	7.00	<0.001
heavy_rain	−1.479	0.688	−2.15	0.032
<i>Lagged weather</i>				
tmax_lag1	−0.278	0.097	−2.87	0.004
tmax_lag2	−0.112	0.111	−1.01	0.312
ppt_lag1	−0.018	0.032	−0.55	0.579
<i>Rolling averages</i>				
volume_roll4_mean	−0.002	0.001	−2.56	0.010
tmax_avg_roll4_mean	0.563	0.269	2.09	0.036
ppt_total_roll4_mean	−0.103	0.099	−1.04	0.296
<i>Market and seasonal</i>				
coverage	−1.473	1.400	−1.05	0.293
n_districts	0.245	0.187	1.32	0.188
month	0.121	0.135	0.90	0.369
week_of_year	−0.012	0.032	−0.38	0.706
constant	7.007	2.032	3.45	<0.001
RMSE = \$3.95. Dependent variable: shipment-weighted FOB price (\$/carton).				

3.9.3 Out-of-Sample RMSE by Year

Table A3: Out-of-sample RMSE by year (\$, expanding window)

Year	Level models (full)			Change models (ΔP)		
	OLS	XGBoost	RF	OLS	XGBoost	RF
2014	2.41	2.74	2.78	2.75	2.80	2.71
2015	3.60	5.63	5.41	3.94	3.97	3.94
2016	2.66	2.79	2.92	3.20	3.32	3.12
2017	4.94	6.18	6.62	5.98	6.06	5.99
2018	5.09	6.20	6.52	6.27	6.16	6.21
2019	4.59	4.83	4.87	5.18	4.89	4.97
2020	4.17	4.87	4.93	4.78	4.43	4.49
2021	3.38	4.15	4.81	3.68	3.33	3.37
2022	7.72	24.37	23.83	6.95	7.25	7.30
2023	4.59	6.86	7.29	5.90	6.25	5.90
2024	4.53	6.57	6.43	5.69	5.89	5.67
2025	4.16	5.22	5.92	4.82	5.33	5.07
2026	4.87	6.35	6.23	7.28	6.84	6.23
Average	4.36	6.67	6.81	5.11	5.12	5.00