Title

The Economic Effects of Weather, Climate Change, and Resource Scarcity on Agricultural Production and Decision Making

Author

Smith, Sarah Caroline

Publication Date

2023

The Economic Effects of Weather, Climate Change, and Resource Scarcity on Agricultural Production and Decision Making

By

SARAH SMITH
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:

Timothy K.M. Beatty, Chair

Julian M. Alston

Rachael E. Goodhue

Pierre Mérel

Committee in Charge

2023

Abstract

In this dissertation, I analyze the economic effects of weather, climate change, adaptation, and natural resource scarcity through the lens of agriculture in California.

Quality is central to much of agriculture because of its role in price determination and contractual arrangements. Nevertheless, prior work on the effect of weather and climate change on agricultural production mostly focuses on staple crop yields. Ignoring quality will likely bias estimates of the impact of weather and climate change on productivity and farm income. In the first essay co-authored with Timothy Beatty, we quantify the effect of temperature exposure and climate change on the revenue of specialty crop growers through two pathways: quality and yield. In contrast to earlier work on irrigated crops, we find extreme temperatures negatively affect both yield and quality leading to reduced grower revenue in a setting where irrigation is the norm. While the yield effect dominates, failing to account for quality significantly underestimates the true effect of temperature exposure on revenue by up to 20%. We predict climate change will significantly reduce yield, quality, and revenue by century's end absent additional adaptation.

In the second essay, I investigate how natural resource scarcity affects farmers' decision making. Using a panel of 3,300 irrigated fields in California, I ask whether farmers respond to water scarcity by changing whether they plant, what they plant and how they produce. Novel data provide evidence of the mechanisms by which growers use water more efficiently – insights unattainable in publicly available or survey data. I find that during a water scarce year, growers are more likely to plant earlier, plant fast-maturing varieties, and preferentially

plant fields equipped with drip irrigation rather than less water-efficient technologies. Water access drives how growers conserve water: growers with low-priority access rely on more costly margins, such as fallowing or changing growing practices in ways that reduce revenue per acre. This demonstrates that agricultural producers engage in water-saving practices more than what has been previously found and that these practices help growers avoid fallowing, a much more costly response.

The first and second essays use proprietary data on California's $1 billion processing tomato industry, which produces nearly one-third of the world's processing tomatoes. Tomatoes are the second-most consumed vegetable in the United States and contain nutrients that are important to human health, but often under consumed. These novel data are collected by a large tomato processor for the purposes of contracting and payment, and include field-level quality, yield, price, and grower practices for many independent farmers operating in California.

The third essay returns to the relationship between quality, weather, and climate for a different specialty crop product: wine. Wine is the most differentiated of all agricultural products, with much of the differentiation based on the combination of wine grape varieties and "terroir": the natural environment in which the grapes are grown. Co-authored with Julian Alston, the objective of this study is analyze the complex relationship between collective reputation, climate, weather, and price premia and quality for varietal wine in California. We find temperatures warmer than the regional norm had negative effects on both Cabernet Sauvignon and Chardonnay wine prices and scores. We also find that wines from premier regions are less influenced by deviations in temperature from regional climate. This supports the notion that producers of higher-value wines intervene more to mitigate the negative effects of weather.

Quality matters for all agricultural producers but some more than others. Prices of processing tomatoes vary plus or minus 20% in my sample whereas wine grape prices vary by a factor of 50, even for the same variety and vintage produced in the same region.

For processing tomatoes, we benefit from detailed and precisely measured data on quality attributes and how they affect grower revenue. Quality variation is arguably more important for wine grapes and wine but also more complex, which makes it more difficult to cleanly link weather effects to attributes that affect prices of grapes and wine. Rather than direct observations of quality attributes, we use expert ratings and recommended retail prices as indicators of wine and wine grape quality. Both papers contribute to our understanding of how weather and climate influence agricultural product quality.

Each of these three papers contributes to our understanding of how economic agents are affected by and respond to weather, climate, and climate change. Most of what we know about agricultural production in the face of climate change comes from research on yields of staple crop. This dissertation contributes by analyzing specialty crops that are understudied despite making up 40% of the total value of U.S. crops. Taken together, the research in this dissertation highlights potential harm from a hotter and drier climate absent adaptation.

Acknowledgments

Many thanks to my committee members Tim Beatty, Julian Alston, Rachael Goodhue, and Pierre Mérel. A special thanks to my advisor Tim for your unwavering encouragement, guidance, and humor. Julian, I feel privileged to have worked with you and I am thankful for our many walks and discussions. Rachael, thank you for your invaluable insights and support. Pierre, I am so appreciative of everything I have learnt from you both in and out of the classroom. Thank you to all the faculty and staff at ARE, the effort you put in is what makes our community so special.

Thank you to all my family and friends in Australia for always cheering me on. I wouldn't have known this path was a option for me without the incredible mentorship of David Ubilava. To my best friends Kirri, Caitlin, Lachlan, and Sam, thank you for bringing so much laughter to my life. I'm so thankful to Mum, Dad, Emma, Ewan, and Lucia for role modelling hard work, perseverance, and love.

To the many friends I have made at Davis, thank you for all the coffee breaks, Arboretum walks, hallway chats, and cheeky beers. A special thanks to Ali, Ashley, Caitlin, Dan, Dasha, Isy, Jeff, Joakim, Laura, Lindsay, Mari, Saloni, Sam, Scott, and TJ. Charlotte Ambrozek and Julian Arteaga deserve a shoutout – I admire you both so much as researchers and as people.

Finally, I am eternally thankful to Tim Whitnall for coaching me through the toughest moments and celebrating with me during the best moments. With you by my side I feel like I could do 20 of these PhD things (let's keep it to one though shall we?)

Contents

List of Figures

List of Tables

Essay 1

Climate Change and Field-Level Crop Quality, Yield, and Revenue

1.1 Introduction

The value of every agricultural product depends on its quality. Grain, meat, and milk are graded according to USDA quality standards, and fresh produce is sorted by size, color, and defects. Products grown under contract often face quality incentives and low quality can violate contractual obligations or make products unmarketable. But work quantifying the economic impacts of extreme weather and climate change on agricultural production focuses almost exclusively on yield (see Carter et al. (2018) or McCarl & Hertel (2018) for summaries). Ignoring quality may bias estimates of the impact of weather and climate change on agricultural productivity and farm income.

We ask three related questions: Has historical weather impacted the revenues of specialty crop producers through both yield and quality? What are the relative magnitudes of the yield and quality effects? To what extent does failing to account for quality bias estimates of weather's impact on farm revenue? We answer these questions using 12,000 field-level observations of processing tomato yield, quality, and grower practices from across California between 2011 and 2020. These data are collected by a large tomato processor for the purposes of contracting and payment and capture the behavior of hundreds of farmers. We use standard gridded weather data from PRISM and the well-established panel specification (Deschênes & Greenstone (2007); Schlenker & Roberts (2009)) to facilitate comparison to earlier work.

We find that extreme weather conditions affect the revenue of growers despite their use of irrigation. Yield responds negatively to exposure to hot temperatures and, to a lesser extent, cool temperatures. An additional 24 hours of exposure above 30°C causes yields to decrease by up to 1.8% relative to 24 hours of average temperatures. Further, quality declines with exposure to hot temperatures, causing growers to receive a lower price. Taken as a whole, we find that, relative to 24 hours of average temperatures, exposure to temperatures in excess of 30°C decreases revenue by up to 2.3%. Exposure to cool temperatures below 10°C causes a statistically significant, but smaller, decrease in revenue. Failing to account for quality effects would bias downward the effect of exposure to heat on revenue by up to 20%.

This paper is among the first to document the effect of weather and climate change on agricultural product quality. Kawasaki & Uchida (2016), Kawasaki (2019), Dalhaus et al. (2020), and Ramsey et al. (2020) all find negative, economically important, effects of weather on grower revenue with quality being a key pathway. Our setting and novel data offer advantages relative to earlier work.

First, quality attributes are precisely measured. California's processing tomato industry has mandatory quality testing by an independent third party, so neither grower nor processor can accidentally or intentionally misstate quality. We observe several individual quality attributes for each field-year observation, which allows us to estimate the impact of weather and climate change on each attribute individually.

Second, our measures of quality have economic significance. Growers are paid a price per ton that depends on observed quality attributes, which introduces variation in price of plus or minus 20%. We cleanly link quality attributes to price using observable contract terms established prior to planting. The contract structure allows us to remove price variation driven by potentially endogenous market conditions and isolate variation in price resulting from variation in quality alone.

Third, in almost all empirical settings, researchers observe the subset of production that exceeds a quality threshold and is selected to be graded or harvested. This selection biases unconditional measures of both quality and yield. As detailed below, selection in California's processing tomato industry is relatively small compared to other settings. We replicate the sample selection problem common in other settings and find that selection bias affects estimates of the effect of weather on quality, yield, and revenue.

Taking a step back, prior work on the effects of weather and climate change on agriculture has mostly focused on yields of staple crops – primarily major grains and oilseeds (Schlenker & Roberts (2009); Lobell et al. (2011); Tack et al. (2015); Chen et al. (2016); Gammans et al. (2017); Shew et al. (2020); Malikov et al. (2020); Schmitt et al. (2022)). In contrast, specialty crops are understudied despite making up 40% of the total value of U.S. crops (USDA NASS, 2017). By focusing on an irrigated specialty crop, we extend a literature that has largely focused on rain-fed staple crops. Irrigated specialty crops have distinct production functions and likely respond differently to weather shocks than rainfed field crops. Prior work finds that irrigated water application essentially eliminates the negative effect of extreme heat and climate change on yields of staple crops (Shaw et al. (2014); Carter et al. (2016); Tack et al. (2017a); Wing et al. (2021)), and agricultural total factor productivity (Ortiz-Bobea et al., 2018). But in a setting where irrigation has long been the rule rather than the exception, we find both yield and quality are affected by exposure to hot temperatures and climate change, leading to lower grower revenue.

Finally, we predict the impact of climate change accounting for climate uncertainty, emissions uncertainty, and regression uncertainty. We find that without additional adaptation, by century's end yield, quality, and revenue will all be reduced by climate change by century's end without additional adaptation. Assuming a middle-of-the-road emissions scenario (SSP2-4.5) and predicting distant future yield and quality based on model parameters applicable to today's technology, by the end of the century we predict a median loss of yield of 13% with a 95% confidence interval of 6% to 31%. We predict losses in quality of 1% to 5% by century's end. Together, we predict declining yield and quality will cause grower revenue to fall.

1.2 The Setting

Tomatoes are the second-most produced fruit or vegetable globally by value (behind potatoes) (FAO, 2023) and the second-most consumed in the United States (USDA ERS, 2020). Tomatoes can be either consumed fresh or processed into paste, ketchup, or a canned product. They contain nutrients like vitamin E, potassium and lycopene that are important to human health but often under consumed (Wu et al., 2022).

Tomatoes destined for processing (henceforth processing tomatoes) are specific varieties, distinct from fresh tomatoes, bred and grown to enhance qualities desirable for processing into paste or for canning. California's $1 billion processing tomatoes industry produces more than 90% of U.S. processing tomato output (California Department of Food and Agriculture, 2019). In California, processing tomatoes are planted between February and June to facilitate continuous harvesting between July and October. They are mostly grown in outdoor fields in the San Joaquin and Sacramento Valleys (California Department of Food and Agriculture, 2019). Processing tomatoes are a warm-season crop – during California's growing season, maximum temperatures average around 30°C and precipitation is scant. Growers irrigate to ensure crops receive enough water and can tolerate high temperatures during the height of summer (Hartz et al., 2008).

Our data are from a large tomato processor that purchases processing tomatoes under contract from growers in the San Joaquin Valley, Sacramento Valley, and Central Coast regions of California. Processors contract with growers because quality of the raw material input used in the processing tomato industry is crucial to produce consistent and high-quality output of processed products. Processors incentivize growers by paying a price that depends on the quality of tomatoes delivered. Contracts are negotiated between individual processors and the California Tomato Growers Association (CTGA) on behalf of all growers. Negotiations establish each processor's seasonal base price, quality adjustments, and bonuses, which processors then offer to growers on a take it or leave it basis.

Table 1.1 summarizes the eight quality attributes we observe and their effect on price. The processor deducts a percentage of the base price for the presence of defects (mold, green tomatoes, worms, material other than tomatoes (MOT) and limited use (LU) tomatoes). The processor employs an incentive program whereby growers receive a bonus (or penalty) if the brix (soluble solids or sugar content) of delivered tomatoes is more (or less) than the average for the same variety in the same county in which they were grown. Quality adjustments are proportional to the quality achieved by growers as measured by the Processing Tomato Advisory Board (PTAB). Finally, the processor values staggered harvesting and delivery to minimize bottlenecks at processing facilities. Producers receive a bonus for delivering tomatoes early or late in the season. Quality incentives are economically important and introduce price variation of plus or minus 20% relative to average prices.

In most empirical settings, researchers observe only a subset of production that exceeds a minimum quality threshold and is selected to be harvested or graded. Our data are unique in that we observe most of the selection and sorting process. First, there is limited opportunity to selectively harvest and grade processing tomatoes. Unlike many specialty crops, processing tomatoes are mechanically harvested[1]. Mechanical harvesters sort in the field and we do not observe what is rejected at harvest. However, the processor implements minimal sorting at harvest. The processor prefers to sort tomatoes at the processing plant because their sorting machines are more accurate than the sorter aboard the harvester. Any selection at harvest is therefore relatively small.

The processor closely manages harvesting logistics because operating processing plants near full capacity is key to profitability. Almost all processing tomato production in California is grown under a contract between a grower and processor (USDA NASS, 2021b).

[1]See Just & Chern (1980) for details on the introduction and widespread adoption of mechanical harvesters in California's processing tomato industry during the 1960s.

Table 1.1: Summary of quality measures

Quality attribute	Explanation	Effect on price
Brix	Brix is a measure of soluble solids or sugar content	Bonus (penalty) if brix is more (less) than the average for the same variety in the same county
Limited use (LU) percent	Tomatoes that are soft, split or squashed and have limited processing use	Deduction from base price in proportion to percentage
Material other than tomatoes (MOT) percent	Mainly dirt and sometimes vines	Deduction from base price in proportion to percentage
Green percent	Tomatoes that are unripe	Deduction from base price in proportion to percentage
Mold percent		Deduction from base price in proportion to percentage
Worm percent		Deduction from base price in proportion to percentage
Color score	Red-ripe fruit are given a color score where low scores indicate a better color	No effect
pH	Higher pH usually indicates the fruit is more ripe	No effect

Contracts are written for specific fields and the processor typically harvests and transports tomatoes to the processing plant. As a consequence, growers have little opportunity to strategically sort prior to sale.

Further, each and every truckload of processing tomatoes in California undergoes mandatory grading, including quality measurement, at an independent state inspection station prior to delivery. Administered by PTAB, the California Processing Tomato Inspection Program was established in 1987 to create and uphold quality standards for California's processing tomatoes. Quality observations are shared with processors and growers and are used for price determination.

Finally, the processor engages in a small degree of sorting in which foreign matter and some very low quality tomatoes are deducted from the quantity used for payment. However, we are able to fully observe this sorting. The processor records total tons prior to sorting as well as paid tons after sorting. In sum, compared with other studies, the institutional setting for our work means the quality and yield observations more accurately reflect conditions in the field and our analysis is consequently less susceptible to selection problems.

The processing tomato agronomic literature (Hartz et al., 2008) finds that maximum temperatures between 25°C and 35°C are ideal for vegetative growth, plant development, and fruit set, so long as plants have sufficient access to water. Hot temperatures without adequate moisture cause tomato plants to become stressed, affecting yield and quality. Temperatures below 10°C slow development and also affect quality. The extent of damage caused by extreme temperature depends on its timing in the phenological cycle. Cool temperatures at the beginning of the growing season are believed to benefit brix, while hot temperatures at the beginning and end of the growing season reduce yield and increase the limited use share (Personal correspondence with the processor, 2020).

Lobell et al. (2007) find that maximum temperatures in April and June explain 58% of yield variability in California's processing tomatoes between 1980 and 2003. Hot temperatures benefit seedling growth during April. However, yields decrease when processing tomato ploants are exposed to maximum temperatures above 32°C in June. Marklein et al. (2020) estimate that 34–87% of land on which tomatoes are grown in California will no longer be suitable by mid-century because summer temperatures will be too hot. This estimate may be an upper bound as they assume hot temperatures translate directly to heat stress. Tomato plants can withstand considerable heat so long as they have access to water (Hartz et al., 2008). Cammarano et al. (2022) project a decrease in global processing tomato production by 2050, driven by temperatures rising above the optimal threshold (28°C) in California and Italian growing regions.

When using proprietary data, there is a tradeoff between internal and external validity. In this case, we obtain a level of detail not available in public data. These data are not from surveys but rather from administrative records of every field contracting with the processor between 2011 and 2020. Our data are at the field-year level and contains a range of information about fields such that we can observe and control for field-specific factors. Detail enhances internal validity but may limit external validity.

One concern is that the hundreds of growers in this proprietary dataset are not representative of the broader processing tomato industry. Several factors argue against this concern. The fields in our sample are geographically dispersed across 18 counties in California and closely match patterns of production locations in California (Figure 1.1). Field-level yields averaged to the county level are 10% higher than county yields reported by NASS. However, the series are highly correlated. Nevertheless, it is always possible that there is selection on unobservables that may affect the type of grower we observe and thus the external validity of our results.

1.3 Data

1.3.1 Field-level dataset

As described above, we use data on all tomatoes grown and sold under contract to a large tomato processor in California, between 2011 and 2020. Observations are at the field-year

(a) Location of fields contracted with processor in 2020

(b) Production in 2020, by county (USDA NASS, 2021a)

Figure 1.1: Map of processing tomato production in California

level ($n = 11,926$) and include information on field acreage, variety, total tons, paid tons, yield, quality attributes, and the latitude and longitude of the field centroid. Observations of quality attributes and tonnage are from PTAB mandatory testing that occurs prior to delivery. The field-specific data also include information about the growing practices, including planting and harvesting dates, irrigation technology, and the crop previously planted on the field. Field-level observations are linked to an unbalanced panel of 438 growers and 247 grower groups. The grower group identifier links growers within the same network or organization[2].

1.3.2 Contract terms

We also observe pricing terms negotiated between the processor and the California Tomato Growers Association for each year. All growers are offered the same contract in a year and

[2] An example of a grower group is four children dividing a family farm. Each child would have a distinct grower id and the four would share a common grower group id.

do not negotiate individual pricing terms with the processor. Every year, a base price is established at the start of the growing season that reflects current and expected market conditions. Contracts also establish bonuses and deductions. As summarized in Table 1.1, there are eight measures of quality of which six are linked to price bonuses or deductions. In addition, growers receive an early (late) season bonus if they deliver tomatoes at the beginning (end) of the growing season.

1.3.3 Outcomes

Yield for field i in year t is defined as total tons divided by acres:

$$\text{yield}_{it} = \frac{\text{total tons}_{it}}{\text{acres}_{it}} \tag{1.1}$$

Next, we calculate price for each field-year observation using observed quality and the schedule of quality bonuses and deductions established at the beginning of the growing season. Processing tomatoes are used to produce storable products and therefore base prices are temporally correlated across years. This could introduce endogeneity into our panel model (see Section 1.4 for more details). Instead, we isolate variation in price driven by quality by applying the estimated adjustments to the 10-year average base price. This removes common price movements driven by potentially endogenous market shocks while preserving common and individual quality shocks. This "quality-adjusted price" is not the real price, rather it is effectively a quality index with weights equal to each quality attribute's effect on price.

$$\text{price}_{it}^{\text{quality adjust}} = \overline{\text{base price}} \times (1 - \text{deducts}_{it}) + \text{bonus}_{it} \tag{1.2}$$

where $\overline{\text{base price}}$ is the average base price over the 10 year sample and deducts_{it} and bonus_{it} are adjustments that depend on observed quality and date of delivery of tomatoes from field i in year t.

Finally, we estimate field-level revenue per acre (henceforth revenue) by multiplying quality-adjusted price by paid tons and dividing by acreage. The processor doesn't count flawed tonnage towards the quantity on which producers are paid. The revenue estimate uses paid tons, which is on average 7% less than total tons used in the yield calculation. The difference between total and paid tons is quantity of tonnage that is not commercially viable and is disposed of by the processor.

$$\text{revenue}_{it} = \frac{\text{paid tons}_{it}}{\text{acres}_{it}} \times \text{price}_{it}^{\text{quality adjust}} \tag{1.3}$$

1.3.4 Weather data

We obtain weather data from PRISM (PRISM Climate Group, Oregon State University, 2020), which publishes daily temperature and precipitation data interpolated to 4km grids for the whole time span. We match weather data to each field-level observation by identifying the PRISM grid in which the field centroid falls. In Section 1.4.1, we explain how we translate daily observations of temperatures into measures of temperature exposure for each field-year observation.

1.3.5 Control variables

We also gather data on several controls. We source information on tomato varieties from AgSeeds (AgSeeds, 2020), which includes key attributes and use categories for each of the 159 varieties in the processor dataset. Finally, we match each field to its major soil type in the National Cooperative Soil Survey (NRCS USDA, 2020).

Table 1.2: Summary statistics

	units	mean	sd	min	max
Area	acres	57.31	43.47	0.3	323.2
Growing days	no.	133.72	9.50	96.0	175.0
Yield	tons/acre	52.49	13.32	6.7	99.9
Quality attributes					
Brix		5.08	0.49	3.5	7.2
LU percent		1.42	1.17	0.0	13.1
MOT percent		1.64	1.23	0.0	13.3
Green percent		3.02	2.33	0.0	24.3
Mold percent		1.66	1.92	0.0	27.5
Worm percent		0.00	0.01	0.0	0.5
Color score		20.98	1.74	13.3	34.4
pH		4.41	0.09	2.9	4.8
Weather					
Average minimum temperature	$^\circ$ C	13.79	1.08	9.8	17.6
Average maximum temperature	$^\circ$ C	31.19	1.62	24.5	35.4
Total precipitation	mm	24.45	25.80	0.0	198.9
Soil type					
Alluvium	prop.	0.96	0.20	0.0	1.0
Eolian	prop.	0.00	0.04	0.0	1.0
Organic material	prop.	0.03	0.17	0.0	1.0
Lacustrine	prop.	0.00	0.03	0.0	1.0
Residuum	prop.	0.01	0.08	0.0	1.0
Irrigation technology					
Drip irrigation	prop.	0.75	0.43	0.0	1.0
Furrow irrigation	prop.	0.13	0.33	0.0	1.0
Missing irrigation tech.	prop.	0.10	0.31	0.0	1.0
Sprinkler irrigation	prop.	0.02	0.14	0.0	1.0
Varietal attributes					
Extended field storage variety	prop.	0.55	0.50	0.0	1.0
Tomato spotted wilt resistant	prop.	0.45	0.50	0.0	1.0
Fusarium Wilt resistant	prop.	0.15	0.36	0.0	1.0
Powdery Mildew resistant	prop.	0.04	0.19	0.0	1.0
Fusarium Crown Rot resistant	prop.	0.00	0.05	0.0	1.0
Bacterial Spot resistant	prop.	0.00	0.05	0.0	1.0
High solids	prop.	0.06	0.23	0.0	1.0
High yield	prop.	0.06	0.23	0.0	1.0
Early	prop.	0.13	0.33	0.0	1.0
Thin	prop.	0.14	0.35	0.0	1.0
Intermediate	prop.	0.23	0.42	0.0	1.0
Thick	prop.	0.58	0.49	0.0	1.0
Pear-shaped	prop.	0.01	0.08	0.0	1.0

1.4 Methods

The aim is to estimate the effect of weather on processing tomato yield, quality, and revenue. Since precipitation during the growing season is scant and growers control the amount of water applied through irrigation, we focus on the effect of temperature exposure and include precipitation as a control.

Since we are interested in the direct and indirect effects of temperature exposure, we take care not to introduce bad controls – variables that are themselves outcome variables. A key example is irrigation volumes, which is itself a function of temperature and also affects the outcome variable. If irrigation volumes were included, the coefficient on temperature exposure could be biased because some effects of temperature may be incorrectly attributed to irrigation volume.

We follow the standard approach proposed by Schlenker & Roberts (2009) and adopted by Gammans et al. (2017) and Shew et al. (2020) among others. Ortiz-Bobea (2021) provides a comprehensive summary. We take an off-the-shelf econometric approach to emphasize that results are driven by differences in focus and setting rather than differences in methodology. We estimate:

$$y_{it} = \int_h f(h)\phi_{it}(h)d(h) + \delta z_{it} + \alpha_{g(i)} + \psi(t) + \epsilon_{it} \tag{1.4}$$

where y_{it} is a log-transformed outcome variable (yield, quality, and revenue) in field i in year t, $\alpha_{g(i)}$ is a grower fixed effect, and $\psi(t)$ is a quadratic year trend. The first term characterizes the relationship between temperature exposure and the outcome variable, where $f(h)$ is the marginal effect of temperature h and $\phi_{it}(h)$ is the growing-season density of exposure at h for field i in year t. This continuous representation is not tractable for estimation, but can be approximated using the restricted cubic spline specification detailed in Section 1.4.2. Field-year specific control variables z_{it} include variety-specific attributes (extended field storage, various disease-resistance traits, high solids, high yield, early, thin, intermediate, thick, and pear-shaped), irrigation technology (drip, sprinkler, furrow), soil type (alluvium, eolian, organic material, lacustrine, and residuum), growing season precipitation, a dummy for planting week, and the difference between actual growing days and estimated growing days specified by the seed manufacturer.

As with any annual crop, growers influence the weather they expect to receive through their choice of planting date. In our setting, tomatoes planted earlier are exposed to cooler temperatures on average, whereas tomatoes planted later are exposed to hotter temperatures on average. The implication is that weather is potentially endogenous and coefficients on temperature exposure may be biased. We include dummies for planting week-of-year to account for this endogeneity in weather. If we fail to control for planting date we obtain biased estimates of effects on quality, but results for yield and revenue are largely unchanged (see Appendix 1.A).

The error term ϵ_{it} is likely heteroskedastic, spatially correlated, and temporally correlated within similar growers over time. We use heteroskedatic robust standard errors two-way clustered by grower group[3] and county by year. We cluster at the grower group level to account for possible dependence among growers within the same grower group. We do not cluster by year as we observe only 10 years of data. Even in a multiway cluster, too few clusters in any cluster group will result in incorrect statistical inference (Cameron et al., 2011). Instead, we cluster county by year to account for spatial correlation. Results are robust to using spatial heteroscedasticity and autocorrelation consistent errors that allow for spatial correlation between nearby fields and serial correlation in panel data (see Appendix 1.B).

We include grower fixed effects $\alpha_{g(i)}$, where individual growers can be associated with multiple fields – on average, each grower is associated with 27 field-year observations. This controls for time-invariant, grower-specific factors that may be related to outcome or explanatory variables. Our preferred specification uses grower fixed effects as regular crop rotation

[3]Recall that the grower group identifier links growers within the same network or organization.

results in an unbalanced panel of field-year observations. Note that results are robust to using field fixed effects in place of grower fixed effects (see Appendix 1.C). Results are also robust to replacing quadratic year trends with (a) linear year trend, and (b) county-specific quadratic trends.

We refrain from using year fixed effects. Adding year fixed effects (akin to adding state-by-year fixed effects since we only observe one state) would absorb much of the useful variation in temperature exposure used to identify the effects of interest. The inclusion of state-by-year fixed effects in Deschênes & Greenstone (2007) is critiqued by Fisher et al. (2012) because "state-by-year [fixed effects] absorb almost all variation and the identification rests on very slim margins, so even small amounts of measurement error will be greatly amplified". Without the inclusion of year fixed effects, we observe considerable variation in temperature exposure (Appendix 1.D shows the distribution of temperature exposure across counties).

Since we do not include year fixed effects, one may be concerned that we are not controlling for temporal correlation in the prices and revenue from common market conditions. Indeed, processing tomatoes are used to produce a storable commodity and therefore base prices are temporally correlated. For example, a negative production shock caused by poor weather in year $t-1$ will induce processors to draw down stores and offer a higher base price in year t. The amount moving in and out of storage is an omitted variable and forms part of the error in the price and revenue regressions. This causes the error to be correlated with weather in the previous year, thus introducing endogeneity bias. We avoid this issue through the way we construct price and revenue, explained in Section 1.3. Before any analysis, we remove variation in price (and revenue) that stems from market conditions including changes in storage and preserve variation caused by quality.

1.4.1 Estimating temperature exposure

We translate daily observations of minimum and maximum temperatures to a measure of temperature exposure for each field-year observation. For each day of the growing season[4],

[4]The growing season starts on the day of planting and ends on the last day of harvest for each field.

we estimate how many hours are spent in 1°C temperature intervals by fitting a sinusoidal curve between each day's minimum and maximum temperature. We then sum over days to estimate the number of days spent in each 1°C temperature interval during the entire growing season. The result is x_{it}, a 1-by-J vector of temperature exposures for field i during the growing season in year t, where J is the number of temperature bins. In our setting, we bin temperatures from 5°C to 41°C, so $J = 37$.[5]

$$x_{it} = \begin{pmatrix} x_{it,5} & x_{it,6} & \dots & x_{it,40} \end{pmatrix} \tag{1.5}$$

where $x_{it,j}$ is the number of days spent between j°C and $(j+1)$°C during the growing season in year t.

This approach has several advantages. First, it addresses the empirical challenge of mixed frequency between regressor and outcome variables. We have many daily observations of minimum and maximum temperatures to match with one annual observation of an outcome variable. Averaging daily temperatures across the growing season would mask differences in exposure to extreme temperatures. The second advantage of this approach is that it preserves the temperature distribution. This allows us to uncover the marginal effect of exposure to different temperatures.

1.4.2 Restricted cubic spline specification

Next, we choose a functional form to characterize the relationship between outcome variables and temperature exposure. Mid-range temperatures are thought to be ideal for yield and quality of processing tomatoes, but these may be reduced by hot or cool temperatures if exposure occurs during key stages of the plant's growth cycle (Hartz et al., 2008). The implication is that the relationship between temperature exposure and outcome variables is nonlinear.

To capture nonlinearity in the response of the outcome variable to temperature, we

[5] Temperatures range from −1°C to 45°C. We aggregate temperature exposure below 5°C and above 41°C to avoid bins with little exposure.

estimate a restricted cubic spline model (otherwise known as a natural cubic spline). The restricted cubic spline model has become popular because it offers several benefits over alternative methods for estimating nonlinear temperature effects (Berry et al. (2014); D'Agostino & Schlenker (2016); Ortiz-Bobea et al. (2019); Blanc & Schlenker (2020); Bucheli et al. (2022)). First, it offers smooth parsimonious semiparametric estimation without needing to define critical temperature thresholds. Second, it imposes a restriction that its tails (i.e. before the first knot and after the last knot) are linear. This reduces overfitting in the data-sparse tails of the temperature distribution, an issue with the polynomial and cubic spline functional forms. We estimate a piecewise linear degree day model as a robustness check. Overall, results from the two specifications – a piecewise linear degree day model and a restricted cubic spline model – are consistent in terms of economic and statistical significance (see Appendix 1.E).

We identify $K = 4$ temperatures that split the distribution of temperature exposure by interval into quintiles. This accounts for the fact that relatively less time is spent at extreme temperatures. Unlike the piecewise linear model, knot placement does not strongly influence the cubic spline results because the marginal effect of exposure is allowed to vary smoothly between knots.

Next, we introduce a basis matrix B and a vector of coefficients Γ. B is J-by-P while Γ P-by-1, where P is the number of parameters on temperature to be estimated, and is directly related to the number of knots K. A restricted cubic spline with $K = 4$ knots results in $P = 3$, which is smaller than the number of temperature exposure bins $J = 37$. This is an advantage of the spline model – it reduces the dimensionality while still allowing for flexible semiparametric estimation. The derivation of the basis matrix B that corresponds to the restricted cubic spline is shown in Appendix 1.F.

Under these assumptions, we can write Equation 1.4 as:

$$y_{it} = x_{it} B\Gamma + \delta z_{it} + \alpha_{g(i)} + \psi(t) + \epsilon_{it} \tag{1.6}$$

Stacking n observations across fields and years gives Equation 1.6 in matrix notation:

$$Y = XB\Gamma + \delta Z + \alpha + \psi + \epsilon \tag{1.7}$$

where X is a n-by-J matrix of temperature exposures, and Y, δZ, α, ψ and ϵ are n-by-1 vectors of outcomes, controls, grower fixed effects, quadratic time trends, and errors respectively.

After estimation, we recover the marginal effect of temperature exposure evaluated at each interval. We pre-multiply the vector of estimated coefficients $\hat{\Gamma}$ by the corresponding B matrix. The resulting J-by-1 vector $\hat{\beta}$ is the marginal effect of spending one additional day at each temperature bin $j = 1, \ldots, J$.

$$\underset{J\times 1}{\hat{\beta}} = \underset{J\times P}{B} \times \underset{P\times 1}{\hat{\Gamma}} \tag{1.8}$$

Last, we derive an estimate of the variance-covariance matrix for $\hat{\beta}$:

$$\underset{J\times J}{\widehat{var}(\hat{\beta})} = \underset{J\times P}{B} \times \underset{P\times P}{\widehat{var}(\hat{\Gamma})} \times \underset{P\times J}{B'} \tag{1.9}$$

1.5 Results

Figure 1.2 displays results for the effects of temperature exposure on our three key outcome variables: yield, quality, and revenue. In each figure, the top graph shows the effect of an additional 24 hours spent in a given temperature interval on the outcome variable relative to 24 hours spent at 26°C. This temperature has the greatest exposure in our sample and represents average temperatures. The 95% confidence intervals account for the possibility of heteroskedasticity, spatial correlation, and temporal correlation in the errors. The gray vertical lines show the positions of the knots. The histogram at the bottom of the frame shows the average exposure to each temperature interval during the growing season across all fields in all years.

For yield, we find that the optimal temperature is around 28°C. Exposure to temperatures above 35°C leads to significantly lower yields. An additional 24 hours of exposure to 40°C decreases yield by almost 2% on average relative to 24 hours of average temperatures. Exposure to cold temperatures below 10°C causes a small but significant decline in yield by 0.7% relative to 24 hours at average temperatures.

We isolate variation in price driven by quality by applying quality adjustments to the 10-year average base price. For quality, we find that the optimal temperature is around 20°C. Quality declines with exposure to hot conditions. An additional 24 hours of exposure above 30°C causes quality to drop by up to 0.2% relative to 24 hours spent at average temperatures. We show results for individual quality defects and bonuses in Appendix 1.G Figures 1.G.1 and 1.G.2. The presence of defects, specifically limited-use tomatoes, material other than tomatoes, green tomatoes, and mold, all increase with exposure to hot temperatures although imprecision in the estimates means we cannot rule out null effects. The presence of limited use tomatoes significantly declines with exposure to cool temperatures relative to average temperatures, resulting in higher quality. The effect of temperature on the solids bonus is imprecise but the point estimate declines with exposure to hot temperatures.

Revenue is maximized with exposure to temperatures around 27°C. An additional 24 hours of exposure to 40°C decreases revenue by 2.3% compared to 24 hours of average temperatures. This is expected since both yield and quality respond negatively to hot conditions. Exposure to cool temperatures below 10°C causes a smaller but significant decrease in revenue by almost 1% relative to 24 hours at average temperatures.

It is plausible that the effect of temperature exposure varies across growers and across years. For example, a high base price increases the quality incentive on a dollar per ton basis which may induce growers to put more effort into raising quality than they would in a low base price year. These estimates average over all responses, which may be heterogeneous.

Appendix 1.H shows the estimated effects of the control variables. Precipitation negatively affects yield (and therefore revenue), but the magnitude is relatively small – a one

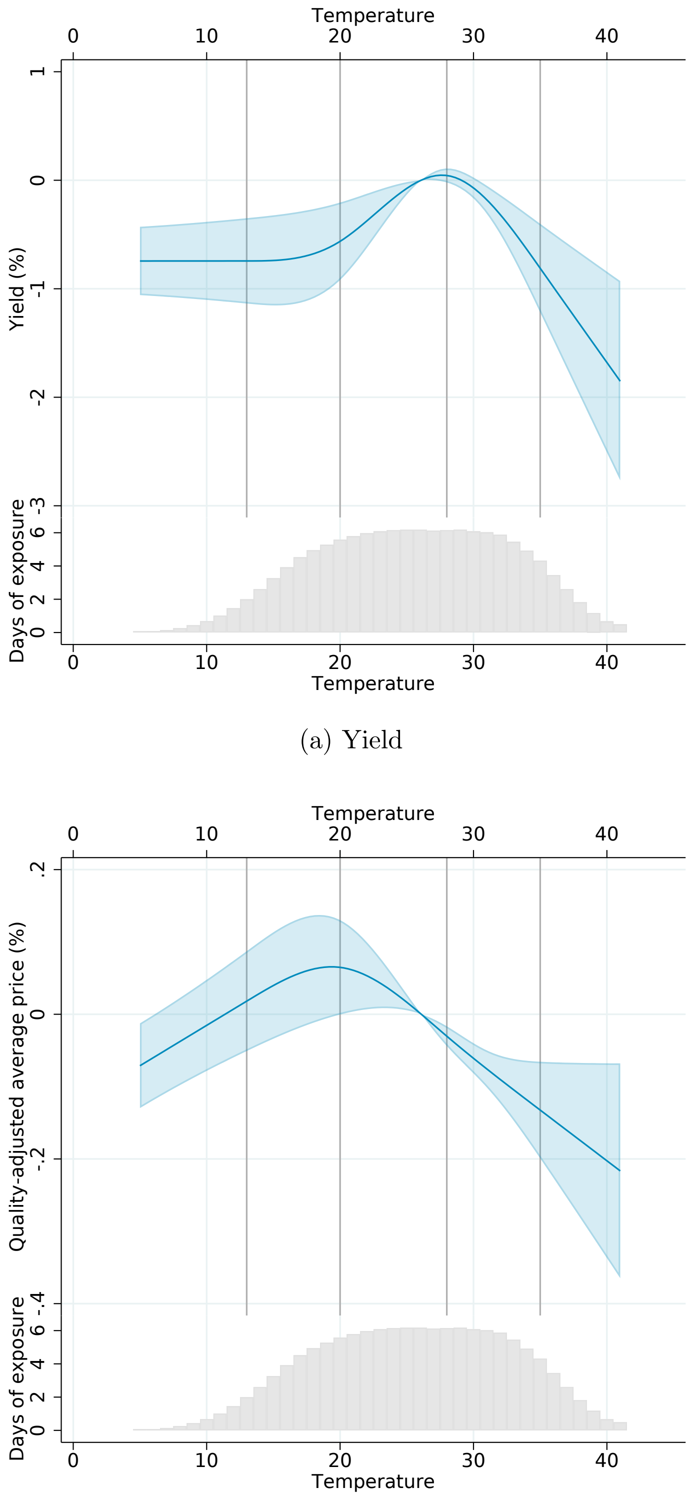

(a) Yield

(b) Quality

Figure 1.2: Restricted cubic spline results

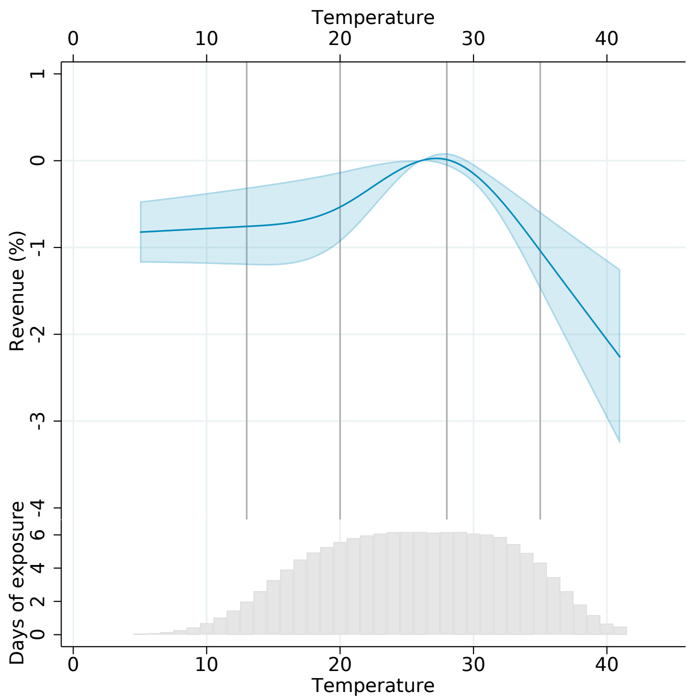

(c) Revenue per acre

Figure 1.2: Restricted cubic spline results

Notes: For each figure, the graph at the top of the frame shows the effect of an additional 24 hours spent at a given temperature interval on the outcome variable relative to 24 hours spent at 26°C. The histogram at the bottom of the frame shows the average exposure to each temperature interval during the growing season across all fields in all years.

standard deviation increase in precipitation decreases yield by 0.2% on average. Fields with drip irrigation are associated with higher yields on average compared to fields using furrow or sprinkler irrigation techniques. Soil type and varietal characteristics are also associated with yield and quality. For example, an early variety, one that requires fewer days to reach maturity and therefore has a shorter season, is associated with lower yields but higher quality on average. The quadratic year trend is insignificant, suggesting that yield and quality were stationary over the course of our 10-year sample.

1.5.1 Decomposition

We find that temperature exposure significantly affects both yield and quality. However, the relative importance of each pathway is unclear a priori. Here, we decompose the effect of temperature exposure on revenue per acre (total effect) into the effect on revenue driven by yield (yield effect) and effect on revenue driven by quality (quality effect). This allows us to answer two questions. What is the relative importance of the yield and quality effects? And, would estimates of yield and quality effects be biased if quality considerations were omitted?

The total effect, which captures both yield and quality pathways, is equal to the effect of temperature exposure on revenue estimated above. Since we use log-transformed variables, the yield effect is similarly equal to the effect of exposure on yield estimated above. The quality effect, however, is slightly larger than the effect of exposure on quality index shown above. Recall revenue is a function of paid tons – the processor won't pay producers for some poor quality tonnage. Therefore, quality can affect revenue via (a) changes in price captured in the quality index, and (b) changes in paid tonnage. We estimate the quality effect as the difference between the yield effect and total effect in Equation 1.10, which captures the effect of temperature on both price and unpaid tons.

$$\begin{aligned}\text{quality effect} &= \text{total effect} - \text{yield effect} \\ &= \hat{\beta}^{\ln(\text{revenue per acre})} - \hat{\beta}^{\ln(\text{yield})}\end{aligned} \tag{1.10}$$

As shown in Figure 1.3, while the yield effect dominates grower revenue, quality also plays

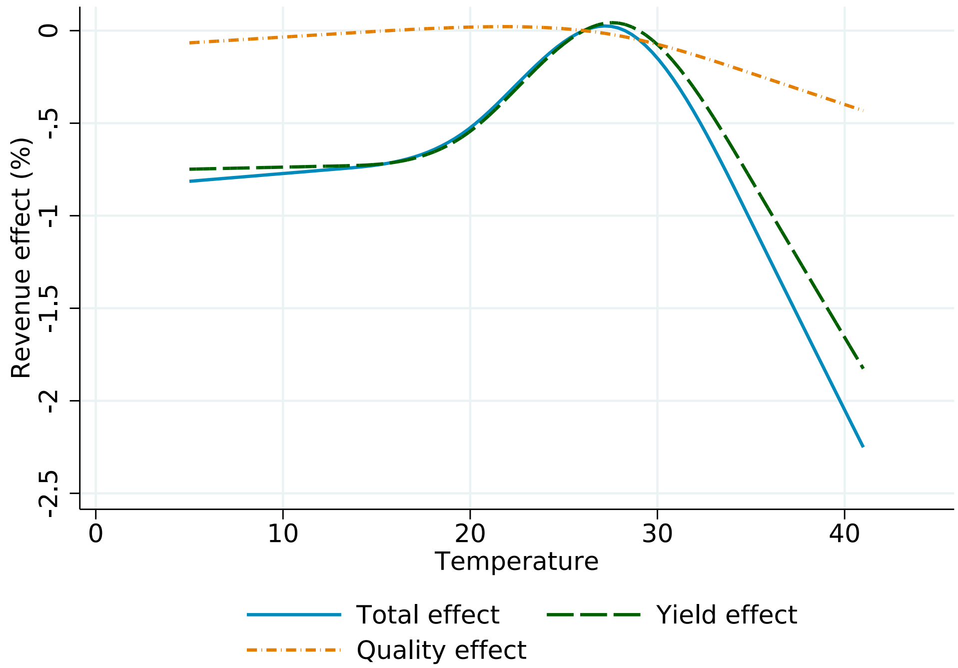

Figure 1.3: Decomposing the effect of exposure on revenue into the yield effect and quality effect

an important role. Without access to data on quality, a researcher can only recover the yield effect: an additional 24 hours of exposure to 40°C decreases revenue by 1.8% compared to 24 hours of average temperatures. This underestimates the effect of exposure on revenue by up to 0.5 percentage points, or 20% of the point estimate at 40°C. Failing to account for quality's effect on revenue biases estimates of temperature on revenue.

1.5.2 Selection

In almost all published work, researchers rely on the assumption that observations of quality and yield accurately reflect conditions at harvest to recover unbiased estimates of temperature and climate change damages. For many agricultural products, data are available only for the subset of production that producers select to be graded. This may bias observations of both quality and yield – e.g. if only high quality product is graded, yield will be underestimated and quality overestimated.

A benefit of our setting is that selection is minimal. Recall from Section 1.2 that ob-

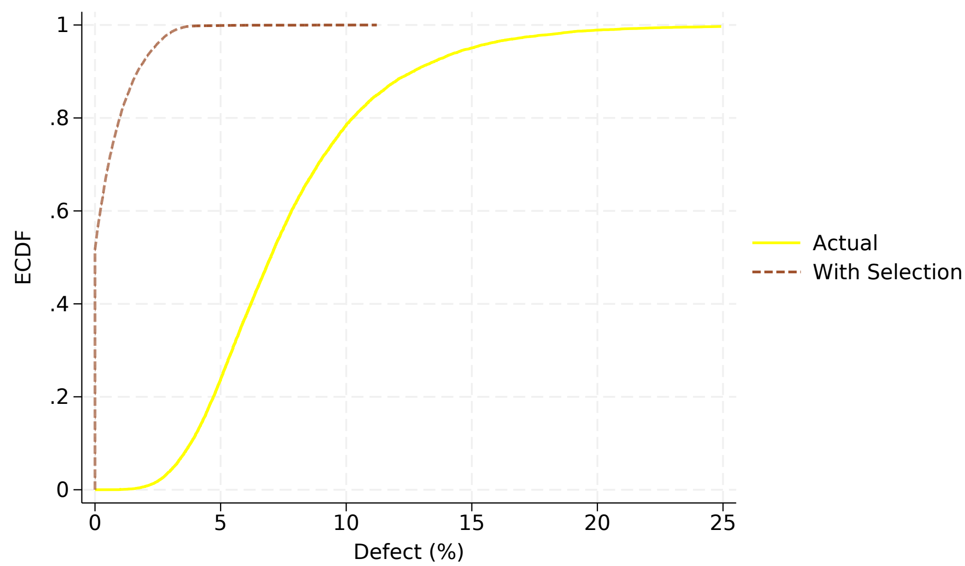

Figure 1.4: Empirical CDF of defects

servations of quality and yield are close to those in the field because of mandatory grading, as well as contracting and harvesting practices. Moreover, the processor engages in sorting in which some defective tonnage isn't counted towards the quantity on which producers are paid. We are able to observe this sorting through differences in total tons (before sorting) and paid tons (after sorting) – paid tons is on average 7% smaller than total tons[6]. This gives us the opportunity consider the consequences of selection for the resulting estimates.

For estimates with selection, we assume that the researcher only observes yield and quality of the paid tons i.e. a portion of total tons is now unobserved. Consistent with the processor sorting defective tonnage, we assume that unobserved tons had defects. Yield under selection is artificially reduced because it is calculated using a smaller tonnage (paid tons) than what is actually harvested from the field (total tons). Quality under selection is artificially improved because some tonnage with defects is no longer observed. Figure 1.4 compares the actual distribution of defects to its distribution with selection.

Our first hypothesis is that the effect of exposure to high temperatures on quality will

[6]Recall that the quantity for which producers are paid does not include some flawed tonnage because it is not commercially viable and is disposed of by the processor.

be biased towards zero under selection. This follows from the result that high temperatures negatively affect quality, and the assumption that lower quality products are being withheld. Our second prediction is that the negative effect on quality will be incorrectly assigned to yield, causing the negative yield effect to increase in magnitude.

We compare estimates from our preferred specification with estimates using observations with added selection. Consistent with our hypothesis, the effect of exposure on yield is biased upwards by up to 10%. The effect of high temperatures on quality is indeed biased towards zero under selection by up to 66%. The upward bias in yield partially offsets the downward bias in quality. However, bias still remains and the effect of exposure to high temperatures on revenue is attenuated by up to 6%. Our results suggest that estimates will be biased in settings with selection but that the magnitude will depend on (a) the actual effect of weather on quality, and (b) how much selection is occurring.

1.6 Climate projection

Armed with the knowledge that processing tomatoes are susceptible to extreme temperatures, a natural follow-up question is: How will climate change affect production of processing tomatoes? We estimate the impact of climate change by comparing predicted outcomes using realized weather with predicted outcomes using future weather projected in climate models.

We select four global climate models – Access CM2, HadGEM 3 GC31-LL, EC Earth 3, and EC Earth 3 - Veg – that are included in Coupled Model Intercomparison Project Phase 6 (CMIP6) and used by the Intergovernmental Panel on Climate Change in their latest assessment report (IPCC, 2023). These models best capture relevant aspects of California's climate (Krantz et al., 2021) and have been statistically downscaled by Pierce et al. (2023) to a 3km resolution and daily time step (available on Cal-Adapt (2023)).

Uncertainty in the impact of climate change on economic outcomes stems from several sources. First, there is statistical uncertainty in the historical relationship between weather variables and the outcomes of interest. Second, it is unclear how much emissions the world will emit into the future. Third, conditional on a particular emissions scenario, there is model

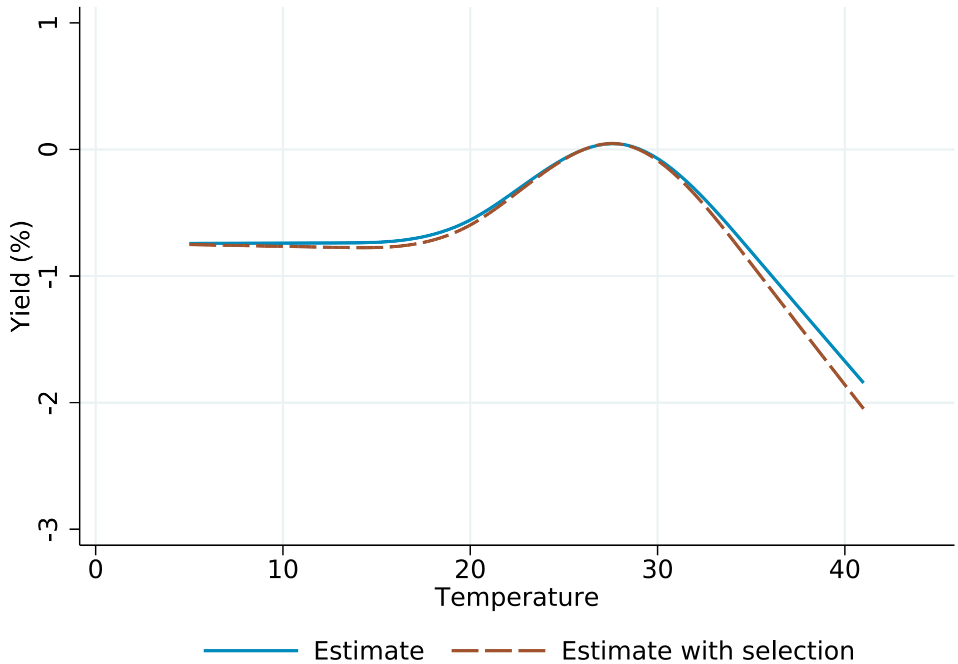

(a) Yield

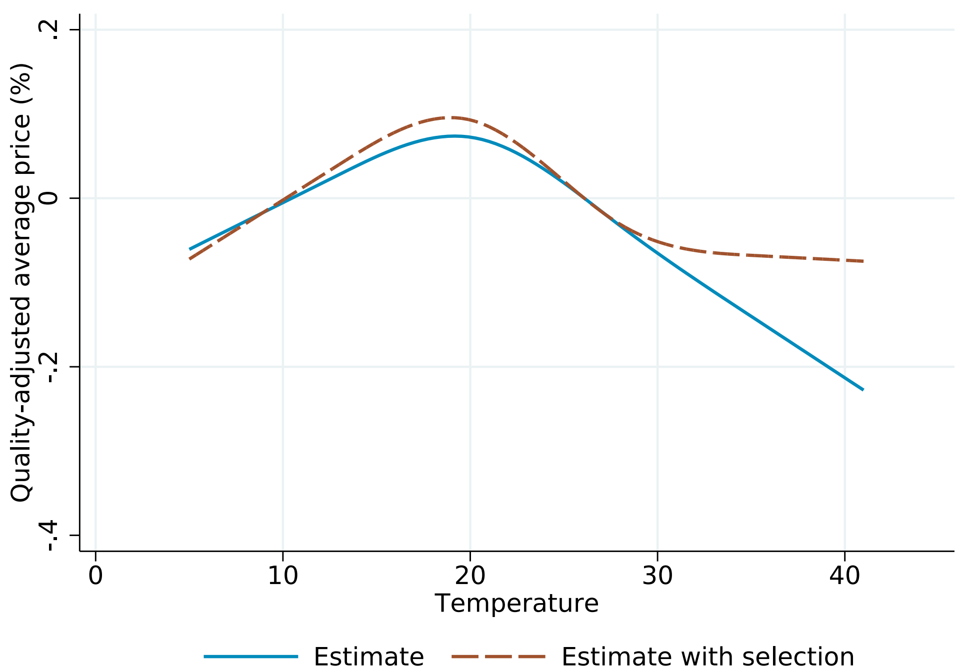

(b) Quality

Figure 1.5: Selection bias

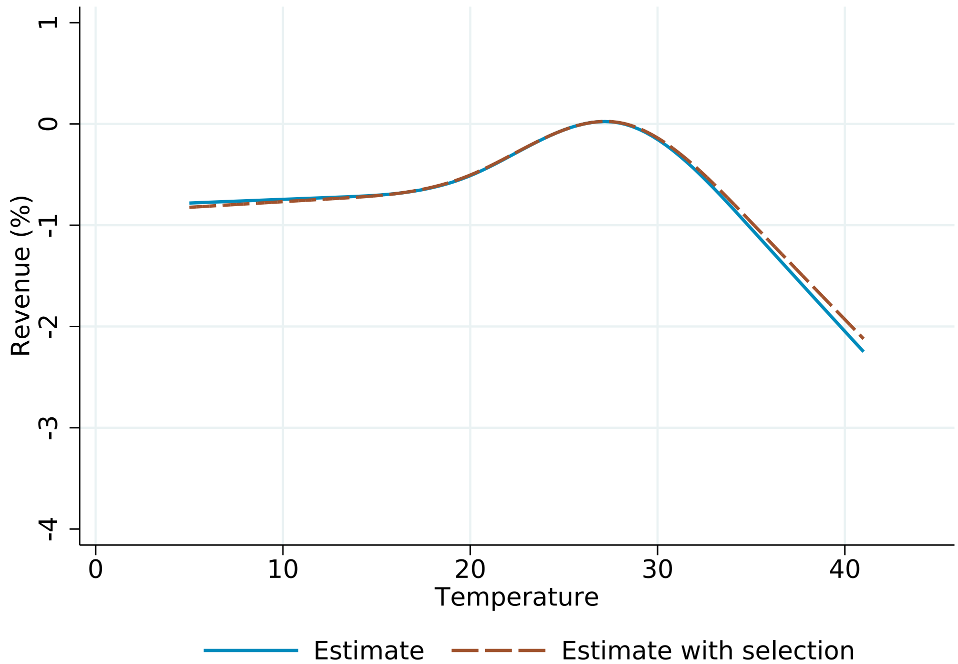

(c) Revenue per acre

Figure 1.5: Selection bias

uncertainty in how a particular level of emissions will change the climate in a particular location.

We estimate the impact of climate change on processing tomato yield, quality, and revenue following Burke et al. (2015) to account for all three sources of uncertainty. To account for emissions uncertainty, we present results for two emissions scenarios in CMIP6: SSP2-4.5, a middle-of-the-road global emissions scenario, and SSP5-8.5, a very-high global emissions scenario.

To account for climate model uncertainty, we implement the model democracy approach routinely used by climate scientists (Burke et al., 2015). We use four climate models and where available, we collect multiple climate projections or “ensembles” from each model that have varying baseline conditions. This yields a total of 10 modeled projections for each emissions scenario. Each projection is given an equal vote in determining the impact estimate.

For each combination of projection and emissions scenario, we do 1,000 wild cluster bootstrap replications to account for statistical uncertainty in the historical relationship between weather variables and the outcomes. The wild cluster bootstrap preserves spatial dependence in the resampled data and performs well even if there are relatively few clusters (Cameron et al., 2011). The wild cluster bootstrap procedure is as follows. First, we obtain residuals $\hat{\epsilon}$ and predicted coefficients after estimating our regression model in Equation 1.7. Next, we generate bootstrap samples denoted by $*$ and indexed by b for each grower group and year cluster g:

$$Y_g^{*b} = X_g B\hat{\Gamma} + \hat{\delta}Z + \hat{\alpha} + \hat{\psi} + \epsilon_g^{*b}, \qquad \epsilon_g^{*b} = w_g^{*b}\hat{\epsilon}_g \tag{1.11}$$

where ϵ_g^{*b} is the randomly reshuffled residuals $\hat{\epsilon}_g$ from the same grower group and year cluster multiplied by a wild weight w_g^{*b} drawn from the Rademacher distribution (i.e. -1 or 1 with equal probability). We then regress Y^{*b} on X to obtain $\hat{\Gamma}^{*b}$. Finally, we predict $\hat{Y}_{itp}^{*b}$ which replaces x_{it}, temperature exposure experienced at field i during year t, with x_{ip}, temperature exposure projected to be experienced at field i in future year p at mid-century ($p = 2041, \ldots 1950$) or end of century ($p = 1991, \ldots, 2100$). All other controls and fixed effects are kept at their year t levels. This simulates outcomes as if the field was exposed to temperatures from future year p instead of the actual temperatures experienced in year t.

We estimate the climate change impact for each bootstrap replication as the difference between the predicted outcome using actual temperatures and predicted outcomes using projected temperatures from two time frames: mid-century, 2041-2050, and end of century, 2091-2100. Finally, we stack the 1,000 bootstrap estimates from each of the 10 projections into a vector of 10,000 impact estimates to form a distribution that accounts for both statistical and climate uncertainty. We construct a confidence interval by taking the 2.5th and 97.5th percentiles to calculate the range containing 95% of impact estimates.

We predict climate change will result in economic damages for processing tomato growers, as shown in Figure 1.6. By mid-century, yield and revenue per acre are predicted to be significantly lower than their 2011-2020 levels. We estimate median losses of yield from 5% to 9% and revenue from 9% to 14% depending on the emissions scenario. The median effect on quality by mid-century is negative but the 95% confidence interval includes zero for both emission scenarios. By the end of the century, yield and quality are expected to both decline significantly by 13% and 2% respectively under a middle-of-the-road emissions scenario, with even larger losses predicted under a very-high emissions scenario (53% and 8% respectively). Revenue per acre is predicted to decline significantly by the end of the century, with median losses of 21% or 80% depending on the emissions scenario. However, this analysis does not account for general equilibrium effects that would likely see an increase in the base price, partially offsetting losses in revenue.

An implicit assumption is that producers, processors, seed manufacturers, and other industry participants respond to future changes in climate in a similar way to how they respond to weather shocks during our sample period. Our results highlight the need for new adaptive responses to reduce the effect of weather and climate change on yield and quality.

1.7 Discussion and Conclusion

Studies of the impacts of climate change on agriculture typically focus on yield (see Ortiz-Bobea (2021)) or farmland values (e.g. Mendelsohn et al. (1994); Bareille & Chakir (2023)). Quality matters because of its role in contractual arrangements and price determination but is typically ignored in earlier work. We use novel data from a large tomato processor to study the effects of temperature exposure on yield, quality, and grower revenue. We find exposure to hot temperatures reduces grower revenue through two channels: yield and quality. If we had failed to account for quality's effect on revenue, our estimates of the effects of temperatures on revenue would have been biased by up to 20%. Our uniquely detailed data gives us a complete and precise picture of how weather and climate change affect product quality.

Prior work uses observational data to analyze effects of weather and climate on the quality of different agricultural products, including grains like rice and wheat (Kawasaki &

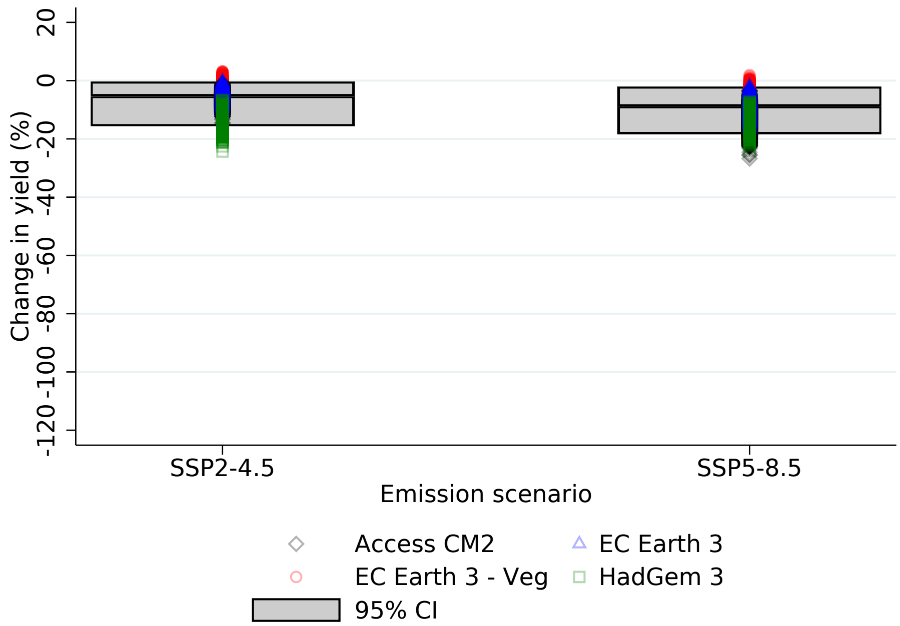

(a) Yield, mid-century

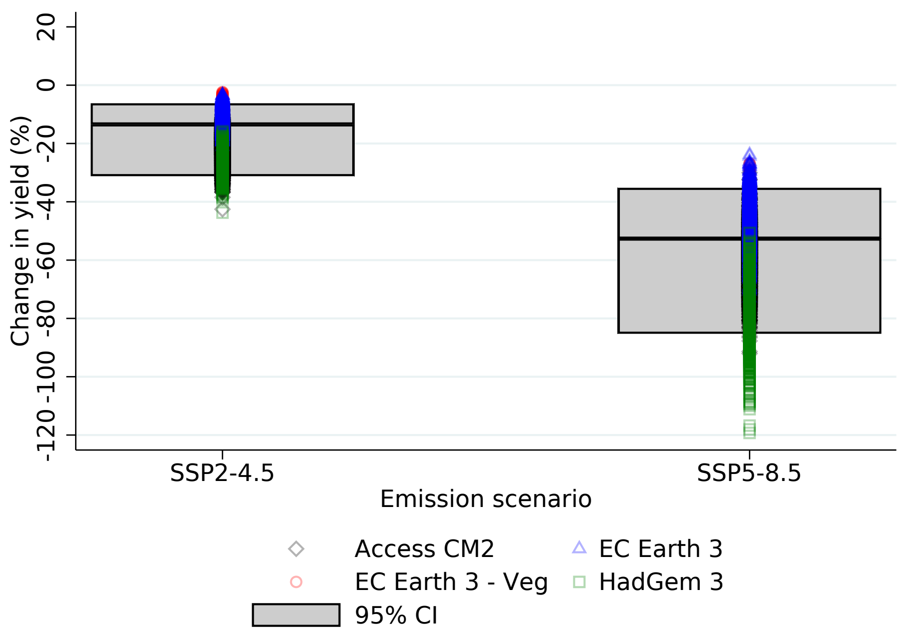

(b) Yield, end of century

Figure 1.6: Projection of climate impacts by mid-century, 2041-2050, and end of century, 2091-2100, relative to a 2011-2020 baseline

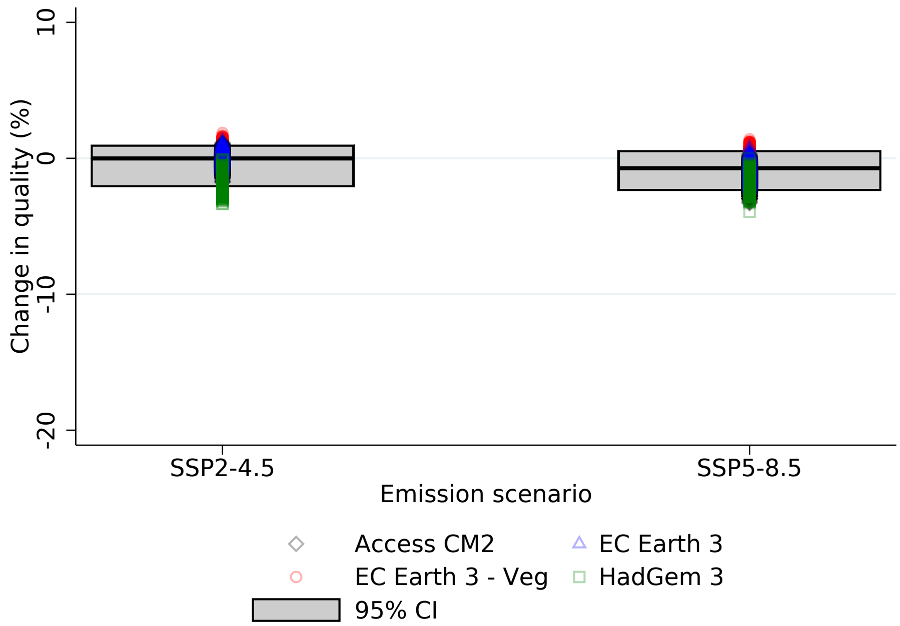

(c) Quality, mid-century

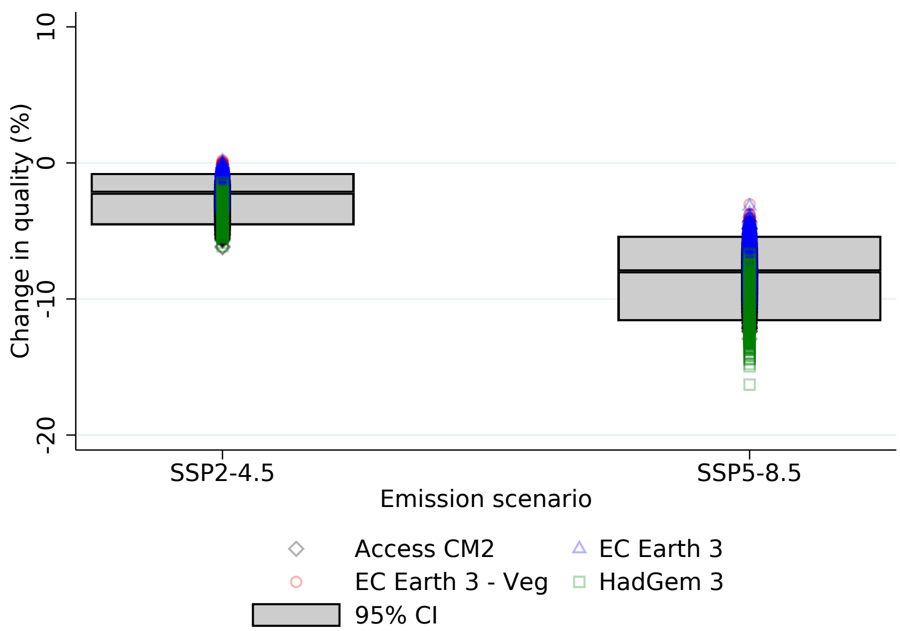

(d) Quality, end of century

Figure 1.6: Projection of climate impacts by mid-century, 2041-2050, and end of century, 2091-2100, relative to a 2011-2020 baseline

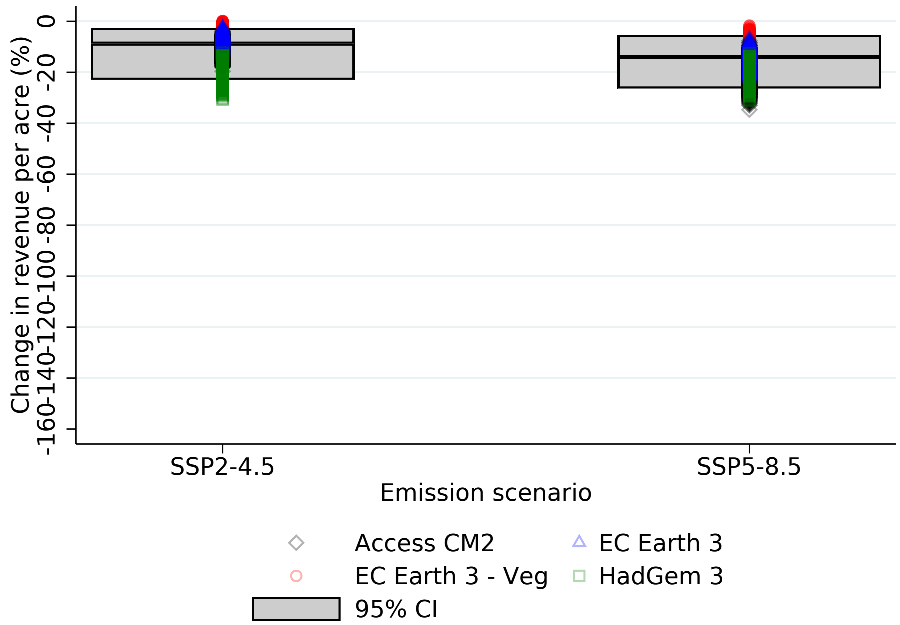

(e) Revenue per acre, mid-century

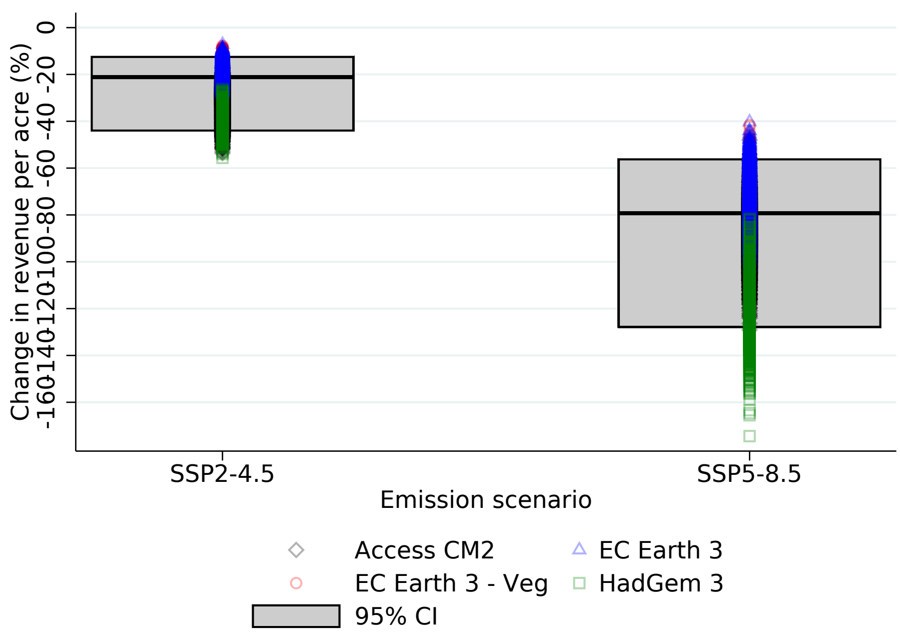

(f) Revenue per acre, end of century

Figure 1.6: Projection of climate impacts by mid-century, 2041-2050, and end of century, 2091-2100, relative to a 2011-2020 baseline

Notes: Each point is an estimate of the projected impact derived from a single combination of projection, emissions scenario, and wild cluster bootstrap replication. The thick black lines represent the median impact estimate and the shaded grey areas represent the 95% confidence intervals that account for statistical and climate uncertainty.

Lichtenberg (2014); Kawasaki & Uchida (2016); Kawasaki (2019)), field crops like peanuts (Ramsey et al., 2020), and specialty crop products like wine (see Ashenfelter & Storchmann (2016)), apples (Dalhaus et al., 2020), and tobacco (Ramsey & Rejesus (2021), although the consequences of climate change are not considered). Of these papers, Kawasaki & Uchida (2016), Kawasaki (2019), Dalhaus et al. (2020) and Ramsey et al. (2020) link changes in quality to grower revenue as required to quantify the economic consequences of weather and climate change. Dalhaus et al. (2020) infer quality from unexplained differences in price. Kawasaki & Uchida (2016) and Kawasaki (2019) use quality grades which inhibits their ability to analyze individual quality attributes. Ramsey et al. (2020) observe one aspect of peanut quality, kernel size, and proxy for its effect on price using value formulas from the Commodity Credit Corporation's loan rates. It is unclear if this captures actual market pricing and all relevant dimensions of quality. In our setting, quality attributes are precisely measured, rather than inferred, and directly linked to price using a schedule of bonuses and deductions established prior to planting.

In addition, studying the effects of climate change on agricultural production with observational data presents identification challenges. Typically, data are available only for the subset of production that growers choose to market because it exceeds an implicit or explicit quality threshold[7]. This has two consequences. First, observations of quality are biased measures of quality at harvest. Kawasaki & Uchida (2016) find low quality rice is less likely to undergo the costly process of being graded and more likely to be withheld for self-consumption or sale on informal markets (such as to local households or for animal feed). Quality measures for hand-harvested crops, such as berries, stone fruit, apples, and leafy greens, may also be biased. Harvesting guidelines (e.g. instructions to only pick high-quality produce), adequate access to labor, and how labor is paid can all affect the observed quality of hand-harvested crops (e.g. Hill & Beatty (2020)).

Second, yield estimates will be biased if weather affects quality and quantities are mea-

[7] An exception is data from field trials such as those used by Ramsey et al. (2020), Tack et al. (2015) and Tack et al. (2017b), which are less likely to suffer from selection bias.

sured after farmers screen out low-quality products. In this situation, quality effects will be attributed falsely to yield and estimates of the effects of weather on both yield and quality will be biased. For example, wet conditions around harvest can result in quality problems in grain that may cause growers to withhold low quality grain from sale. If this withheld grain is not counted towards production totals, an outside observer will incorrectly describe the effects of wet conditions around harvest as reducing yield instead of quality. The overall effect on estimates of revenue or profit is ambiguous since bias in estimated yield is at least partially offset by bias in estimated quality. This concern is not limited to papers focusing on quality but also applies to *all* studies estimating the effects of weather on yield or revenue that don't address selection.

In our setting, processing tomatoes are always grown under contract, mechanically harvested, and graded at an independent state inspection station. These institutional factors mean that our analysis is unlikely to suffer from significant selection bias. It also gives us the opportunity to estimate selection as experienced in other settings and quantify the magnitude of the bias. We find that selection biases estimates of the effect of temperature on quality, yield, and revenue. These results illustrate the consequences of data limitations that should be carefully considered in all empirical work.

Finally, continued warming predicted under climate change is a cause for concern. Growers cannot fully mitigate damages caused by extreme heat, indicating that the processing tomato industry will be susceptible to the harm from the continued warming predicted as the climate changes. Absent additional adaptation, we predict climate change will cause both yield and quality of California's processing tomatoes to decline by the end of the century, inducing a significant loss of grower revenue. This stands in contrast to previous work that finds irrigation can mitigate the effects of heat on yield. Rather, our results suggest that irrigated agriculture can be susceptible to climate change. These findings reinforce the need for investment in research into and development of heat-tolerant varieties and related technologies.

A caveat of our research is its focus on a single processor operating in one specific agricultural industry. While our narrow focus comes at the expense of potentially limiting the external validity of our results, we benefit from detailed and reliable observations of production by individual farmers. This provides new insights into the effect of weather and climate on individual commercial agricultural producers.

References

AgSeeds (2020). Season Review. Available at http://www.agseeds.com/.

Ashenfelter, O. & Storchmann, K. (2016). Climate change and wine: A review of the economic implications. *Journal of Wine Economics*, 11(1), 105–138.

Bareille, F. & Chakir, R. (2023). The impact of climate change on agriculture: A repeat-ricardian analysis. *Journal of Environmental Economics and Management*, (pp. 102822).

Berry, S. T., Roberts, M. J., & Schlenker, W. (2014). Corn production shocks in 2012 and beyond: Implications for harvest volatility. In *The Economics of Food Price Volatility* (pp. 59–81). University of Chicago Press.

Blanc, E. & Schlenker, W. (2020). The use of panel models in assessments of climate impacts on agriculture. *Review of Environmental Economics and Policy*.

Bucheli, J., Dalhaus, T., & Finger, R. (2022). Temperature effects on crop yields in heat index insurance. *Food Policy*, 107, 102214.

Burke, M., Dykema, J., Lobell, D. B., Miguel, E., & Satyanath, S. (2015). Incorporating climate uncertainty into estimates of climate change impacts. *Review of Economics and Statistics*, 97(2), 461–471.

Cal-Adapt (2023). Data derived from LOCA2 Downscaled CMIP6 Climate Projections. Cal-Adapt website developed by University of California at Berkeley's Geospatial Innovation Facility under contract with the California Energy Commission. Available at https://cal-adapt.org/.

California Department of Food and Agriculture (2019). 2018-2019 California Agricultural Statistics Review. Available at https://www.cdfa.ca.gov/statistics/PDFs/2018-2019AgReportnass.pdf.

Cameron, A. C., Gelbach, J. B., & Miller, D. L. (2011). Robust inference with multiway clustering. *Journal of Business & Economic Statistics*, 29(2), 238–249.

Cammarano, D., Jamshidi, S., Hoogenboom, G., Ruane, A. C., Niyogi, D., & Ronga, D. (2022). Processing tomato production is expected to decrease by 2050 due to the projected increase in temperature. *Nature Food*, (pp. 1–8).

Carter, C., Cui, X., Ghanem, D., & Mérel, P. (2018). Identifying the economic impacts of climate change on agriculture. *Annual Review of Resource Economics*, 10, 361–380.

Carter, E. K., Melkonian, J., Riha, S. J., & Shaw, S. B. (2016). Separating heat stress from moisture stress: analyzing yield response to high temperature in irrigated maize. *Environmental Research Letters*, 11(9), 094012.

Chen, S., Chen, X., & Xu, J. (2016). Impacts of climate change on agriculture: Evidence from china. *Journal of Environmental Economics and Management*, 76, 105–124.

Conley, T. G. (1999). GMM estimation with cross sectional dependence. *Journal of econometrics*, 92(1), 1–45.

D'Agostino, A. L. & Schlenker, W. (2016). Recent weather fluctuations and agricultural yields: implications for climate change. *Agricultural economics*, 47(S1), 159–171.

Dalhaus, T., Schlenker, W., Blanke, M. M., Bravin, E., & Finger, R. (2020). The effects of extreme weather on apple quality. *Scientific reports*, 10(1), 1–7.

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

FAO (2023). Food and Agriculture Organization Corporate Statistical Database (FAOSTAT). Available at `https://www.fao.org/faostat/en/`.

Fisher, A. C., Hanemann, W. M., Roberts, M. J., & Schlenker, W. (2012). The economic impacts of climate change: evidence from agricultural output and random fluctuations in weather: comment. *American Economic Review*, 102(7), 3749–3760.

Gammans, M., Mérel, P., & Ortiz-Bobea, A. (2017). Negative impacts of climate change on cereal yields: statistical evidence from France. *Environmental Research Letters*, 12(5), 054007.

Hartz, T., Miyao, G., Mickler, J., Lestrange, M., Stoddard, S., Nuñez, J., & Aegerter, B. (2008). Processing tomato production in California. *UC Vegetable Research & Information Center*.

Hill, A. E. & Beatty, T. K. M. (2020). Haste makes waste: Evidence on speed and quality tradeoffs in the workplace. Working paper.

Hsiang, S. M. (2010). Temperatures and cyclones strongly associated with economic production in the Caribbean and Central America. *Proceedings of the National Academy of sciences*, 107(35), 15367–15372.

IPCC (2023). *Sixth Assessment (AR6) Synthesis Report: Climate Change 2023*. Cambridge University Press.

Just, R. E. & Chern, W. S. (1980). Tomatoes, technology, and oligopsony. *The Bell Journal of Economics*, (pp. 584–602).

Kawasaki, K. (2019). Two harvests are better than one: double cropping as a strategy for climate change adaptation. *American Journal of Agricultural Economics*, 101(1), 172–192.

Kawasaki, K. & Lichtenberg, E. (2014). Econometric analysis of grading standards: The ordered fractional approach. *American Journal of Agricultural Economics*, 96(1), 345–365.

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

Krantz, W., Pierce, D., Goldenson, N., & Cayan, D. (2021). Memorandum on Evaluating Global Climate Models for Studying Regional Climate Change in California. *Scripps Institution of Oceanography, UC San Diego*.

Lobell, D. B., Cahill, K. N., & Field, C. B. (2007). Historical effects of temperature and precipitation on California crop yields. *Climatic change*, 81(2), 187–203.

Lobell, D. B., Torney, A., & Field, C. B. (2011). Climate extremes in California agriculture. *Climatic change*, 109(1), 355–363.

Malikov, E., Miao, R., & Zhang, J. (2020). Distributional and temporal heterogeneity in the climate change effects on us agriculture. *Journal of Environmental Economics and Management*, 104, 102386.

Marklein, A., Elias, E., Nico, P., & Steenwerth, K. (2020). Projected temperature increases may require shifts in the growing season of cool-season crops and the growing locations of warm-season crops. *Science of The Total Environment*, 746, 140918.

McCarl, B. A. & Hertel, T. W. (2018). Climate change as an agricultural economics research topic. *Applied Economic Perspectives and Policy*, 40(1), 60–78.

Mendelsohn, R., Nordhaus, W. D., & Shaw, D. (1994). The impact of global warming on agriculture: a Ricardian analysis. *The American economic review*, (pp. 753–771).

NRCS USDA (2020). Soil Survey Geographic (SSURGO) Database for California. Available at `https://websoilsurvey.nrcs.usda.gov/`.

Ortiz-Bobea, A. (2021). The empirical analysis of climate change impacts and adaptation in agriculture. In *Handbook of agricultural economics*, volume 5 (pp. 3981–4073). Elsevier.

Ortiz-Bobea, A., Knippenberg, E., & Chambers, R. G. (2018). Growing climatic sensitivity of us agriculture linked to technological change and regional specialization. *Science advances*, 4(12), eaat4343.

Ortiz-Bobea, A., Wang, H., Carrillo, C. M., & Ault, T. R. (2019). Unpacking the climatic drivers of us agricultural yields. *Environmental Research Letters*, 14(6), 064003.

Personal correspondence with the processor (2020). 2020-07-06.

Pierce, D., F., K. J., & Cayan, D. (2023). Localized Construction Analogs. Scripps Institution of Oceanography.

PRISM Climate Group, Oregon State University (2020). PRISM Climate Data. Available at `http://prism.oregonstate.edu`.

Ramsey, A. F. & Rejesus, R. M. (2021). Bayesian hierarchical models for measuring varietal improvement in tobacco yield and quality. *Journal of Agricultural and Applied Economics*, 53(4), 563–586.

Ramsey, A. F., Tack, J. B., & Balota, M. (2020). Double or nothing: Impacts of warming on crop quantity, quality, and revenue. *Journal of Agricultural and Resource Economics*, (pp. 1–22).

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

Schmitt, J., Offermann, F., Söder, M., Frühauf, C., & Finger, R. (2022). Extreme weather events cause significant crop yield losses at the farm level in German agriculture. *Food Policy*, 112, 102359.

Shaw, S. B., Mehta, D., & Riha, S. J. (2014). Using simple data experiments to explore the influence of non-temperature controls on maize yields in the mid-West and Great Plains. *Climatic change*, 122(4), 747–755.

Shew, A. M., Tack, J. B., Nalley, L. L., & Chaminuka, P. (2020). Yield reduction under climate warming varies among wheat cultivars in South Africa. *Nature communications*, 11(1), 1–9.

Snyder, R. L. (1985). Hand calculating degree days. *Agricultural and forest meteorology*, 35(1-4), 353–358.

Stone, C. J. & Koo, C.-Y. (1985). Additive splines in statistics. In *Proceedings of the Statistical Computing Section ASA. Washington, DC* (pp. 45–48).

Tack, J., Barkley, A., & Hendricks, N. (2017a). Irrigation offsets wheat yield reductions from warming temperatures. *Environmental Research Letters*, 12(11), 114027.

Tack, J., Barkley, A., & Nalley, L. L. (2015). Effect of warming temperatures on US wheat yields. *Proceedings of the National Academy of Sciences*, 112(22), 6931–6936.

Tack, J., Lingenfelser, J., & Jagadish, S. K. (2017b). Disaggregating sorghum yield reductions under warming scenarios exposes narrow genetic diversity in US breeding programs. *Proceedings of the National Academy of Sciences*, 114(35), 9296–9301.

USDA ERS (2020). Ag and Food Statistics: Charting the Essentials, February 2020. Available at `https://www.ers.usda.gov/webdocs/publications/96957/ap-083.pdf?v=1384.6`.

USDA NASS (2017). 2017 Census of Agriculture. Available at `www.nass.usda.gov/AgCensus`.

USDA NASS (2021a). 2020 California Processing Tomato County Estimates. Available at `https://www.nass.usda.gov/Statistics_by_State/California/Publications/County_Estimates/2021/PTOMCounty_0321.pdf`.

USDA NASS (2021b). 2021 California Processing Tomato Report. Available at `https://www.nass.usda.gov/Statistics_by_State/California/Publications/Specialty_and_Other_Releases/Tomatoes/2021/PTOMInt_0121.pdf`.

Wing, I. S., De Cian, E., & Mistry, M. N. (2021). Global vulnerability of crop yields to climate change. *Journal of Environmental Economics and Management*, 109, 102462.

Wu, X., Yu, L., & Pehrsson, P. R. (2022). Are processed tomato products as nutritious as fresh tomatoes? Scoping review on the effects of industrial processing on nutrients and bioactive compounds in tomatoes. *Advances in Nutrition*, 13(1), 138–151.

1.A Endogeneity in weather

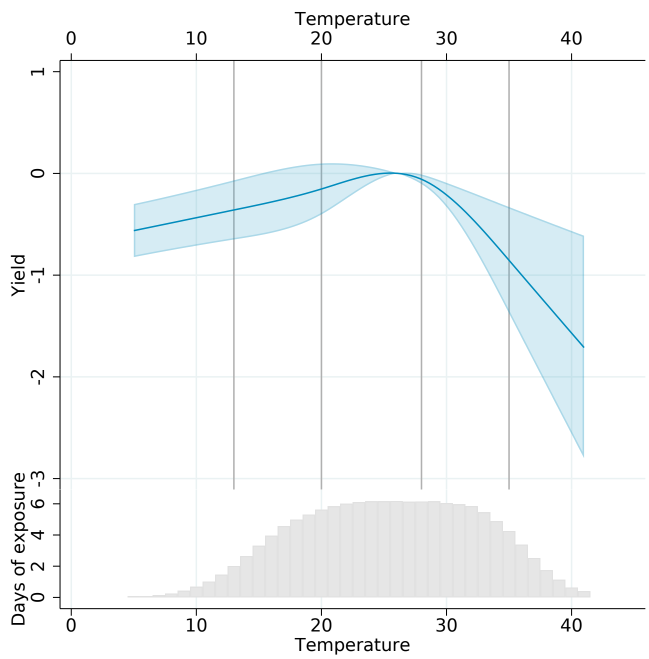

(a) Yield

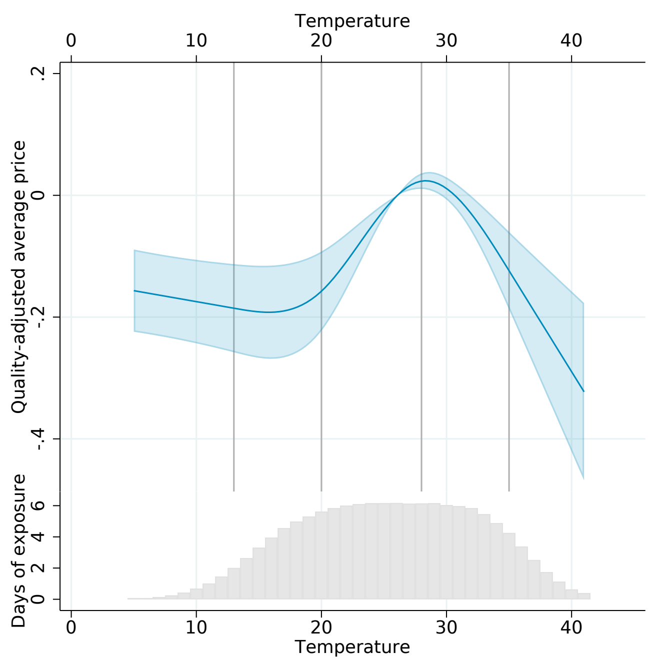

(b) Quality

Figure 1.A.1: Results without a planting date control

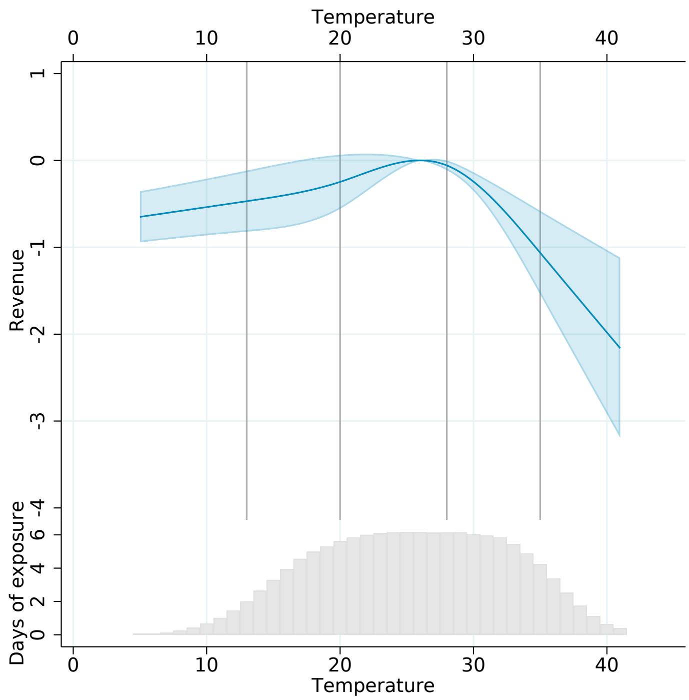

(c) Revenue per acre

Figure 1.A.1: Results without a planting date control

Notes: For each figure, the graph at the top of the frame shows the effect of an additional 24 hours spent at a given temperature interval on the outcome variable relative to 24 hours spent at 26°C. The histogram at the bottom of the frame shows the average exposure to each temperature interval during the growing season across all fields in all years.

1.B Spatial Heteroscedasticity and Autocorrelation Consistent Errors

Our preferred specification clusters standard errors by grower group and county by year to account for the possibility of heteroskedasticity, spatial correlation, and temporal correlation in the errors. An alternative method to correct for possible dependence in standard errors is to estimate spatial heteroscedasticity and autocorrelation consistent (HAC) errors that allow for spatial correlation and serial correlation in panel data (Conley, 1999). Using code from Hsiang (2010), we allow for spatial correlation for field observations that are within 200km (124 miles) of each other. The correlation between observations is assumed to decay linearly with distance.

Results are robust to using spatial HAC errors, however the appropriateness of using this method on an unbalanced panel remains an open question.

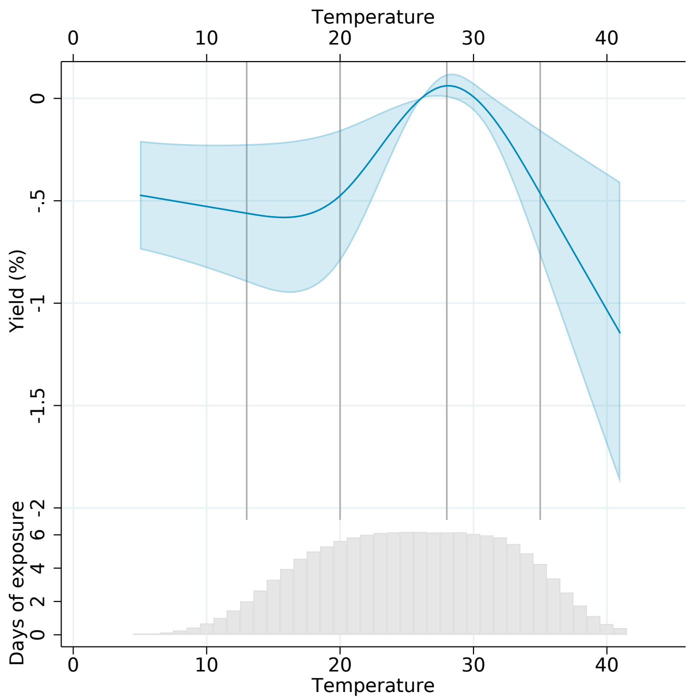

(a) Yield

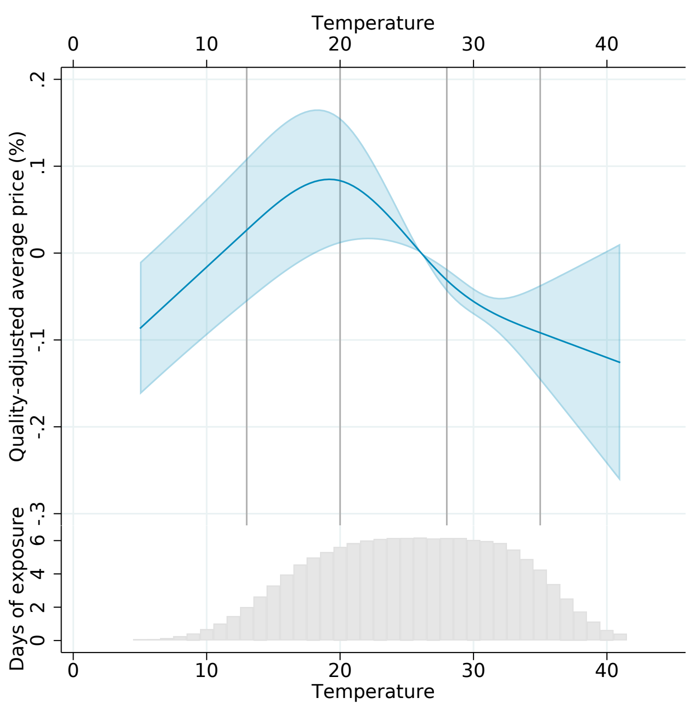

(b) Quality

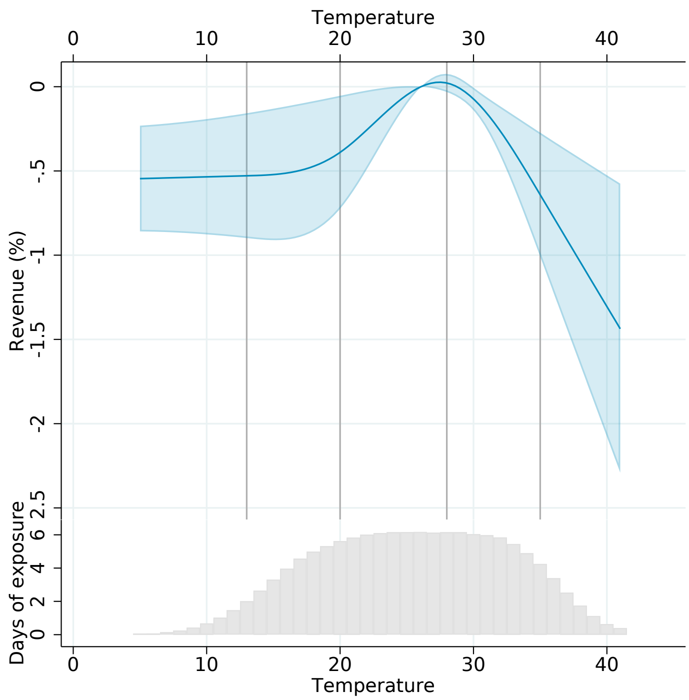

(c) Revenue per acre

Figure 1.B.0: Results using spatial HAC errors

Notes: For each figure, the graph at the top of the frame shows the effect of an additional 24 hours spent at a given temperature interval on the outcome variable relative to 24 hours spent at 26°C. The histogram at the bottom of the frame shows the average exposure to each temperature interval during the growing season across all fields in all years.

1.C Robustness checks

Our preferred specification uses grower fixed effects. Grower fixed effects capture time-invariant characteristics of the growers and characteristics of their respective fields that are both time-invariant and common across fields. One might be concerned that fields differ in ways that are correlated with weather, which would introduce omitted variables bias into our estimation. To alleviate this concern, we estimate Equation 1.15 using field fixed effects instead of grower fixed effects. Some fields do not appear multiple times in our sample because of crop rotation. We drop around 30% of field-year observations because they do not make a field-level panel. The results from this estimation are similar to those from the estimation using grower fixed effects.

We also replace the quadratic year trend with a (a) linear year trend, and (b) county-specific quadratic year trends. The results are robust to the choice of functional form for the time trend.

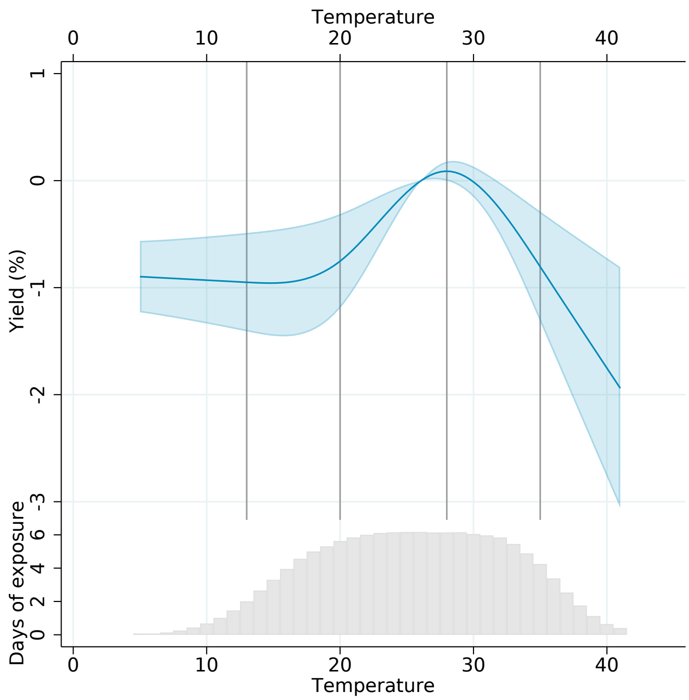

(a) Restricted spline with field fixed effects

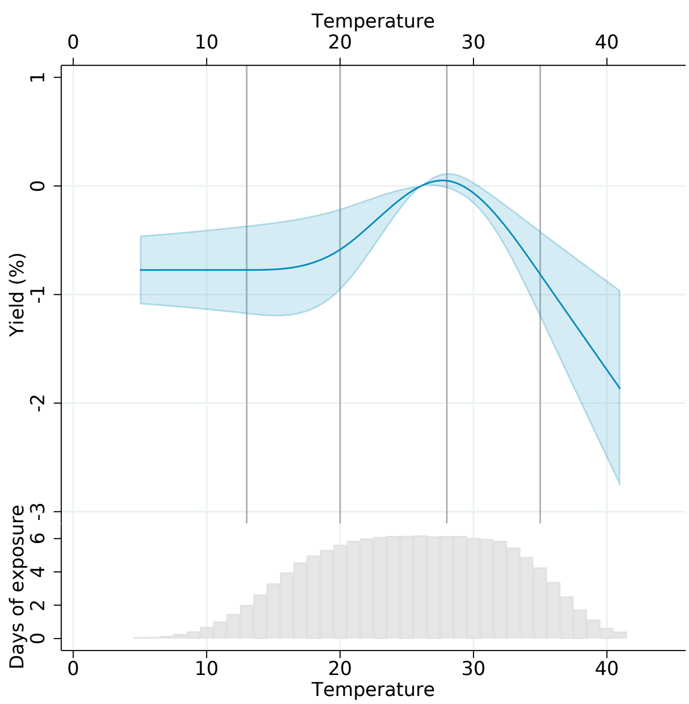

(b) Restricted spline with linear year trend

Figure 1.C.1: Yield, robustness checks

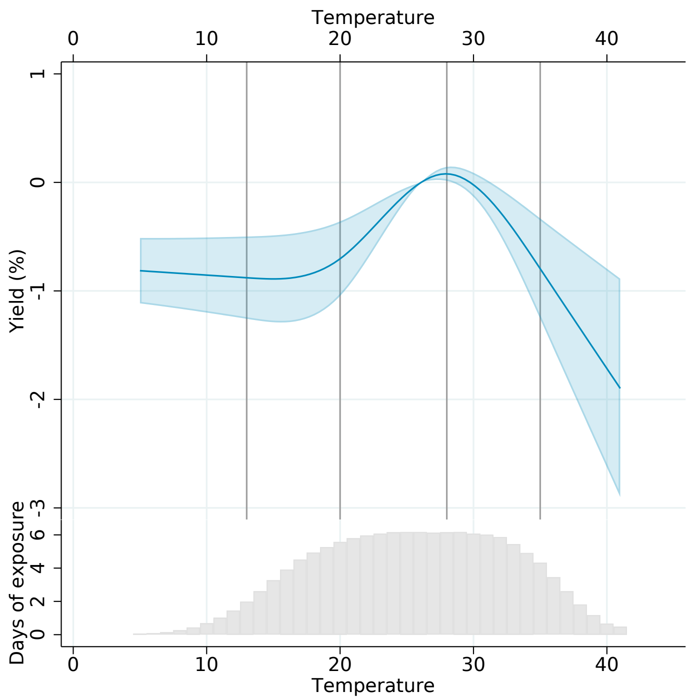

(c) Restricted spline with county-specific quadratic year trends

Figure 1.C.1: Yield, robustness checks

Notes: For each figure, the graph at the top of the frame shows the effect of an additional 24 hours spent at a given temperature interval on the outcome variable relative to 24 hours spent at 26°C. The histogram at the bottom of the frame shows the average exposure to each temperature interval during the growing season across all fields in all years.

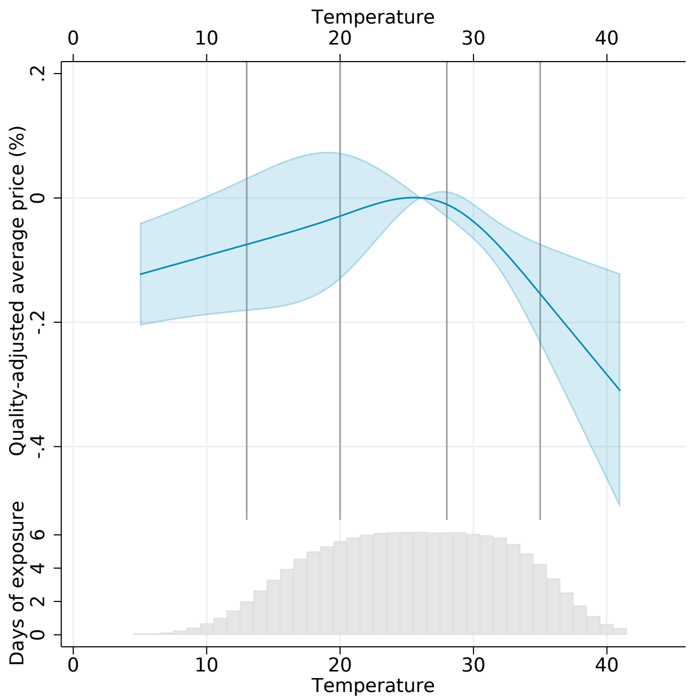

(a) Restricted spline with field fixed effects

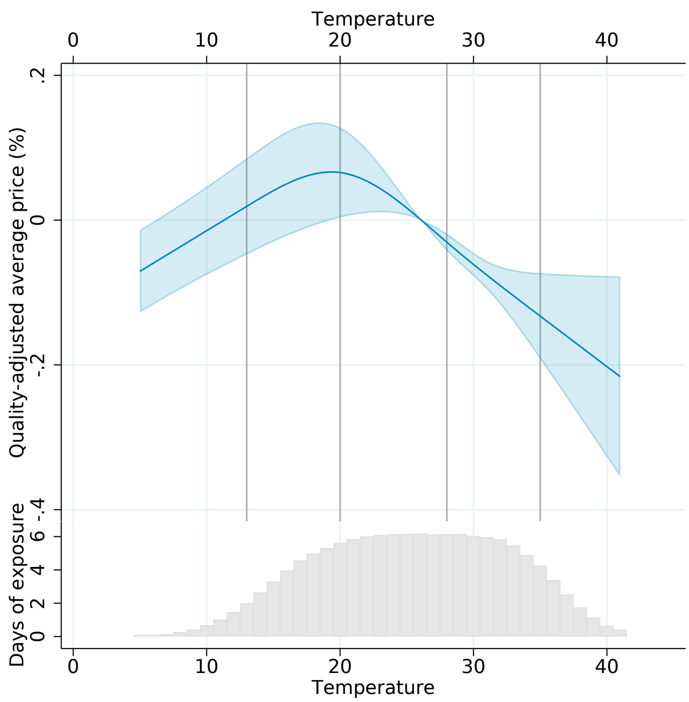

(b) Restricted spline with linear year trend

Figure 1.C.2: Quality, robustness checks

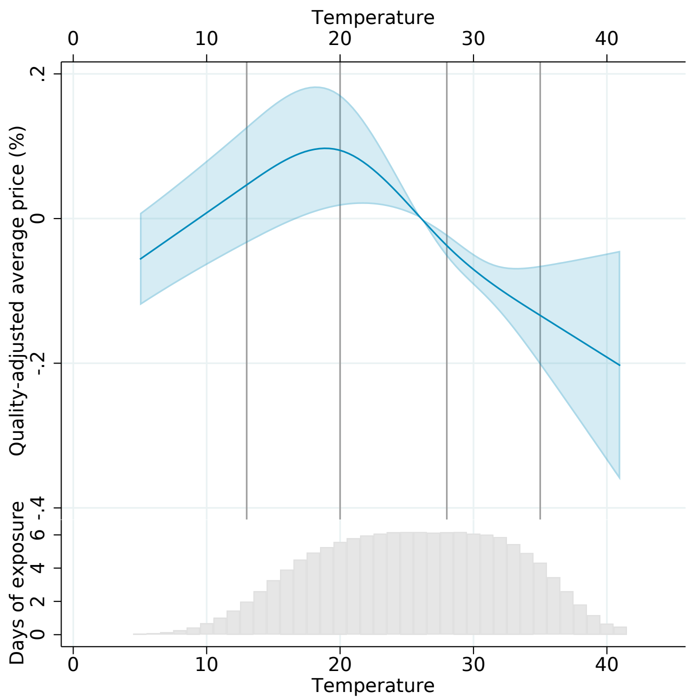

(c) Restricted spline with county-specific quadratic year trends

Figure 1.C.2: Quality, robustness checks

Notes: For each figure, the graph at the top of the frame shows the effect of an additional 24 hours spent at a given temperature interval on the outcome variable relative to 24 hours spent at 26°C. The histogram at the bottom of the frame shows the average exposure to each temperature interval during the growing season across all fields in all years.

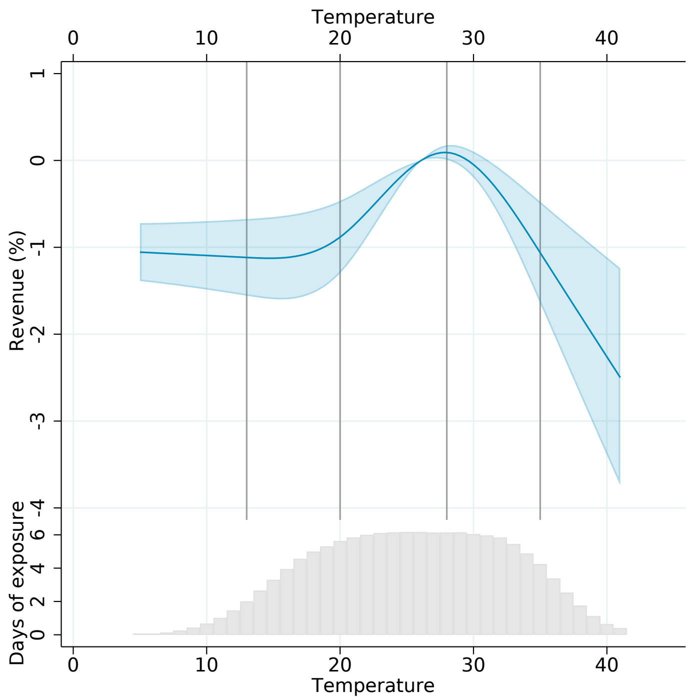

(a) Restricted spline with field fixed effects

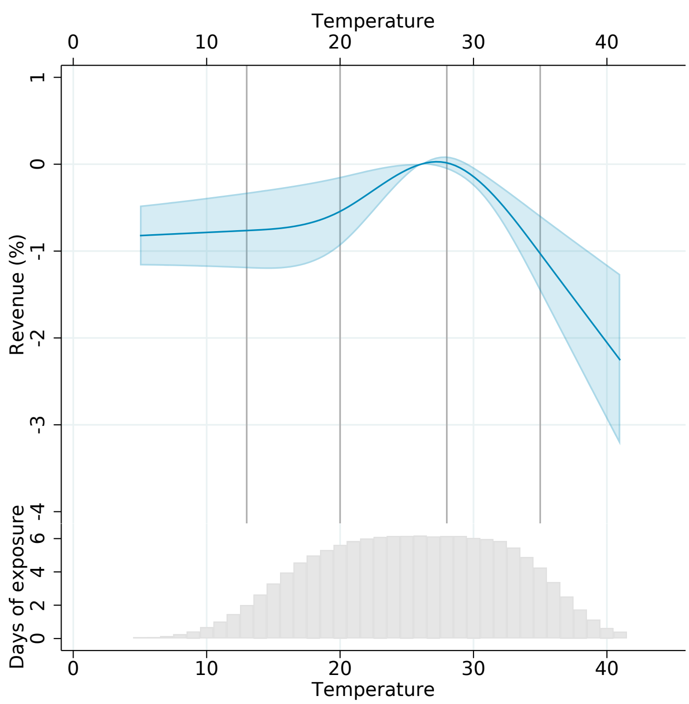

(b) Restricted spline with linear year trend

Figure 1.C.3: Revenue per acre, robustness checks

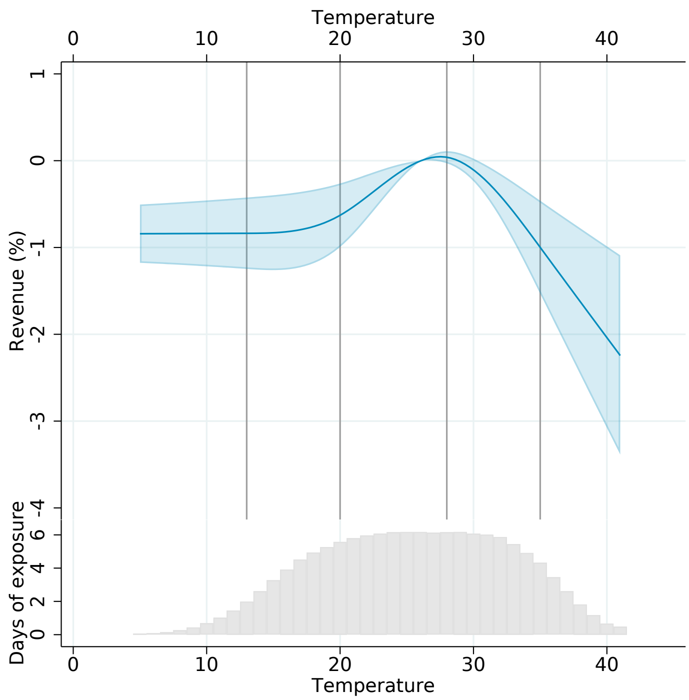

(c) Restricted spline with county-specific quadratic year trends

Figure 1.C.3: Revenue per acre, robustness checks

Notes: For each figure, the graph at the top of the frame shows the effect of an additional 24 hours spent at a given temperature interval on the outcome variable relative to 24 hours spent at 26°C. The histogram at the bottom of the frame shows the average exposure to each temperature interval during the growing season across all fields in all years.

1.D Variation in weather

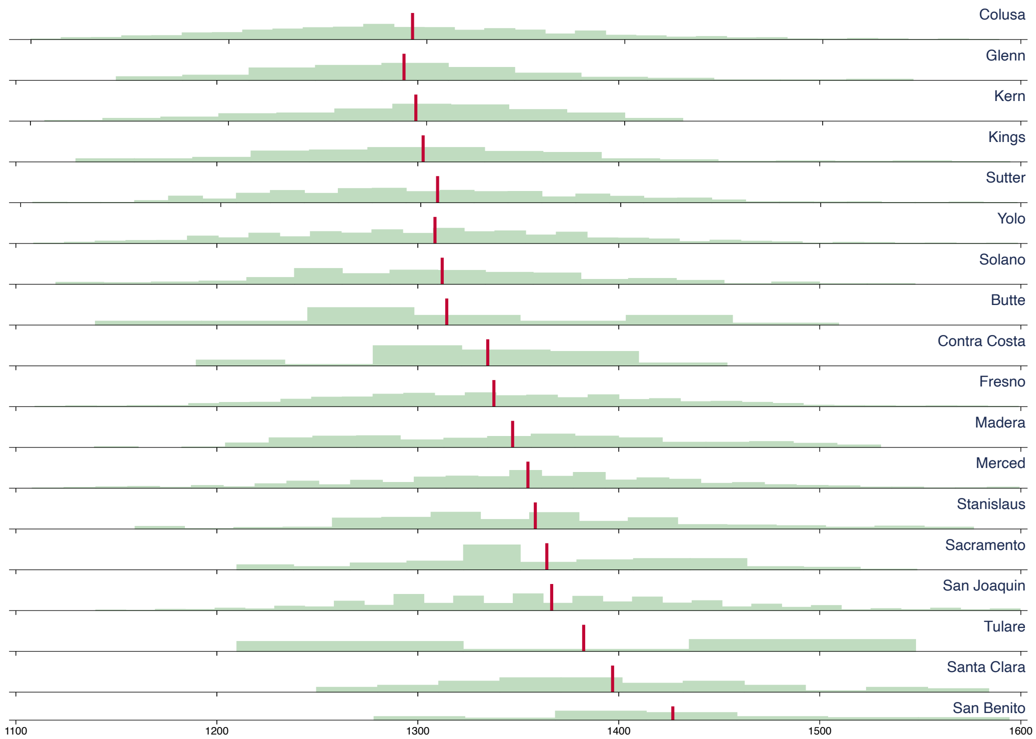

Figure 1.D.1: Histogram of degree days below 10°C by county, with county-average in red

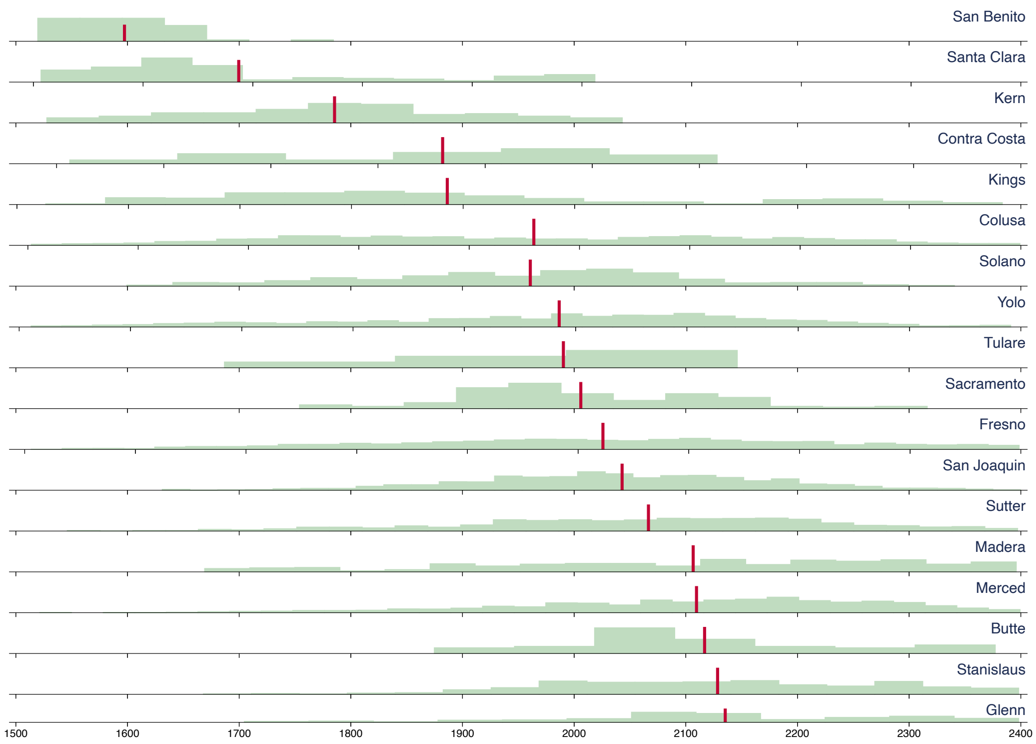

Figure 1.D.2: Histogram of degree days between 10°C and 35°C by county, with county-average in red

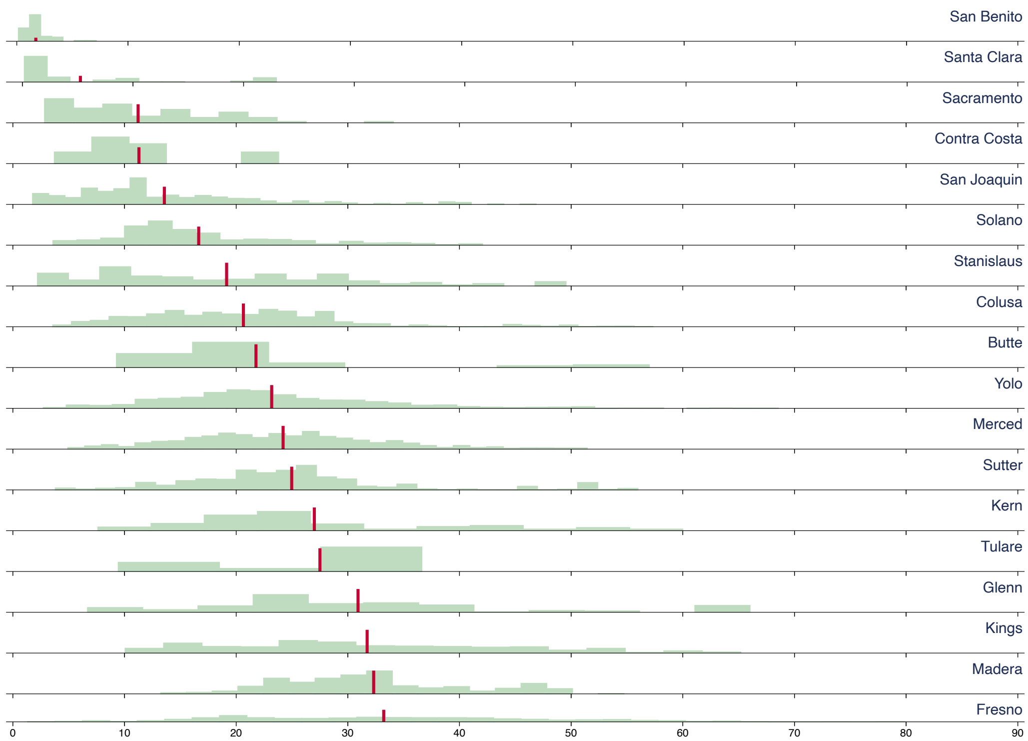

Figure 1.D.3: Histogram of degree days above 35°C by county, with county-average in red

1.E Linear piecewise degree day specification

The piecewise linear degree day model is widely used in the agronomic and agricultural economic literature. It imposes more structure than flexible, semiparametric models. It also relies on the econometrician to correctly choose knot locations where the marginal effects change. However, it is less likely to overfit the data and has been shown in some contexts to perform better out-of-sample (Schlenker & Roberts, 2009).

To implement the piecewise linear degree day functional form, we first need to calculate degree days. Degree days are related to, but different from, temperature exposure. Temperature exposure measures how long is spent in a given temperature interval. Degree days measure how long and by how much temperatures exceed the lower bound of a temperature interval while being truncated at an upper bound (Snyder, 1985). When the temperature interval is small (e.g. 1°C), the difference between the two methods is relatively small because the "how much" dimension is unimportant relative to the "how long" dimension. When the temperature interval is large, as is the case in a piecewise linear model, the difference between the two methods will be large. For example, if we use temperature exposure, we assume that the damage of one day of exposure at 35°C is equal to the damage of one day at 40°C. If we use degree days, we assume the damage of five days at 35°C is equal to the damage of one day at 40°C. The underlying assumption of degree days is that the effect of temperature exposure increases linearly with temperature between the lower and upper bounds.

Degree days can be computed from the temperature exposure vector x_{it}. The expression for calculating degree days between a lower bound of $\underline{h}$ and upper bound of $\overline{h}$ is:

$$DD_{it,[\underline{h},\overline{h}]} = \sum_{j=\underline{h}}^{\overline{h}-1} x_{it,j} \times (j - \underline{h} + 1) \tag{1.12}$$

Next, we choose knot locations. In the first set of results, we use knot locations suggested by the agronomic literature. Mid-range temperatures are ideal for yield and quality outcomes, but these outcomes may be damaged by hot (greater than 35°C) or cool (less than 10°C) temperatures (Hartz et al., 2008). Accordingly, we choose two knots at $\kappa_1 = 10$°C and $\kappa_2 = 35$°C. We estimate degree days using Equation 1.12 for each of the three "segments": below 10°C, between 10°C and 35°C, and above 35°C.

Equation 1.4 can then be modelled as:

$$y_{it} = \beta_0 + \beta_1 DD_{it,(-\infty,10]} + \beta_2 DD_{it,[10,35]} + \beta_3 DD_{it,[35,\infty)} + \delta z_{it} + \alpha_{g(i)} + \psi(t) + \epsilon_{it} \tag{1.13}$$

In the second set of results, we use knot locations that correspond to the turning points in the restricted cubic spline estimates.

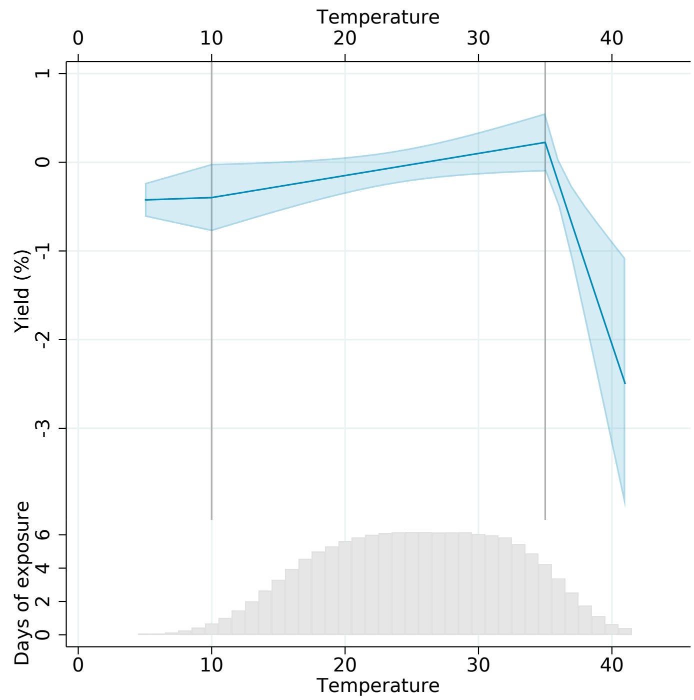

(a) Yield

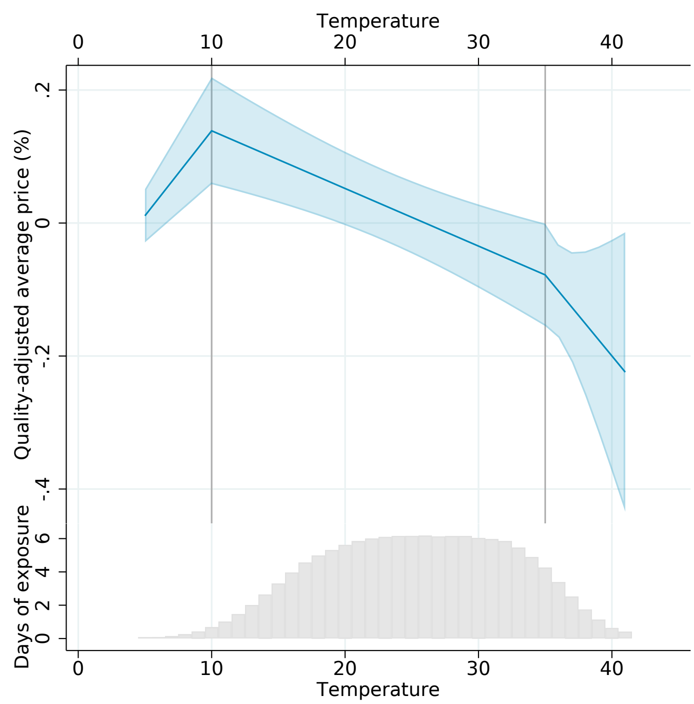

(b) Quality

Figure 1.E.1: Piecewise linear degree day results, agronomic knots

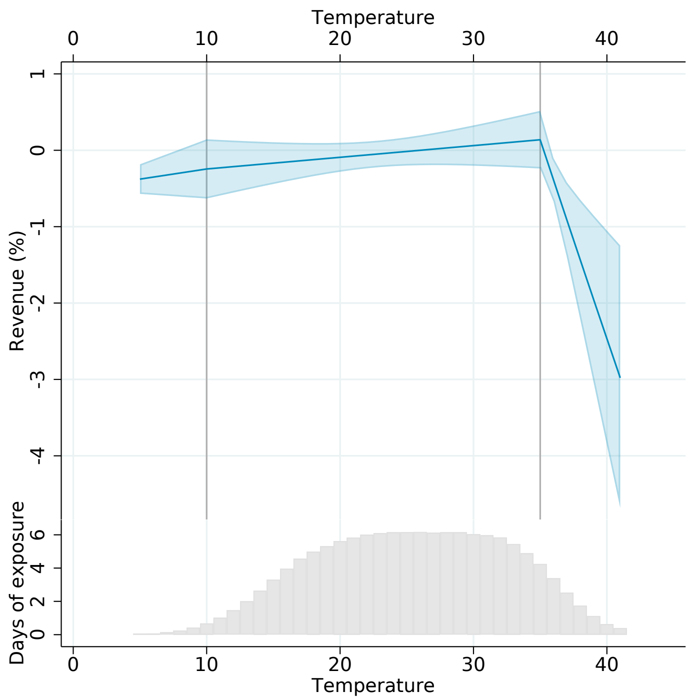

(c) Revenue per acre

Figure 1.E.1: Piecewise linear degree day results, agronomic knots

Notes: For each figure, the graph at the top of the frame shows the effect of an additional 24 hours spent at a given temperature interval on the outcome variable relative to 24 hours spent at 26°C. The histogram at the bottom of the frame shows the average exposure to each temperature interval during the growing season across all fields in all years.

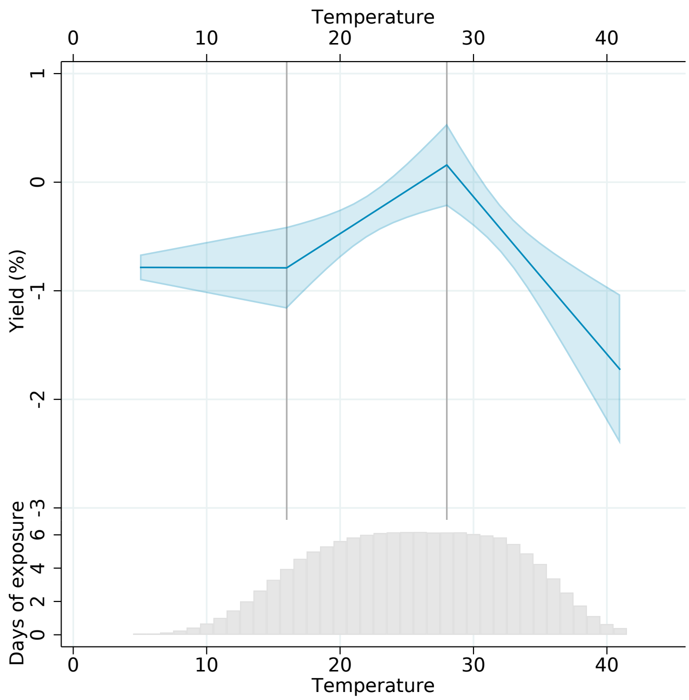

(a) Yield

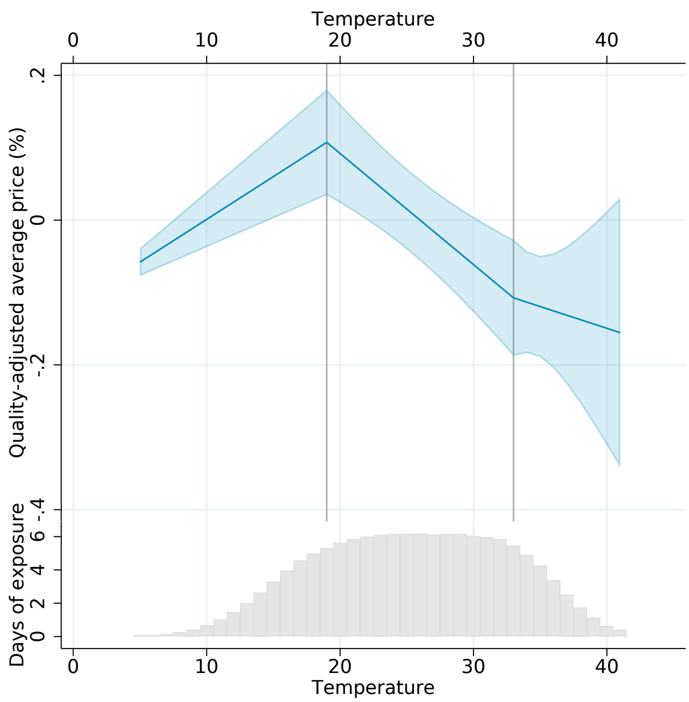

(b) Quality

Figure 1.E.2: Piecewise linear degree day results, spline knots

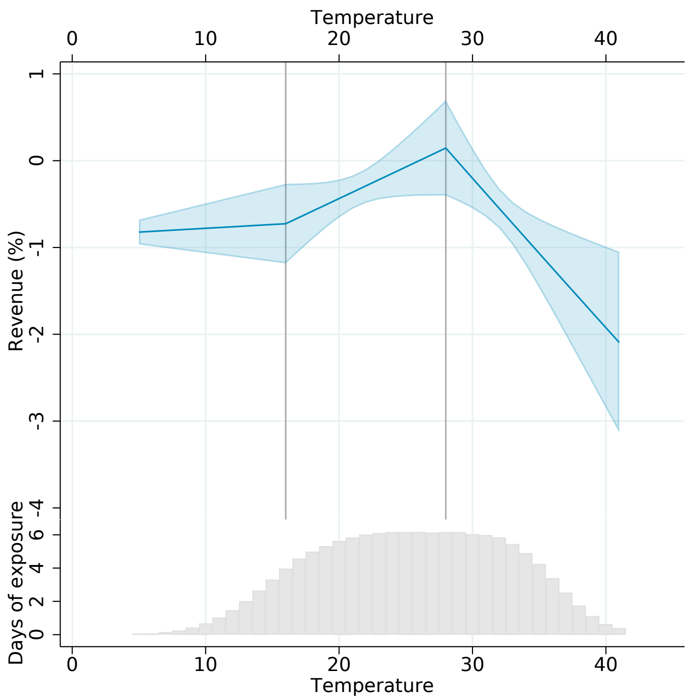

(c) Revenue per acre

Figure 1.E.2: Piecewise linear degree day results, spline knots

Notes: For each figure, the graph at the top of the frame shows the effect of an additional 24 hours spent at a given temperature interval on the outcome variable relative to 24 hours spent at 26°C. The histogram at the bottom of the frame shows the average exposure to each temperature interval during the growing season across all fields in all years.

1.F Details on restricted cubic spline specification

The B matrix for a cubic spline combines the B matrix for the cubic polynomial and a $K \times J$ matrix Z. First, define $1 \times J$ temperature vector $W = \begin{pmatrix} 5 & 6 & 7 & \dots & 41 \end{pmatrix}$. Z is the resulting matrix after applying the cubic spline basis function to W with K knots at $\kappa_1, \dots, \kappa_K$.

$$\begin{aligned} B_{cubic.spline} &= \begin{pmatrix} B_{cubic.polynomial} & Z \end{pmatrix} \\ &= \begin{pmatrix} 5 & 25 & 125 & (5-\kappa_1)_+^3 & (5-\kappa_2)_+^3 & \dots & (5-\kappa_K)_+^3 \\ 6 & 36 & 216 & (6-\kappa_1)_+^3 & (6-\kappa_2)_+^3 & \dots & (6-\kappa_K)_+^3 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \\ 41 & 41^2 & 41^3 & (41-\kappa_1)_+^3 & (41-\kappa_2)_+^3 & \dots & (41-\kappa_K)_+^3 \end{pmatrix} \end{aligned} \tag{1.14}$$

Stone & Koo (1985) use the linearity constraints to develop a restricted cubic spline function. Using this function, a restricted cubic spline with K knots requires the estimation of only $K-1$ parameters on temperature (as opposed to $K+3$ parameters for the cubic spline). Equation 1.15 is the restricted spline function that we apply to W for $i = 1, 2, \dots, K-2$ [8].

$$\begin{aligned} V_1 &= W \\ V_{i+1} &= \frac{(W-\kappa_i)_+^3 - (\kappa_K - \kappa_{K-1})^{-1}\{(W-\kappa_{K-1})_+^3(\kappa_K - \kappa_i) - (W-\kappa_K)_+^3(\kappa_{K-1} - \kappa_i)\}}{(\kappa_K - \kappa_1)^2} \\ &\quad \text{for } i = 1, 2, \dots, K-2 \end{aligned} \tag{1.15}$$

Together, W and V_2 to V_{K-1} make up the B matrix for the restricted cubic spline, where $V_{i,j}$ is the j-th element of the V_i vector.

[8] For more details, see the Stata manual for the function mkspline

$$
\begin{aligned}
B_{rest.cubic.spline} &= \begin{pmatrix} W & V_2 & \dots & V_{K-1} \end{pmatrix} \\
&= \begin{pmatrix} 5 & V_{2,0} & \dots & V_{K-1,0} \\ 6 & V_{2,1} & \dots & V_{K-1,1} \\ 7 & V_{2,2} & \dots & V_{K-1,2} \\ \vdots & \vdots & \ddots & \vdots \\ 41 & V_{2,45} & \dots & V_{K-1,45} \end{pmatrix}
\end{aligned} \tag{1.16}
$$

1.G Results for individual quality attributes

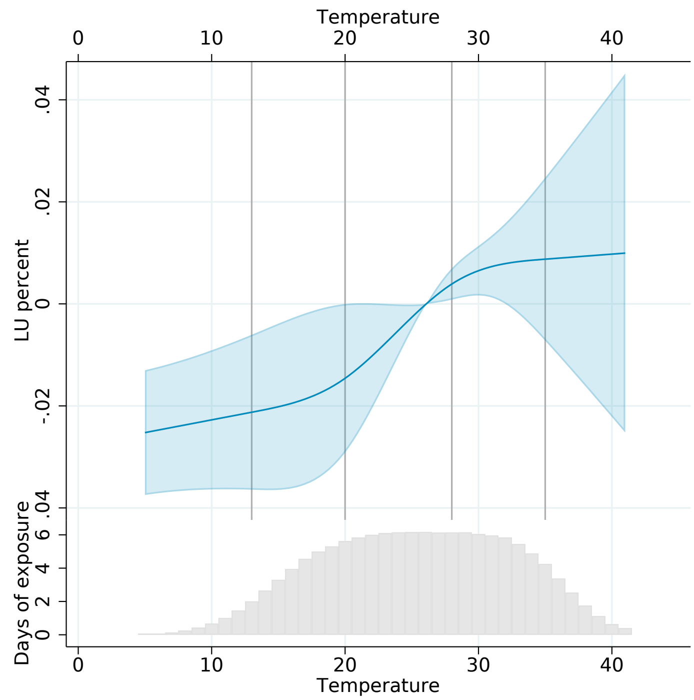

(a) Limited use (LU)

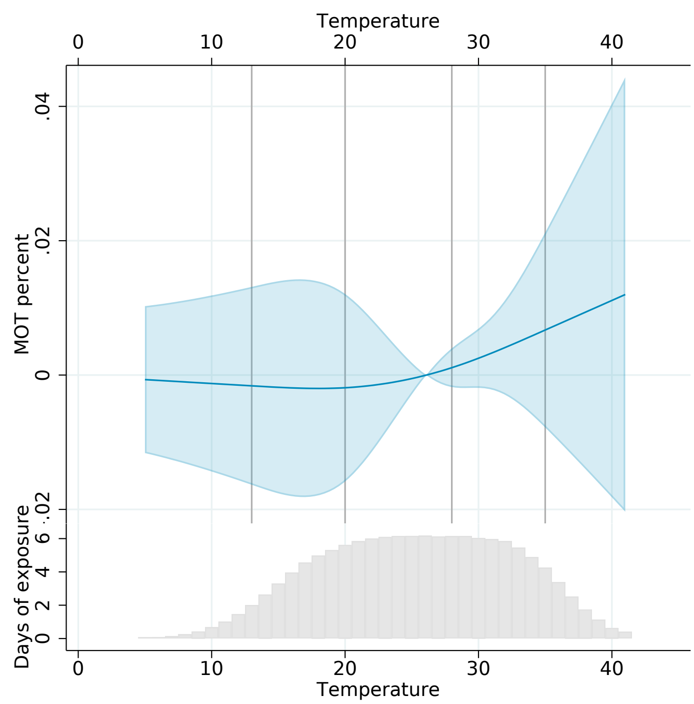

(b) Material other than tomatoes (MOT)

Figure 1.G.1: Restricted cubic spline results for individual quality defects

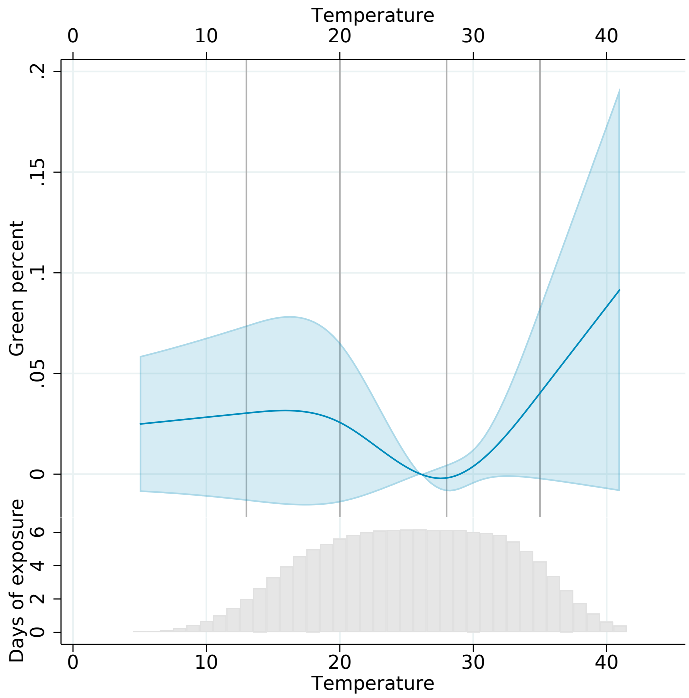

(c) Green

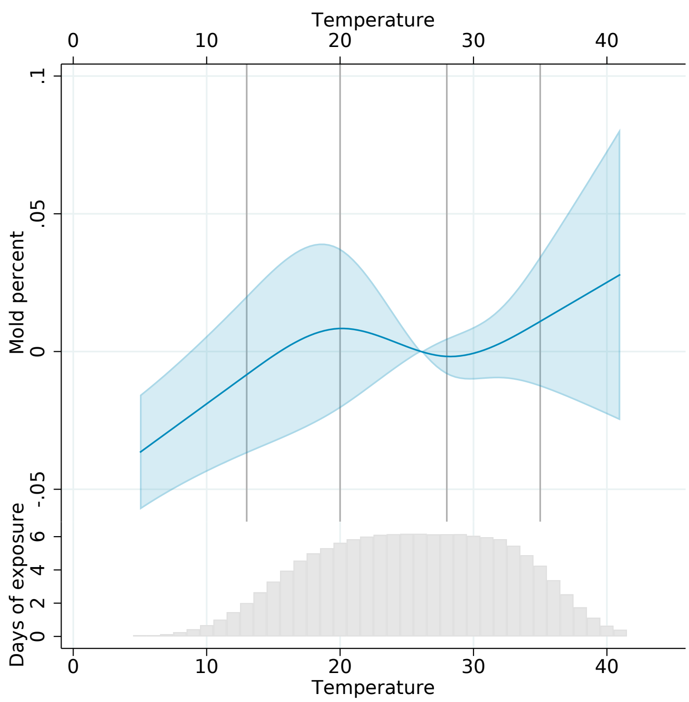

(d) Mold

Figure 1.G.1: Restricted cubic spline results for individual quality defects

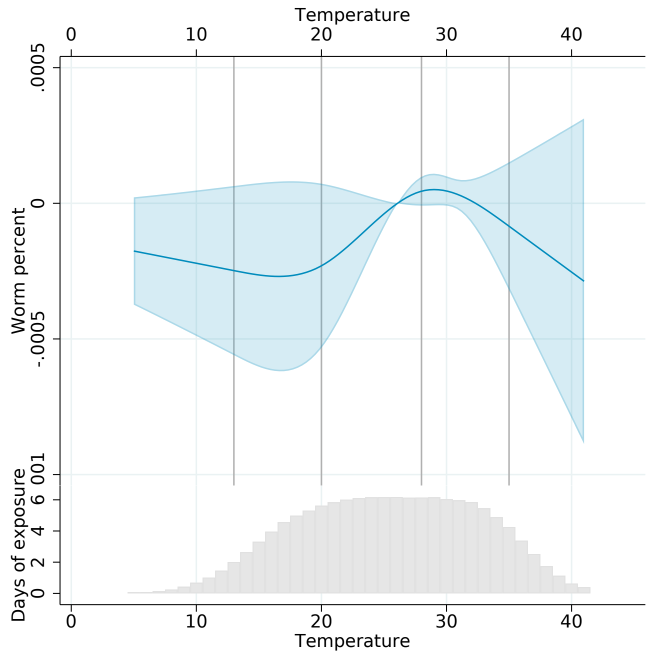

(e) Worm

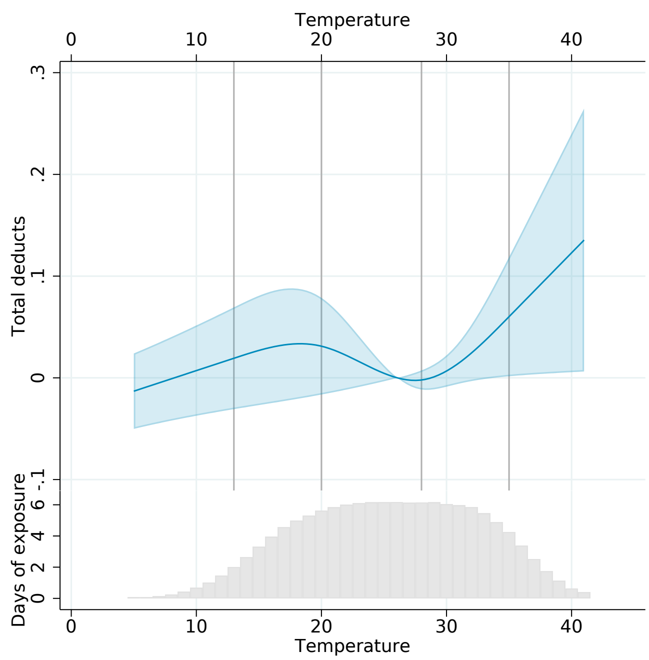

(f) Total defects

Figure 1.G.1: Restricted cubic spline results for individual quality defects

Notes: For each figure, the graph at the top of the frame shows the effect of an additional 24 hours spent at a given temperature interval on the outcome variable relative to 24 hours spent at 26°C. The histogram at the bottom of the frame shows the average exposure to each temperature interval during the growing season across all fields in all years.

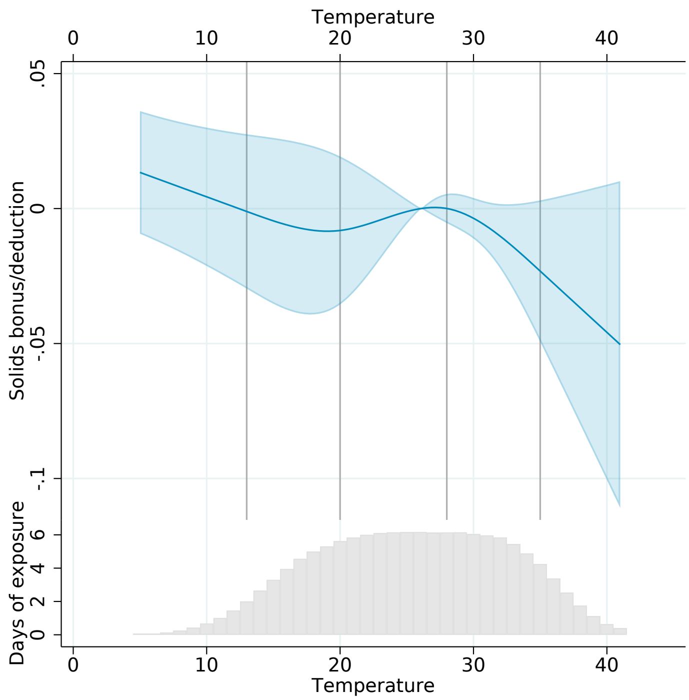

(a) Solids bonus

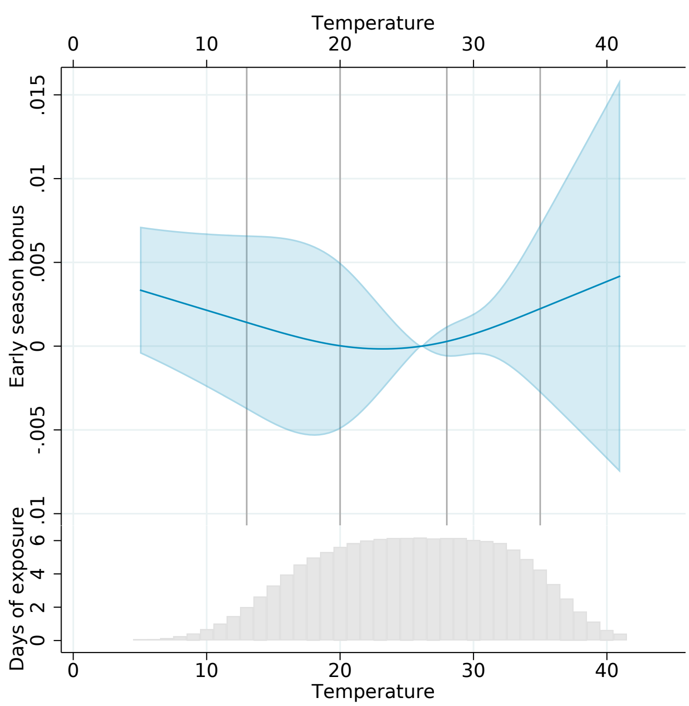

(b) Early season bonus

Figure 1.G.2: Restricted cubic spline results for individual quality bonuses

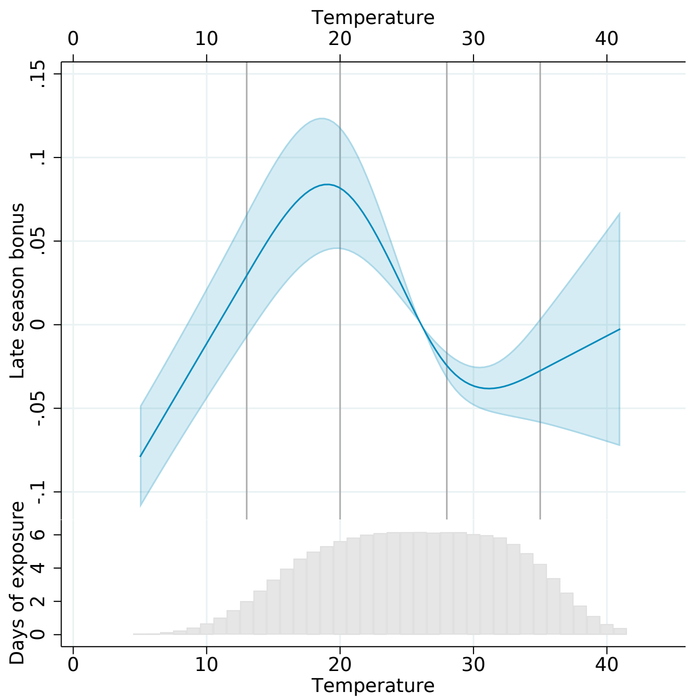

(c) Late season bonus

Figure 1.G.2: Restricted cubic spline results for individual quality bonuses

Notes: For each figure, the graph at the top of the frame shows the effect of an additional 24 hours spent at a given temperature interval on the outcome variable relative to 24 hours spent at 26°C. The histogram at the bottom of the frame shows the average exposure to each temperature interval during the growing season across all fields in all years.

1.H Estimated effect of control variables

Table 1.H.1: Coefficient estimates on control variables

	(1) log yield	(2) log quality	(3) log revenue
Precipitation	-0.0008*	-0.0001	-0.0009***
	(0.0003)	(0.0001)	(0.0003)
Soil type			
Eolian	0.0558	-0.0042	0.0581
	(0.1151)	(0.0070)	(0.1152)
Organic material	-0.0478***	0.0057*	-0.0384**
	(0.0136)	(0.0027)	(0.0136)
Lacustrine	-0.0069	0.0445***	0.0629
	(0.1485)	(0.0065)	(0.1475)
Residuum	0.0168	-0.0043	0.0063
	(0.0519)	(0.0050)	(0.0544)
Varietal attributes			
Extended field storage variety	-0.0063	0.0221***	0.0313**
	(0.0092)	(0.0029)	(0.0111)
Tomato spotted wilt resistant	0.0047	0.0015	0.0036
	(0.0081)	(0.0024)	(0.0099)
High solids	0.0221	-0.0194	-0.0086
	(0.0485)	(0.0127)	(0.0498)
Fusarium Wilt resistant	-0.0161	-0.0075*	-0.0335*
	(0.0144)	(0.0037)	(0.0167)
Powdery Mildew resistant	-0.0457	0.0204***	-0.0044
	(0.0245)	(0.0049)	(0.0258)
High yield	0.0115	0.0283*	0.0552
	(0.0497)	(0.0133)	(0.0522)
Fusarium Crown Rot resistant	-0.0149	0.0008	-0.0163
	(0.0362)	(0.0110)	(0.0507)
Bacterial Spot resistant	-0.1381*	0.0136	-0.1279*
	(0.0544)	(0.0116)	(0.0583)
Early	-0.0512**	0.0221***	-0.0242
	(0.0167)	(0.0037)	(0.0178)
Thick	-0.0240*	0.0012	-0.0204
	(0.0113)	(0.0027)	(0.0126)
Thin	-0.0530***	0.0051	-0.0471**
	(0.0140)	(0.0035)	(0.0141)
Pear-shaped	-0.0940*	0.0041	-0.0921*
	(0.0433)	(0.0065)	(0.0431)
Irrigation technology			
Drip irrigation	0.0495*	-0.0054	0.0496*
	(0.0207)	(0.0034)	(0.0210)
Furrow irrigation	-0.0478*	0.0084	-0.0341
	(0.0231)	(0.0053)	(0.0238)
Sprinkler irrigation	-0.1102***	0.0081	-0.1148***
	(0.0199)	(0.0083)	(0.0315)
Harvesting early	-0.0017*	0.0001	-0.0018*
	(0.0008)	(0.0002)	(0.0008)
Year trend	0.0040	0.0008	0.0037
	(0.0137)	(0.0036)	(0.0166)
Year trend sqrd	-0.0006	0.0000	-0.0003
	(0.0012)	(0.0003)	(0.0013)

Essay 2

Doing More with Less: Margins of Response to Water Scarcity in Irrigated Agriculture in California

2.1 Introduction

Irrigation is used to produce 40 percent of global calories (FAO, 2023) and is a key adaptive response to warming temperatures under climate change. But water supplies are predicted to become more variable and scarce in many important growing regions of the world, including California's $50 billion agricultural zone (Arias et al., 2021). During periods of water scarcity, a key margin of response to climate–applying more water–is constrained. The future of agriculture will depend on farmers' ability to respond to the challenges of a changing climate and to "do more with less."

Existing work finds little evidence that U.S. farmers rainfed growing staple crops are adapting to a changing climate (Schlenker & Roberts (2009); Burke & Emerick (2016); Moore et al. (2017)). Focusing on outcome variables that implicitly embed producers' response to weather and climate, such as yield or profit, allows researchers to estimate the net effect of adaptive behavior yet it does not allow them to unpack the processes by which farmers respond. Shedding light on current and future adaptation actions requires opening the black box of farmer decision-making. While recent work has made progress to fill this gap (Taraz (2017); Jagnani et al. (2021); Aragón et al. (2021)), many margins and settings remain understudied.

In this paper, I analyze if and how farmers respond to water scarcity along the extensive and intensive margins, and uncover the mechanisms through which farmers save water. I use a fixed effects model to causally identify the effect of water availability on farmer behavior on 3,300 fields in California between 2011 and 2021. California is an ideal setting to study this question. First, it is a highly productive yet water constrained agricultural region that relies on irrigation to produce crops. Second, California's system of surface water allocations introduces variation in water supplies that are plausibly exogenous to growers' cropping and input decisions. I bring together public data on crop choice and water supplies and novel, spatially-detailed data on growing practices and irrigation. These data capture extensive and intensive margin decisions made by individual farmers.

This paper adds to the literature on adaptation to water scarcity in agriculture. One well-studied response is along the extensive margin: changes in land use. Previous work has found fallowing increases during periods of water scarcity (or higher water prices) (Moore et al. (1994); Schoengold et al. (2006); Hendricks & Peterson (2012); Manning et al. (2017); Hagerty (2021)). But fallowing is a blunt instrument–it is costly to growers who forgo income when land is left idle. A reliance on fallow implies that long-term water scarcity will force widespread fallowing and potential exit from agriculture. A focus on the extensive margin ignores other interventions producers can employ to use available water resources more efficiently.

Earlier work is inconclusive about the role of intensive margin adjustments, in other words growing practices that influence water use per acre. Using econometric analysis (Schoengold et al. (2006); Hendricks & Peterson (2012); Drysdale & Hendricks (2018)) and programming models (Cortignani & Severini (2009); Medellín-Azuara et al. (2012); Graveline &

Mérel (2014); Sapino et al. (2022)), some papers find that growers meaningfully reduce their water intensity. But other evidence suggests that California growers do not adjust along the intensive margin in response to changes in surface water availability (Hagerty, 2021) or groundwater pumping costs (Burlig et al., 2021). I bring new, remotely-sensed and spatially-detailed evapotranspiration data to directly test whether California growers adjust along the intensive margin.

I also unpack the mechanisms by which growers save water. Growing practices are rarely systematically observed but I am able to fill this gap using detailed proprietary data on every field contracted with a large tomato processor operating in California's $1 billion processing tomato industry. These field-level panel data are collected for the purposes of contracting and payment and does not rely on grower self-reports. Some growing practice decisions are made jointly by growers and the processor, however I control for the processor's preferences and focus on how growers respond. These data provide insights into farmer behavior at a level of detail and accuracy that cannot be obtained in publicly available data or survey data.

I first estimate the effect of water availability on the extensive margin using satellite imagery of land use between 2011 and 2021 from USDA NASS (2022). Consistent with earlier work, I find that growers fallow a greater proportion of their acreage during periods of water scarcity. Water scarcity also affects crop choice: I find that growers are more likely to maintain acreage in low-water crops and reduce acreage of high-water crops.

In addition to changing whether they plant and what they plant, farmers also change *how* they produce. On average, growers plant 1.6 days earlier during a water scarce year (a one-standard deviation decrease in surface water supply). Planting earlier exposes plants to cooler temperatures on average which in turn reduces the plant's water demands. Growers are more likely to plant varieties that require fewer days to reach maturity and are more likely to use fields with drip irrigation than fields with less efficient irrigation systems.

I also find that water access priority drives how growers choose to conserve water. During a water-scarce year, growers with high-priority water access experience smaller declines in their water supply and adjust by planting low-water crops in place of high-water crops. By comparison, growers with low-priority access experience large declines in their water availability during a water-scarce year. These growers fallow acreage that would otherwise be planted to high-water crops. In addition to extensive margin adjustments, they also adjust their growing practices to conserve water on their remaining crops. Growers with low-priority access are more likely to plant fast-maturing varieties of processing tomatoes earlier in the season. However, these practices come at the cost of lower yield, which leads to less revenue per acre.

Using a back-of-the-envelope calculation, I find that intensive margin response to water scarcity cost California processing tomato growers $20 million in lost revenue in a water-scarce year. While adapting growing practices may cause a loss of revenue, it helps farmers avoid fallowing, which is significantly more costly.

2.2 The Setting

Details on processing tomatoes can be found in 1.2. Here, I expand on agriculture and water in California more broadly.

California is one of the most agriculturally productive regions in the world yielding $50 billion in agriculture (CDFA, 2021), of which more than 70 percent is specialty crop output (USDA NASS, 2017). With very little rainfall during key growing months, virtually all agricultural crops in California are irrigated.

In California, water used for irrigation comes from two main sources. First, surface water comes from snow that falls on nearby mountain ranges and then melts and flows through rivers and streams. Around 40 percent of surface water is diverted to farmland via a system of reservoirs and canals. The second source is groundwater pumped from underground aquifers. Groundwater contributes around 40 percent of California's water supply in a typical year and up to 60 percent during a drought year (DWR, 2022).

Most farmers in California source surface water from an irrigation district–a term I use to designate any organization that acquires and distributes surface water to farmers within its jurisdiction. Irrigation districts choose how to distribute water to their farmers. In keeping with historical convention, many irrigation districts divide surface water acquisitions evenly across farmland in the district and allow farmers to make decisions on how to use water allocated to them.

Irrigation districts in California may hold long-term water rights called "appropriative rights:" the right to divert and use a specific quantity of surface water from rivers or streams. Many irrigation districts also hold contracts with state and federal water projects. Districts in the Sacramento Valley and San Joaquin Valley may hold contracts with the California state government's State Water Project (SWP) and/or the federal government's Central Valley Project (CVP). Under each project, a district's contract specifies a water entitlement and contract type. However, a district is not guaranteed to receive its full entitlement. Each year, contract-holders are permitted to use a proportion of their entitlement or "allocation" that ranges from 0 to 100 percent. Allocations are set by the relevant government agencies and vary year-to-year based on environmental conditions affecting surface water availability, including winter precipitation stored as snowpack and prevailing reservoir levels.

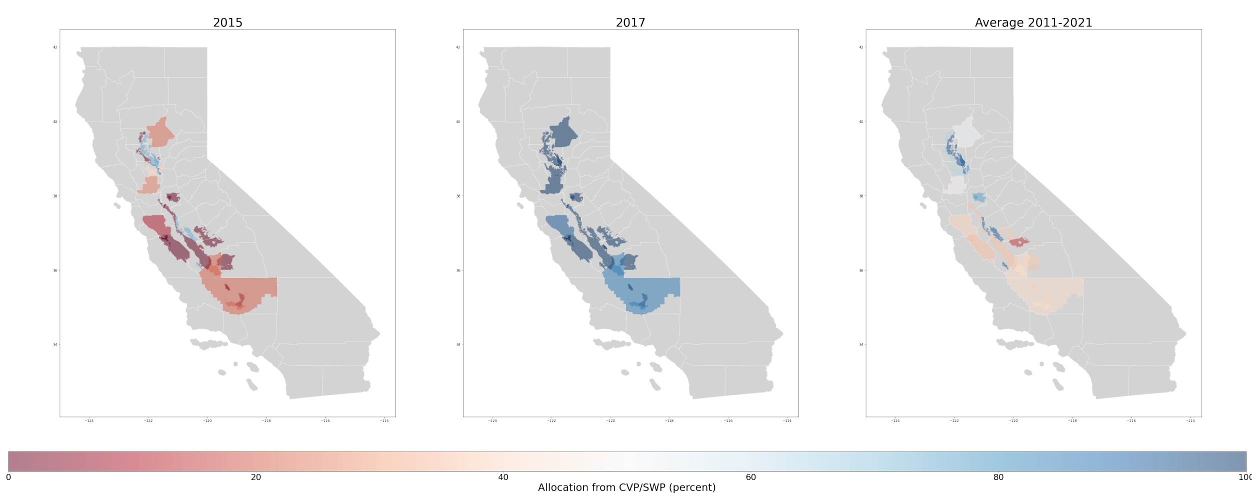

Figure 2.1: Spatial and temporal variation in allocations from water projects

Figure 2.1 shows how allocations from water projects vary both temporally and spatially. In years of ample water supplies, such as 2017, all districts are allocated (close to) 100 percent of their entitlement. In water scarce years, a district is assigned a proportion of its entitlement that depends on the contract type held by the district and degree of water scarcity. For example, during the drought in 2015, districts with high-priority contracts experienced smaller declines in their allocation compared to districts with low-priority contracts.

Information about surface water availability is salient to farmers when they make planting decisions in the spring. Preliminary allocations are released by government agencies in winter and revisions, if any, are typically small[1]. Therefore, the environmental conditions that affect allocations are largely known before the growing season begins.

Climate change is projected to affect future water availability in many regions around the world (Arias et al., 2021). In California, climate change is predicted not only to reduce snow accumulation but also to cause earlier and faster melting of the snowpack, leading to a decrease in summer surface water availability (Hayhoe et al., 2004). Declines in surface water may be partially offset by groundwater pumping. However, changes in precipitation patterns may also affect groundwater recharge. Excessive pumping drops the groundwater levels and increases the cost of water extraction. In extreme cases, wells run dry because groundwater levels fall below the installed capacity of existing wells, for example during the 2014 drought (Howitt et al., 2014). While environmental conditions strongly influence water supply, 80 percent of US irrigation districts do not have a formal plan for responding to future water scarcity (Wallander et al., 2022).

The State of California introduced a legislation package with the long-term goal of sustainably managing groundwater by 2042 called the Sustainable Groundwater Management Act (2014). Groundwater sustainability agencies for medium and high-priority groundwater basins are obligated to prepare and implement groundwater management plans. Groundwater management plans were required by January 2022 for the Department of Water Resources to evaluate. Importantly for this study, the plans were not binding during the period of anal-

[1] From 2011 to 2021, the largest downward revision in an irrigation district's allocation from the Central Valley Project was in 2013 when the allocation for Agricultural Contractors South of the Delta declined from the 25% (announced on February 25$^{\text{th}}$) to 20% (revised on March 22$^{\text{nd}}$).

ysis from 2011 to 2021.

2.3 Data

Estimating the effect of water scarcity on growers' decisions requires spatially-detailed data on water availability and potential margins of response. I collect public data on water rights and allocations, weather, and land use. Comparable data on the intensive margin are scant. To fill this gap, I use novel proprietary data on processing tomato growers and satellite-derived estimates of irrigation. The proprietary data include the latitude and longitude of each field centroid which I use to match fields to observations of estimated irrigation, water supplies, and land use during years when fields were not contracted with the processor. Figure 2.2 summarizes the key datasets, their sources, and linkages across datasets. Summary statistics for the merged data are in Tables 2.1 and 2.2.

2.3.1 Explanatory Variable: Surface Water Supply

Irrigation districts play a key role in acquiring and delivering surface water to fields within their jurisdiction. Most districts acquire surface water from two sources: appropriative rights and project contracts from state and federal water projects. I match each district to the quantity of water it is permitted to divert from rivers and streams under its appropriative right using data from the State Water Resources Control Board's Electronic Water Rights Information Management System.

Districts often hold contracts with the State Water Project (SWP) and/or Central Valley Project (CVP). These contracts establish a district's fixed maximum entitlement from which a proportion or "allocation" will be available to the district in a given year. In a year of ample water supplies, districts are allocated close to 100 percent of their entitlement. In a water scarce year, districts are allocated a proportion of their entitlement that depends on the district's contract type and overall degree of water scarcity. Allocations vary from 0 percent to 100 percent across years and contract types in the sample. I collected data on entitlements, contract types, and allocations from the California Department of Water Resources and the U.S. Bureau of Reclamation.

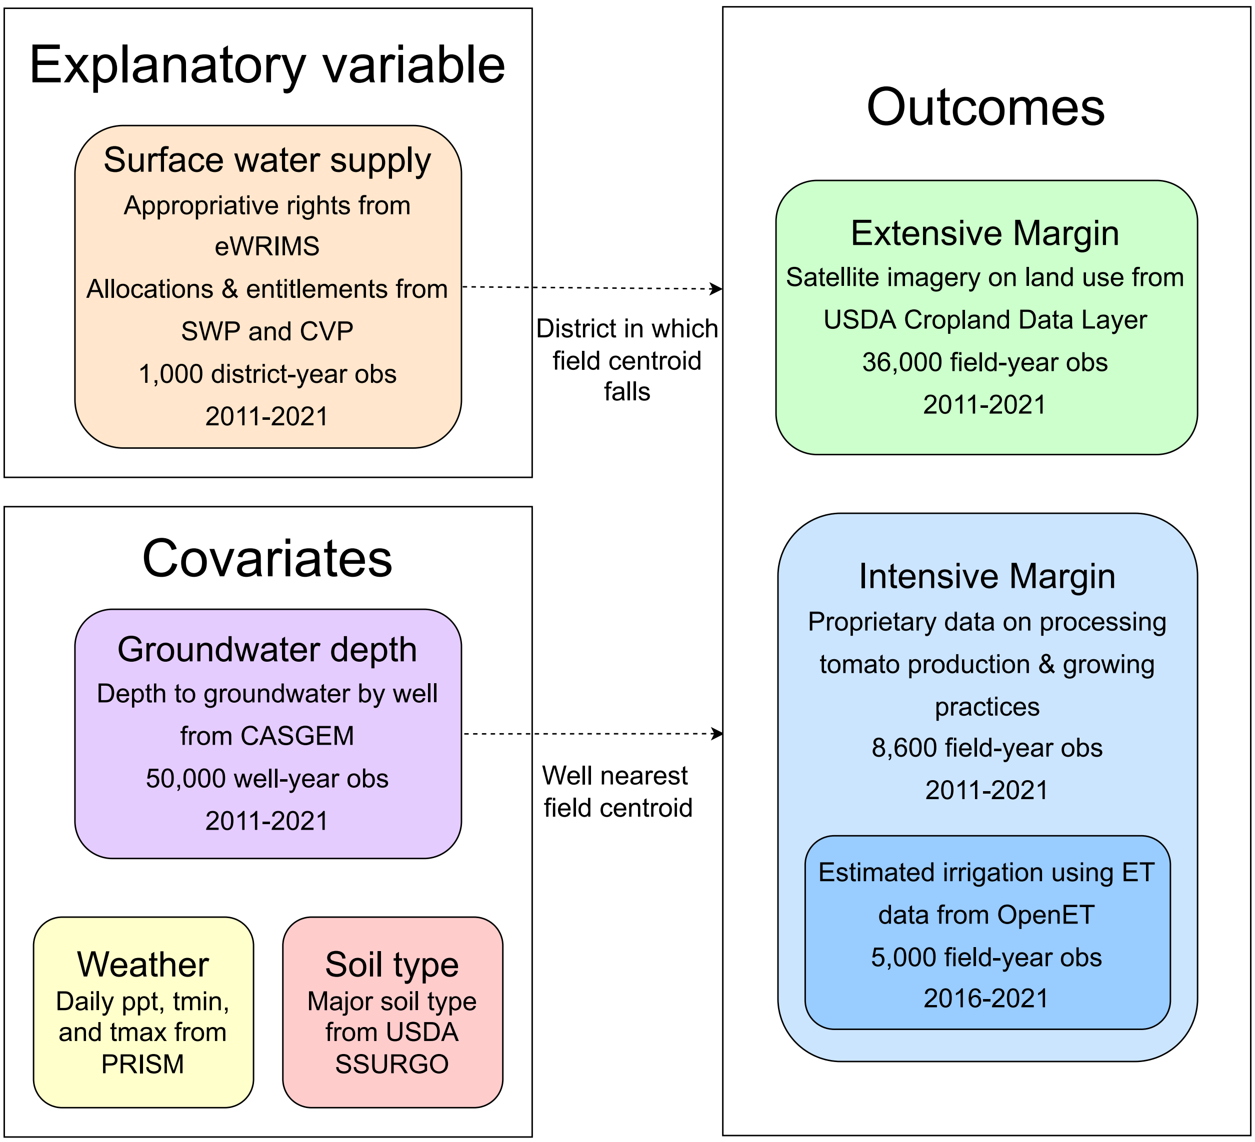

Figure 2.2: Diagram of key datasets and links between datasets

I defined surface water supply as the proportion of total surface water supplies available to the water district d in year t:

$$\text{surface water}_{dt} = \\ = \frac{\text{appropriative}_d + \text{entitlement}_d^{\text{CVP}} \times \text{allocation}_{dt}^{\text{CVP}} + \text{entitlement}_d^{\text{SWP}} \times \text{allocation}_{dt}^{\text{SWP}}}{\text{appropriative}_d + \text{entitlement}_d^{\text{CVP}} + \text{entitlement}_d^{\text{SWP}}} \tag{2.1}$$

where appropriative_d is the fixed annual appropriative right for district d, $\text{entitlement}_d^{\text{CVP}}$ and $\text{entitlement}_d^{\text{SWP}}$ are d district's entitlement from CVP and SWP respectively. The allocations $\text{allocation}_{dt}^{\text{CVP}}$ and $\text{allocation}_{dt}^{\text{SWP}}$ are the annual allocations for district d in year t from CVP and SWP respectively that depend on the contract types held by the district.

Finally, I spatially matched fields to the irrigation district(s) in which they are located using district boundaries from the California Department of Water Resources mapped in Figure 2.3. Irrigation districts can overlap and so a field may match with multiple districts. I assigned these fields a surface water supply equal to the average surface water supply across the multiple districts.

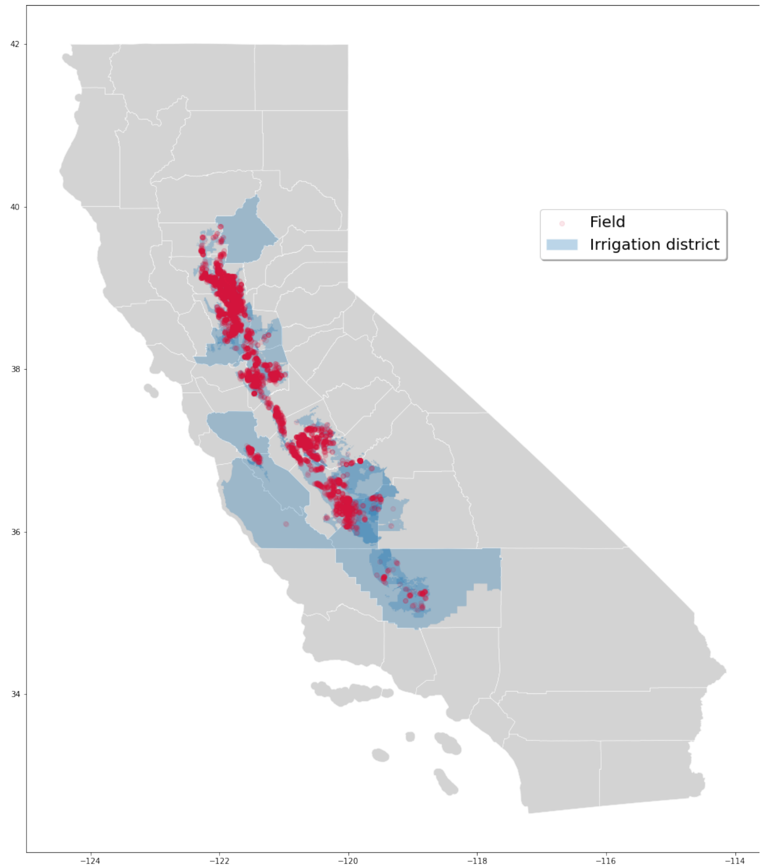

Figure 2.3: Irrigation districts and field locations

2.3.2 Outcomes: Extensive Margin

Growers can respond to water scarcity by changing whether they fallow or plant a crop, and which crop they plant. I use the USDA's Cropland Data Layer (USDA NASS, 2022) to identify the land use on a sample of fields predominately in the Sacramento Valley and San Joaquin Valley in California. I select the sample of fields to match fields in the proprietary data i.e. for at least one season, these fields grow processing tomatoes and I observe each field's acreage and centroid.

The Cropland Data Layer is an annual raster where each 30-by-30 meter pixel is identified as a crop or other land use category. In lieu of field boundaries, I create a circular buffer around each field's centroid that encloses an area equal to 20 percent of the field's acreage. For every field-year observation, I identify the land use that covers more than 50 percent of the pixels within the buffer[2]. This generates a panel of around 36,000 field-year observations of land use from 2011 to 2021. The most common land use is processing tomatoes (34 percent of field-year observations), followed by grain (8 percent wheat and 8 percent other grain), fallow (9 percent), cotton (9 percent), and alfalfa (7 percent).

I classify crops as high-water or low-water crops based on their average applied water per acre from the University of California Agriculture and Natural Resources crop production guides (UC ANR, 2022). High-water crops are those that use more than 3 acre-feet per acre on average, such as cotton, alfalfa, nut trees, and pasture. Low-water crops include tomatoes, wheat, melons, and most other grains and vegetables.

2.3.3 Outcomes: Intensive Margin

Processing Tomato Data

In addition to extensive margin adjustments, growers can respond to water scarcity along the intensive margin by changing how they grow. However, information about growing practices is rarely systematically collected and published. To fill this gap, I use detailed proprietary data on every field contracted with a large tomato processor from 2011–2021. These data are not from surveys but rather from administrative records collected by the processor for the purposes of contracting and payment. Fields are linked to anonymized grower identifiers so I can track growers across years. Crucially, they include the coordinates of each field centroid which I use to merge the datasets.

Growers can change their practices in several ways to conserve water. The first growing practice I observe is the date a grower plants a field. Planting date affects expected temperature exposure: planting earlier exposes plants to cooler temperatures on average which decreases the plant's water demands. I convert planting date to planting day of year so that timing can be compared across years. The average planting day of year is 107 or mid-April, with the earliest fields planted in mid-February and the latest planted at the end of June.

[2]If the modal pixel is less than 50 percent, the field is classified as having missing data. Further, fields identified as developed, open water, wetlands, shrubland, and forest are also classified as having missing data.

The second intensive margin adjustment I observe is the season length: the number of days between planting and the final day of harvest[3]. Since tomato plants require adequate access to water every day they are in the ground, a shorter season uses less water. The average season length is 134 days but this varies from a low of 98 days to a high of 175 days.

The third intensive margin adjustment I observe is the variety of tomato grown in a field. There are many varieties of processing tomatoes identified in the processor data with varying characteristics to suit particular growing practices or end uses. I collect additional information about each variety's attributes from AgSeeds (AgSeeds, 2020). I focus on varieties characterized as "early" – a variety that requires fewer days to reach maturity and therefore has a shorter season. On average, around 15 percent of acreage is planted with a fast-maturing variety.

The final growing practice I observe is the irrigation technology used on each field. I measure the likelihood tomatoes are grown using drip irrigation technology rather than sprinkler or furrow irrigation technologies in a given year. Most fields are equipped with drip irrigation (87% of field-year observations). The remaining fields are fitted with less water-efficient furrow (10 percent) or sprinkler (3 percent) irrigation systems. Growers infrequently change irrigation systems on their fields: only 100 fields (2 percent of all fields) change irrigation systems over the 11 year period, of which 80 percent were upgraded to drip irrigation.

This detailed dataset also includes field-level observations of tonnage of tomatoes harvested, acreage, quality attributes of fruit, and how each individual quality attribute affects price (see 1.3 for more details). I calculate yield as total tons divided by harvested acres for each field-year observation. Following 1.3, I calculate an economically meaningful measure of quality by isolating the effect of quality on price. Finally, I calculate field-level revenue per acre by multiplying quality-adjusted price by paid tons and dividing by harvested acres.

[3]For large fields, harvest can occur over several days.

Estimated Irrigation

I estimate irrigation applied to fields growing processing tomatoes using novel remotely-sensed evapotranspiration (ET) data from OpenET, a joint venture between NASA, the Desert Research Institute, and the Environmental Defense Fund. These recently released data offer several advantages over other data on water use. First, they do not rely on data disclosed by growers that may mistakenly or intentionally be misreported. Second, OpenET publish ET at a very fine spatial resolution of 30 meters by 30 meters—0.4 percent of the average field size in my sample. OpenET assures its data quality by undergoing extensive ground-truth validation and peer-review processes. A disadvantage of these data are their span: they are currently only available from 2016 onward[4].

The variable of interest is applied irrigation, which is the amount of water a producer applies to the plants. ET measures a component of applied irrigation known as consumptive water use—water used by the plants to facilitate crop growth and cooling. As shown in Figure 2.4, applied irrigation (blue) is not equal to consumptive water use (red) because (a) plants consume water from both irrigation and precipitation, and (b) not all irrigated water applied is used by the plant because of irrigation inefficiencies. I estimate applied irrigation from consumptive water use by first deducting precipitation from consumptive water use to calculate consumptive irrigation water use. Unreliable meteorological data may bias estimates of precipitation and introduce error into estimates of irrigation volumes (Foster et al., 2020). This is of little concern in my setting. Meteorological data for California are high quality and there is very little precipitation during the summer growing season to estimate.

The final consideration is nonconsumptive irrigation water use or irrigation inefficiency: water applied through irrigation that is not used by the plants because of run-off or deep percolation. I later show that grower-level time-invariant heterogeneity in irrigation inefficiency is controlled through the inclusion of grower fixed effects. This allows irrigation inefficiency

[4]OpenET plan to release data on earlier years in the near future.

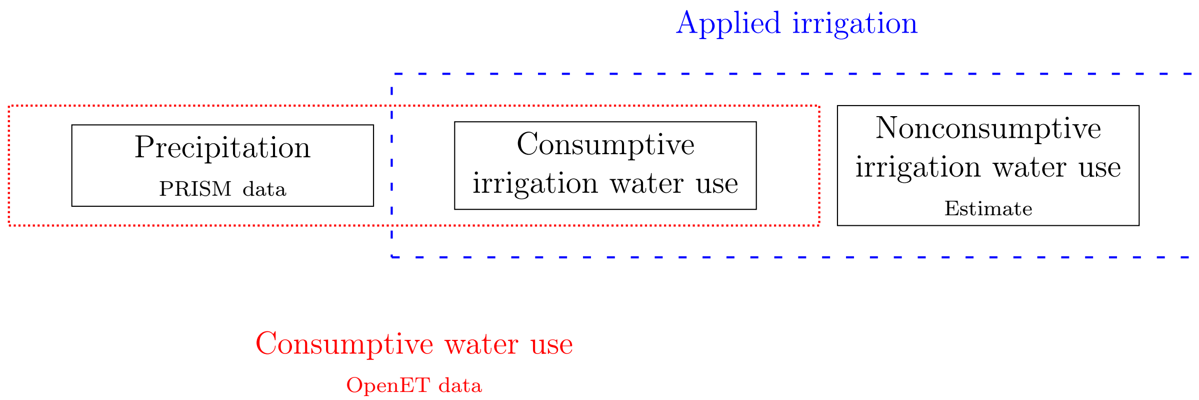

Figure 2.4: Estimating applied irrigation

to vary across growers but assumes it remains fixed over the six year period[5].

2.3.4 Covariates

Finally, I compile publicly available data on covariates that may be correlated with water scarcity.

Groundwater Supply

Unlike surface water, groundwater pumping was effectively unregulated between 2011 and 2021[6]. However, groundwater pumping is energy intensive and groundwater depth is a critical factor driving pumping costs. I use depth to groundwater in feet below the ground's surface as a proxy for groundwater availability.

Data on well locations and periodic groundwater depth measurements come from the Department of Water Resources via the California Natural Resources Agency. These include data collected through the California Statewide Groundwater Elevation Monitoring Program and Sustainable Groundwater Management Act Portal's Monitoring Network Module. For each year (2011-2021) and each well, I calculate the average depth to groundwater in feet below the ground's surface for the duration of the processing tomato growing season (Febru-

[5] An alternative is to allow irrigation inefficiency to vary across fields rather than growers. I later show that the two approaches produce similar results.

[6] The State of California introduced a new groundwater management legislation package in 2014 but it was not binding during the period of analysis

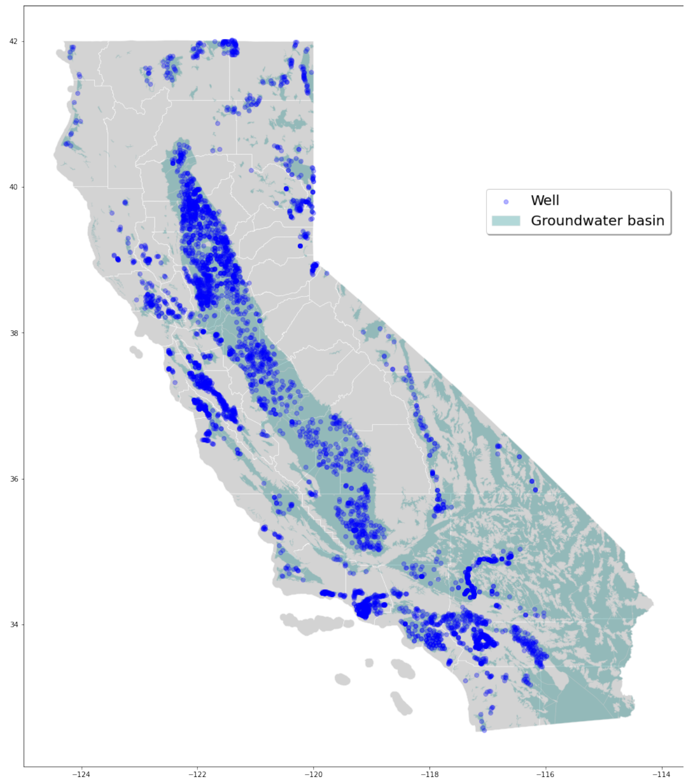

Figure 2.5: Groundwater basins and well locations

ary to October). I exclude wells missing two or more years of data. This generates a panel of 50,906 well-year observations for 4,744 wells shown in Figure 2.5. I match each field to the groundwater level at the well nearest the field centroid[7]. Multiple fields can match to the same well: indeed, the 3,300 fields match to 578 unique wells.

Weather

I collect gridded, daily temperature and precipitation data from PRISM (PRISM Climate Group, Oregon State University, 2020). I match weather data to each field-level observation by identifying the 4km PRISM grid in which the field centroid falls.

I convert daily observations of minimum and maximum temperature into degree days: a measure of accumulated exposure to heat pioneered by Snyder (1985) and popularized by Schlenker & Roberts (2009) (see Ortiz-Bobea (2021) for details). Mid-range temperatures

[7]An alternative is to match fields to the average groundwater level for the basin in which a field is located. I use this measure as a robustness check (see Appendix 2.B.)

are ideal for processing tomato growth while exposure to hot temperatures (above 35°C) and cool temperatures (below 10°C) slows growth and potentially damages the plant (e.g. Hartz et al. (2008)). For each field-year observation, I calculate growing season degree days for three intervals: below 10°C, between 10°C and 35°C, and above 35°C. This method has several advantages over averaging temperature across the growing season. First, degree days preserves the distribution of temperature exposure throughout the day and thus captures exposure to extreme temperatures. Second, since each interval enters the model separately, this method allows for possible nonlinearity in the effect of temperature exposure as suggested by Hartz et al. (2008).

Soil type

A farm may consist of fields of varying quality. One observable aspect of field quality is soil type. I collect data on soil composition from the National Cooperative Soil Survey (NRCS USDA, 2020). I spatially match each field centroid to its major soil type. Most fields have alluvial soils with the remaining fields being categorized as either eolian, lacustrine, residuum, or high in organic material.

Table 2.1: Summary statistics: extensive margin

	units	count	mean	sd	min	max
Surface water supply	%	36162	64.99	37.42	0.00	100.00
Land use						
High-water crop	%	36162	30.35	45.98	0.00	100.00
Alfalfa	%	36162	7.09	25.66	0.00	100.00
Cotton	%	36162	9.11	28.78	0.00	100.00
Garlic	%	36162	1.85	13.47	0.00	100.00
Grapes	%	36162	0.49	6.99	0.00	100.00
Onions	%	36162	1.49	12.13	0.00	100.00
Other grain	%	36162	4.74	21.26	0.00	100.00
Tree crop	%	36162	5.58	22.94	0.00	100.00
Low-water crop	%	36162	60.99	48.78	0.00	100.00
Corn	%	36162	4.17	19.99	0.00	100.00
Dry Beans	%	36162	1.96	13.85	0.00	100.00
Melons	%	36162	1.48	12.09	0.00	100.00
Other grain	%	36162	3.53	18.44	0.00	100.00
Other vegetable	%	36162	1.82	13.36	0.00	100.00
Peppers	%	36162	0.16	4.01	0.00	100.00
Sunflowers	%	36162	5.80	23.38	0.00	100.00
Tomato	%	36162	34.23	47.45	0.00	100.00
Tree crop	%	36162	0.03	1.83	0.00	100.00
Wheat	%	36162	7.81	26.83	0.00	100.00
Fallow	%	36162	8.67	28.13	0.00	100.00
Covariates						
Groundwater level	feet	36162	97.94	99.89	-42.50	625.50
Average min. temp.	° C	36162	11.46	0.98	7.60	14.94
Average max. temp.	° C	36162	27.84	1.24	22.15	30.61
Total precipitation	mm	36162	151.36	100.75	11.52	682.62

Notes: Summary statistics for the merged data from 2011-2021 and weighted by field acreage. Weather variables are for the February to October growing season.

Table 2.2: Summary statistics: intensive margin

	units	count	mean	sd	min	max
Growing practices						
Planting day of year	days	8617	107.31	22.98	48.00	179.00
Growing days	days	8617	134.46	9.47	98.00	175.00
Early variety	%	8617	14.59	35.31	0.00	100.00
Drip irrigation	%	8617	87.18	33.43	0.00	100.00
Furrow irrigation	%	8617	9.67	29.55	0.00	100.00
Sprinkler irrigation	%	8617	3.16	17.48	0.00	100.00
Consumptive irrigation water use	Acre-feet per acre	4905	2.00	0.29	0.59	3.04

Notes: Summary statistics for the merged subset of observations in processing tomatoes data from 2011-2021 and weighted by field acreage.

2.4 Methods

The conceptual model in Appendix 2.A suggests that a grower may respond to water availability along different margins of adjustment. First, I ask if growers respond to water availability along the extensive margin: whether to plant and what to plant. Outcomes of interest are the probability of fallowing, probability of planting a low-water crop, and probability of planting a high-water crop. Second, I ask if and how growers respond along the intensive margin, conditional on growing a low-water crop. The intensive margin model uses the subset of observations of processing tomatoes under contract with a large processor. For these observations, I observe detailed data on growing practices that make up the intensive margin responses: planting day of year, season length, varietal choice, irrigation technology choice, and estimated irrigation. Finally, I analyze the net effect of intensive margin adjustments induced by water scarcity on processing tomato yield, quality, and revenue.

The preferred specification is:

$$\text{outcome}_{it} = \beta \,\text{surface water supply}_{d(i)t} + \gamma \,\text{controls}_{it} + \alpha_{g(i)} + \lambda_t + \epsilon_{it}, \tag{2.2}$$

where surface water supply$_{d(i)t}$ is the proportion of surface water supply available in year t to district d in which field i falls. I weight observations by field acreage which yields an estimate representative of statewide effects[8]. I control for the groundwater depth in year t at the well nearest the field i, temperature exposure and precipitation for field i in year t, and the major soil type for field i in controls$_{it}$. I also include grower fixed effects $\alpha_{g(i)}$, where $g(i)$ is the grower associated with field i, and year fixed effects λ_t. Standard errors ϵ_{it} are heteroskedastic robust and clustered by grower and district-year to account for the possibility of both temporal and spatial dependence.

The inclusion of two-way fixed effects removes both time and grower-invariant confounders. Year fixed effects capture state-wide factors common to all fields in a year, including factors affecting the agricultural economy and the processor's preferences. Some decisions made jointly by the grower and processor, such as the variety and planting date. Since I use data from one processor, I can control for the effect of the processor's preferences through the inclusion of year fixed effects. For example, the year fixed effect would absorb the processor shifting its entire season schedule in response to seasonal conditions. Conditioning on the processor's response, the remaining variation reflects grower decisions. Year fixed effects also absorb the statewide average profitability of particular land-use choices, an important driver of land-use decisions.

Grower fixed effects[9] subsume average surface water availability to the grower so identifying variation in surface water comes from year-to-year variation in allocations. The estimate β captures how the average grower responds to a change in their water availability.

2.4.1 Identifying Assumptions

Giving β in Equation 2.2 a causal interpretation requires that surface water supply is quasi-random, conditional on average supply. This is reasonable since surface water allocations are set by government agencies in response to exogenous environmental conditions. Districts differ in their average surface water availability because those with high-priority water access

[8] As a robustness check, I also estimate the model without weights. If the coefficients change dramatically, it would be evidence of misspeficiation or heterogeneity. The results are largely robust to weighting decisions (see Appendix 2.B).

[9] An alternative is to use field fixed effects that would absorb unobservable differences in field quality within each grower's farm. Estimates using field fixed effects are consistent in terms of sign and magnitude but the intensive margin results are more imprecise (see Appendix 2.B).

experience smaller declines in their allocation than those with low-priority access. However, this is not a concern for identification since a grower's average allocation is captured by grower fixed effects. Critically, a grower cannot influence variation in their allocation to individual fields.

One may be concerned that grower fixed effects do not adequately capture omitted variables for growers who have fields that span multiple irrigation districts. As a robustness check, I estimate a model with grower-by-district fixed effects to account for growers who span multiple water districts. I later show that most results are virtually unchanged (see Appendix 2.B).

The greatest threat to identification is time-varying unobservable characteristics that are correlated with both water availability and outcomes. For example, regions may experience different trends in crop profitability that are linked to local factors like processing capacity (e.g. Sayre (2022)). Indeed, crop choices have progressively switched from annuals to perennials, particularly almonds, in parts of California (Carman, 2019). As a robustness check, I include a linear regional time trend that captures potential divergence in agronomic or economic conditions and results are virtually identical in terms of magnitude and significance (see Appendix 2.B).

2.4.2 Heterogeneity

While the main specification in Equation 2.2 recovers the average response to water scarcity, it may hide important heterogeneity in how growers respond. Recall that even within the same year, water allocations can vary dramatically across California due to differences in water access priority: districts with high-priority access experience smaller declines in their water allocation compared to districts with low-priority contracts. The conceptual model in Appendix 2.A suggests that investment in water-saving practices depends on the stringency of water allocations, implying that growers' response to water scarcity depends on their surface water access. To test this, I interact surface water supply with a dummy variable “high priority$_{d(i)}$” that identifies districts with high-priority water access:

$$\begin{aligned}\text{outcome}_{it} =& \kappa_1 \, \text{surface water supply}_{d(i)t} + \kappa_2 \, \text{high priority}_{d(i)} \\ &+ \kappa_3 \left(\text{high priority}_{d(i)} \times \text{surface water supply}_{d(i)t}\right) + \gamma \, \text{controls}_{it} + \alpha_{g(i)} + \lambda_t + \epsilon_{it},\end{aligned} \tag{2.3}$$

where κ_1 is the response of growers with low-priority water access and $\kappa_1 + \kappa_3$ is the response of growers with high-priority water access.

Since priority status only varies in the cross-section, it is possible there are features of the environment that vary spatially and are correlated with (but not caused by) priority status, for example land quality. While I include variables like soil type to control for such characteristics, care should be taken not to interpret these estimates as the causal effect of changing a grower's priority status. That said, this analysis does shed light on how growers' response to water scarcity depends on their priority status.

2.5 Results

Table 2.3 reports estimates of the effect of surface water supply on extensive margin outcomes. I find that growers respond to water scarcity by planting fewer acres to crops, leaving more acres fallow. A one standard deviation decline in surface water supply significantly increases the share of land fallowed by 2.9 percentage points. This is indicative of a reasonably large change in planting decisions given the average share of fallowed acreage is only 9 percent. Overall, my result is comparable to prior work looking at the effects of water scarcity on planting decisions (Hagerty (2021); Manning et al. (2017)).

Water scarcity also affects the types of crops growers choose to plant. A one standard deviation decrease in surface water supply significantly decreases the share of acreage planted to high-water crops like cotton and pasture by 4.3 percentage points. This is consistent with the conceptual model that shows high-water crops are not the most profitable land use when water allocations are low (see Appendix 2.A). Growers are also slightly more likely to plant low-water crops when water is scarce, however the effect is not statistically significant. I also estimate the effect on processing tomato acreage given its relative importance for this sample of fields. Similar to the results for low-water crops (of which processing tomatoes is a part), the estimated impact of water scarcity on processing tomato acreage is small and insignificant. Together, these results plausibly suggest that growers maintain acreage in low-water crops and choose to fallow fields that would otherwise be planted to high-water crops.

Table 2.3: Effect of water supply on the extensive margin

	(1) Fallow	(2) High-water crop	(3) Low-water crop	(4) Processing tomato
	prop.	prop.	prop.	prop.
Surface water supply	-0.0008***	0.0011***	-0.0004	-0.0002
	(0.0002)	(0.0002)	(0.0002)	(0.0003)
Marginal effect	0.029***	-0.043***	0.014	0.008
	(0.0078)	(0.0087)	(0.0090)	(0.0096)
Grower fixed effects	✓	✓	✓	✓
Year fixed effects	✓	✓	✓	✓
Observations	36161	36161	36161	36161

Notes: Each column shows the results from a separate regression model for the outcome variable identified in the column header, using the full panel of field-year observations from 2011-2021. The reported estimate is the effect of a 1 percentage point increase in surface water supply. The marginal effect is the estimated effect multiplied by negative one standard deviation in surface water supply. Additional controls in all models include: groundwater depth, temperature exposure, precipitation, and major soil type. Standard errors (in parentheses) are two-way clustered by grower and district-year. Significance: * p<0.05, ** p<0.01, *** p<0.001. Estimates are weighted by field acreage.

In addition to extensive margin adjustments, growers also respond to water scarcity by changing growing practices as reported in Table 2.4. A one standard deviation decrease in surface water supply causes processing tomato growers to plant fields earlier by 1.6 days on average. All else equal, planting earlier exposes tomato plants to cooler temperatures and reduces water demand. The effect is significant but small relative to the wide planting window. Recall planting dates are negotiated between growers and processors and processors value spread in planting dates to avoid processing bottlenecks. A small effect is consistent with the fact that growers set their planting date in conjunction with the processor.

Growers are more likely to plant varieties that require fewer days to reach maturity during periods of water scarcity. A one standard deviation decrease in surface water supply significantly increases the share of processing tomato acreage planted to early maturing varieties by 2.9 percentage points. I find that the growing season length decreases during a water-scarce year, but the effect is small and not statistically significant.

During periods of water scarcity, growers prefer to use drip irrigation over less water efficient irrigation systems. A one standard deviation decrease in surface water supply significantly increases the share of planted acreage using drip irrigation by 2 percentage points. This is a relatively small magnitude compared with the large proportion of acreage using drip irrigation (87 percent). Two-thirds of growers are only ever observed using drip irrigation on their fields which limits further adaptation along this margin.

There are two main mechanisms through which growers change the proportion of tomato acreage grown with drip irrigation. The first is an enduring transition from sprinkler or furrow irrigation to drip irrigation on a particular field. However, growers infrequently change irrigation systems on a given field, and when they do it is more likely a long-term decision than a short-run response to water scarcity. Drip irrigation results are no longer significant with the inclusion of field fixed effects, suggesting that growers are not updating field-level irrigation systems in response to water scarcity. The second and more likely mechanism is a reallocation of planted and fallowed acreage across fields. During a water scarce year, growers are more likely to allocate fields equipped with drip irrigation to processing tomatoes. Fields equipped with less efficient irrigation systems are either growing other crops or left fallow, although the latter is more likely given the extensive margin adjustments I observe during water scarce years.

I consider the effect of water availability on the log of estimated irrigation between 2016 and 2021 – a shorter panel than the other models reflecting OpenET's current data availability. Recall that irrigation inefficiency — a fixed percentage of any applied water lost due to run-off or percolation — is unobserved. I include grower fixed effects to absorb time-invariant heterogeneity in irrigation inefficiency. Taking the log of irrigation is important to account for the proportional nature of irrigation inefficiency, for example if a grower consistently loses 10% of applied water to percolation, the fixed effect will only control for this inefficiency if the dependent variable is in logs.

Despite growers adjusting their growing practices, I find little evidence that water scarcity causes growers to decrease the amount of water used to grow processing tomatoes. The effects of surface and groundwater availability on estimated irrigation are both precise zeroes. In a robustness check, I replace grower fixed effects with field fixed effects that allows irrigation inefficiency to vary across fields rather than growers and the results are unchanged.

Table 2.4: Effect of water supply on the intensive margin

	(1) Planting day	(2) Season length	(3) Early variety	(4) Drip irrigation	(5) ln(irrigation)
	day of year	no. days	prop.	prop.	
Surface water supply	0.0439***	0.0006	-0.0008***	-0.0005*	0.0001
	(0.0117)	(0.0014)	(0.0002)	(0.0002)	(0.0002)
Marginal effect	-1.630***	-0.022	0.029***	0.020**	-0.004
	(0.4337)	(0.0528)	(0.0085)	(0.0090)	(0.0067)
Grower fixed effects	✓	✓	✓	✓	✓
Year fixed effects	✓	✓	✓	✓	✓
Observations	8609	8609	8609	8609	4896

Notes: Each column shows the results from a separate regression model for the outcome variable identified in the column header, using the subset of observations in the processing tomatoes data from 2011-2021. Results in column (5) are for 2016-2021. The reported estimate is the effect of a 1 percentage point increase in surface water supply. The marginal effect is the estimated effect multiplied by negative one standard deviation in surface water supply. Additional controls in all models include: groundwater depth, temperature exposure, precipitation, and major soil type. Standard errors (in parentheses) are two-way clustered by grower and district-year. Significance: * p<0.05, ** p<0.01, *** p<0.001. Estimates are weighted by field acreage.

The net effect of surface water supply on processing tomato yield, quality, and revenue is reported in Table 2.5. While the coefficients indicate that yield (and therefore revenue) decline during periods of water scarcity, imprecision in the estimates means I cannot rule out no effect.

Table 2.5: Effect of water supply on yield, quality, and revenue

	(1) ln(yield)	(2) ln(quality)	(3) ln(revenue)
Surface water supply	0.0003	0.0000	0.0003
	(0.0003)	(0.0001)	(0.0003)
Marginal effect	-0.011	-0.001	-0.013
	(0.0102)	(0.0023)	(0.0111)
Grower fixed effects	✓	✓	✓
Year fixed effects	✓	✓	✓
Observations	8609	8609	8609

Notes: Each column shows the results from a separate regression model for the outcome variable identified in the column header, using the subset of observations in the processing tomatoes data from 2011-2021. The reported estimate is the effect of a 1 percentage point increase in surface water supply. The marginal effect is the estimated effect multiplied by negative one standard deviation in surface water supply. Additional controls in all models include: groundwater depth, temperature exposure, precipitation, and major soil type. Standard errors (in parentheses) are two-way clustered by grower and district-year. Significance: * $p<0.05$, ** $p<0.01$, *** $p<0.001$. Estimates are weighted by field acreage.

Growers respond to water scarcity differently depending on their water access priority, as shown in Tables 2.6 and 2.7. I find that growers with high-priority water access rely exclusively on crop switching to manage reductions in water availability. In an effort to conserve water, growers decrease their acreage in high-water crops by 6 percentage points but this is completely offset with an increase in low-water crop acreage. There is no significant evidence that they increase fallowing or adjust their growing practices, however imprecision in some estimates means adjustment along these margins cannot be completely ruled out.

In contrast, growers with low-priority water access rely on more costly margins of response. These growers respond to water scarcity by fallowing acreage that would otherwise be dedicated to high-water crops. The share of acreage in low-water crops like processing tomatoes is effectively unchanged. But growers adjust *how* they grow processing tomatoes. During a water-scarce year, growers with low-priority water access are more likely to plant faster maturing varieties earlier in the season.

Table 2.6: Effect of water supply on the extensive margin by water priority

	(1) Fallow	(2) High-water crop	(3) Low-water crop	(4) Processing tomato
	prop.	prop.	prop.	prop.
Surface water supply	-0.0008***	0.0009***	-0.0001	-0.0001
	(0.0002)	(0.0002)	(0.0002)	(0.0003)
High priority=1 × Surface water supply	0.0015*	0.0009	-0.0025*	-0.0023**
	(0.0007)	(0.0007)	(0.0010)	(0.0007)
Surface water supply + High priority=1 × Surface water supply	0.001	0.002**	-0.003**	-0.002***
	(.0007)	(.0007)	(.001)	(.0008)
Grower fixed effects	✓	✓	✓	✓
Year fixed effects	✓	✓	✓	✓
Observations	36161	36161	36161	36161

Notes: Each column shows the results from a separate regression model for the outcome variable identified in the column header, using the full panel of field-year observations from 2011-2021. The reported estimate is the effect of a 1 percentage point increase in surface water supply for different subgroups. The first row "Surface water supply" is the estimate for growers with low-priority water access, and the third row "Surface water supply + High priority=1 × Surface water supply" is the estimate for growers with high-priority water access. Additional controls in all models include: groundwater depth, temperature exposure, precipitation, and major soil type. Standard errors (in parentheses) are two-way clustered by grower and district-year. Significance: * $p<0.05$, ** $p<0.01$, *** $p<0.001$. Estimates are weighted by field acreage.

As shown in Table 2.8, growers with low-priority water access incur a yield penalty when growing processing tomatoes during a water-scarce year. This captures the net effect of changes along the intensive margin. As a result, their revenue is 3.4 percent smaller relative to a year with average water availability.

Together, these findings are consistent with the conceptual model. Growers can manage small declines in their allocation by switching from high-water to low-water crops. The incentive to invest water-saving practices emerges when growers experience large declines in their allocation. However, these practices incur the cost of a lower yield and therefore revenue per acre.

Table 2.7: Effect of water supply on the intensive margin by water priority

	(1) Planting day	(2) Season length	(3) Early variety	(4) Drip irrigation	(5) ln(irrigation)
	day of year	no. days	prop.	prop.	
Surface water supply	0.0280*	0.0003	-0.0006*	-0.0002	0.0001
	(0.0132)	(0.0016)	(0.0003)	(0.0002)	(0.0002)
High priority=1 × Surface water supply	-0.0361	-0.0012	0.0007	0.0009	0.0002
	(0.0339)	(0.0028)	(0.0007)	(0.0008)	(0.0006)
Surface water supply + High priority=1 × Surface water supply	-0.008	-0.001	0.000	0.001	0.000
	(.0366)	(.0036)	(.0008)	(.0009)	(.0007)
Grower fixed effects	✓	✓	✓	✓	✓
Year fixed effects	✓	✓	✓	✓	✓
Observations	8609	8609	8609	8609	4896

Notes: Each column shows the results from a separate regression model for the outcome variable identified in the column header, using the subset of observations in the processing tomatoes data from 2011-2021. Results in column (5) are for 2016-2021. The reported estimate is the effect of a 1 percentage point increase in surface water supply for different subgroups. The first row "Surface water supply" is the estimate for growers with low-priority water access, and the third row "Surface water supply + High priority=1 × Surface water supply" is the estimate for growers with high-priority water access. Additional controls in all models include: groundwater depth, temperature exposure, precipitation, and major soil type. Standard errors (in parentheses) are two-way clustered by grower and district-year. Significance: * p<0.05, ** p<0.01, *** p<0.001. Estimates are weighted by field acreage.

Table 2.8: Effect of water supply on yield, quality, and revenue by water priority

	(1) ln(yield)	(2) ln(quality)	(3) ln(revenue)
Surface water supply	0.0009**	-0.0001	0.0009**
	(0.0003)	(0.0001)	(0.0003)
High priority=1 × Surface water supply	-0.0004	0.0003	0.0001
	(0.0008)	(0.0002)	(0.0008)
Surface water supply + High priority=1 × Surface water supply	0.001	0.000	0.001
	(.0007)	(.0002)	(.0007)
Grower fixed effects	✓	✓	✓
Year fixed effects	✓	✓	✓
Observations	8609	8609	8609

Notes: Each column shows the results from a separate regression model for the outcome variable identified in the column header, using the subset of observations in the processing tomatoes data from 2011-2021. The reported estimate is the effect of a 1 percentage point increase in surface water supply for different subgroups. The first row "Surface water supply" is the estimate for growers with low-priority water access, and the third row "Surface water supply + High priority=1 × Surface water supply" is the estimate for growers with high-priority water access. Additional controls in all models include: groundwater depth, temperature exposure, precipitation, and major soil type. Standard errors (in parentheses) are two-way clustered by grower and district-year. Significance: * p<0.05, ** p<0.01, *** p<0.001. Estimates are weighted by field acreage.

2.6 Discussion and Conclusion

Growers respond to water scarcity by changing whether they plant and what they plant. During a water-scarce year, I find that growers fallow a greater share of their acreage and reduce the share of acreage in high-water crops. Water access priority drives differences in how growers respond. A water-scarce year causes growers in districts with low-priority water access to fallow instead of planting high-water crops. By comparison, growers in districts with high-priority water access respond by switching from high-water to low-water crops and do not significantly increase fallowing.

I also analyze the effect of water supply on margins that earlier work could not consider because of data constraints. In contrast to other studies that argue California farmers do not adapt along the intensive margin (Hagerty (2021); Burlig et al. (2021)), I find significant evidence that growers change growing practices in response to water scarcity. Growers who contract processing tomatoes with a large tomato processor react to water scarcity by planting earlier and choosing fast-maturing varieties. Most of the intensive margin response is driven by growers with low-priority access to water. Growers are more likely to use drip irrigation over less efficient irrigation methods but the effect is small. Two-thirds of growers are only ever observed using drip irrigation on their fields which limits further adaptation along this margin.

Decreased water supply under climate change will create challenges for agriculture but the extent of damages will depend on how growers and the broader agricultural industry adapt. I show that growers respond to water scarcity along the extensive and intensive margins but that these responses are not without cost. I find that growers in low-priority water districts earn 3.4 percent less revenue during a water-scarce year (one standard deviation decline in water availability) because of reduced yield alone. To scale up this loss of revenue, I do a partial equilibrium, back-of-the-envelope calculation, assuming growers in my sample are representative of the broader California processing tomato industry. I apply the 3.4 percent loss of revenue to the share of California's $1 billion processing tomato industry produced in low-priority water districts. This calculation suggests that intensive margin adjustments in response to water scarcity cost processing tomato growers $20 million in lost revenue in a water-scarce year.

These findings have important implications for industry and policy makers. First, my research shows growers engage in water-saving practices. This reinforces the need to support and create opportunities for growers to invest in climate-smart production.

Finally, my findings indicate large returns from research and development of cultivars that can be profitably grown using less water. While fallowing effectively conserves water, it also fails to generate income for farmers and agricultural products for a growing global population. Investment in new cultivars may assist growers conserve water without fallowing.

References

AgSeeds (2020). Season Review. Available at `http://www.agseeds.com/`.

Aragón, F. M., Oteiza, F., & Rud, J. P. (2021). Climate change and agriculture: Subsistence farmers' response to extreme heat. *American Economic Journal: Economic Policy*, 13(1), 1–35.

Arias, P., Bellouin, N., Coppola, E., Jones, R., Krinner, G., Marotzke, J., Naik, V., Palmer, M., Plattner, G.-K., Rogelj, J., et al. (2021). *Climate Change 2021: The Physical Science Basis. Contribution of Working Group14 I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change; Technical Summary*. IPCC.

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

Burlig, F., Preonas, L., & Woerman, M. (2021). *Energy, Groundwater, and Crop Choice*. Technical report, National Bureau of Economic Research.

Carman, H. F. (2019). California's changing land use patterns for crop production, 1959-2017. *ARE Update*, 23(2), 9–11. University of California Giannini Foundation of Agricultural Economics.

CDFA (2021). 2020-2021 California Agricultural Statistics Review.

Cortignani, R. & Severini, S. (2009). Modeling farm-level adoption of deficit irrigation using positive mathematical programming. *Agricultural water management*, 96(12), 1785–1791.

Drysdale, K. M. & Hendricks, N. P. (2018). Adaptation to an irrigation water restriction imposed through local governance. *Journal of Environmental Economics and Management*, 91, 150–165.

DWR (2022). Groundwater Wells. Available at `https://water.ca.gov/Programs/Groundwater-Management/Wells`.

FAO (2023). Food and Agriculture Organization Corporate Statistical Database (FAOSTAT). Available at `https://www.fao.org/faostat/en/`.

Foster, T., Mieno, T., & Brozovic, N. (2020). Satellite-based monitoring of irrigation water use: assessing measurement errors and their implications for agricultural water management policy. *Water Resources Research*, (pp. e2020WR028378).

Graveline, N. & Mérel, P. (2014). Intensive and extensive margin adjustments to water scarcity in france's cereal belt. *European Review of Agricultural Economics*, 41(5), 707–743.

Hagerty, N. (2021). Adaptation to water scarcity in irrigated agriculture. Working paper.

Hartz, T., Miyao, G., Mickler, J., Lestrange, M., Stoddard, S., Nuñez, J., & Aegerter, B. (2008). Processing tomato production in California. *UC Vegetable Research & Information Center*.

Hayhoe, K., Cayan, D., Field, C. B., Frumhoff, P. C., Maurer, E. P., Miller, N. L., Moser, S. C., Schneider, S. H., Cahill, K. N., Cleland, E. E., et al. (2004). Emissions pathways, climate change, and impacts on California. *Proceedings of the national academy of sciences*, 101(34), 12422–12427.

Hendricks, N. P. & Peterson, J. M. (2012). Fixed effects estimation of the intensive and extensive margins of irrigation water demand. *Journal of Agricultural and Resource Economics*, (pp. 1–19).

Howitt, R., Medellín-Azuara, J., MacEwan, D., Lund, J. R., & Sumner, D. (2014). *Economic analysis of the 2014 drought for California agriculture.* Center for Watershed Sciences University of California, Davis, CA.

Jagnani, M., Barrett, C. B., Liu, Y., & You, L. (2021). Within-season producer response to warmer temperatures: Defensive investments by kenyan farmers. *The Economic Journal*, 131(633), 392–419.

Manning, D. T., Goemans, C., & Maas, A. (2017). Producer responses to surface water availability and implications for climate change adaptation. *Land Economics*, 93(4), 631–653.

Medellín-Azuara, J., Howitt, R., & Harou, J. (2012). Predicting farmer responses to water pricing, rationing and subsidies assuming profit maximizing investment in irrigation technology. *Agricultural water management*, 108, 73–82.

Moore, F. C., Baldos, U., Hertel, T., & Diaz, D. (2017). New science of climate change impacts on agriculture implies higher social cost of carbon. *Nature communications*, 8(1), 1–9.

Moore, M. R., Gollehon, N. R., & Carey, M. B. (1994). Multicrop production decisions in western irrigated agriculture: the role of water price. *American Journal of Agricultural Economics*, 76(4), 859–874.

NRCS USDA (2020). Soil Survey Geographic (SSURGO) Database for California. Available at `https://websoilsurvey.nrcs.usda.gov/`.

Ortiz-Bobea, A. (2021). Climate, agriculture and food. *arXiv preprint arXiv:2105.12044*.

PRISM Climate Group, Oregon State University (2020). PRISM Climate Data. Available at `http://prism.oregonstate.edu`.

Sapino, F., Pérez-Blanco, C. D., Gutiérrez-Martín, C., García-Prats, A., & Pulido-Velazquez, M. (2022). Influence of crop-water production functions on the expected performance of water pricing policies in irrigated agriculture. *Agricultural Water Management*, 259, 107248.

Sayre, J. E. (2022). Farm to firm: Clustering and return to scale in agricultural value chains. Working paper.

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

Schoengold, K., Sunding, D. L., & Moreno, G. (2006). Price elasticity reconsidered: Panel estimation of an agricultural water demand function. *Water resources research*, 42(9).

Schoengold, K. & Zilberman, D. (2007). The economics of water, irrigation, and development. *Handbook of agricultural economics*, 3, 2933–2977.

Snyder, R. L. (1985). Hand calculating degree days. *Agricultural and forest meteorology*, 35(1-4), 353–358.

Taraz, V. (2017). Adaptation to climate change: Historical evidence from the indian monsoon. *Environment and Development Economics*, 22(5), 517–545.

UC ANR (2022). Agricultural Production. Available at `http://anrcatalog.ucdavis.edu/`.

USDA NASS (2017). 2017 Census of Agriculture. Available at `www.nass.usda.gov/AgCensus`.

USDA NASS (2022). Cropland Data Layer. Available at `https://nassgeodata.gmu.edu/CropScape/`.

Wallander, S., Hrozencik, A., & Aillery, M. (2022). *Irrigation Organizations: Drought Planning and Response*. Technical report, USDA Economic Research Service.

2.A Conceptual framework

Consider a profit-maximizing farmer who relies on surface water to irrigate crops[10]. The farmer chooses to plant H acres of a crop with the outside option to fallow for zero profits. I assume all planted acres are harvested for a price p known at the time of planting[11]. The farmer faces the following profit maximization problem:

$$\begin{aligned} \max_{x \geq 0, e \in (0,1)} \quad & \pi = pHf(w(x,e);\theta) - cH - p_e e \\ \text{s.t.} \quad & x \leq A \end{aligned} \tag{2.4}$$

where f is per acre output or yield of the crop that is a function of effective water w and weather θ. Effective water w is the amount of water that is used by the plant, which differs from the applied irrigation x because of irrigation inefficiencies – excess water can run off or percolate. I introduce an efficiency parameter $e \in (0,1)$, and the farmer can invest in irrigation efficiency to cause a greater proportion of applied irrigation to be translated into effective water for a cost of p_e. It costs c per acre to produce the crop. The farmer is allocated A acre-feet per acre of water that varies year-to-year. The constraint is that the farmer may not irrigate more than their allocation.

Next, let's consider the relationship between crop yield and effective water w. I use a Von Liebig function of water productivity, which assumes that water exhibits constant returns until it reaches an upper bound after which additional water provides no additional benefits. This is reasonable in settings where overwatering or flooding is unlikely (Schoengold & Zilberman, 2007). Yield is capped at the maximum attainable yield $\overline{f}$. I assume applied irrigation x and irrigation efficiency e act as substitutes. For example, a farmer could achieve f^1 in Figure 2.A.1 by applying a lot of water with low efficiency or by applying less water with high efficiency.

[10] This is a static model and does not include dynamics that would be needed to capture profit-maximizing choices for perennial crops like almonds.

[11] While this model does not account for price risk, it is reasonable to assume that price risk is relatively small given the widespread use of contracts that stipulate price in California agriculture.

Investment in irrigation efficiency is costly, whereas applied water is unpriced but constrained by the allocation A. Applied water has a shadow price $\tilde{p}_A$ that is a consequence of degree of water scarcity imposed by the allocation. Figure 2.A.1 shows how the optimal level of irrigation efficiency changes with the allocation given a fixed price of investing in irrigation efficiency p_e. For the allocation A_1, the farmer will apply the entire allocation so that the constraint holds with equality $x = A_1$. The farmer chooses the optimal irrigation efficiency e_1^* that equates the relative prices $(\frac{\tilde{p}_{A_1}}{p_e})$ with the marginal rate of substitution, as shown in panel (a) of Figure 2.A.1. An increase in the allocation to A_2 decreases the scarcity of water relative to the price of investing in irrigation efficiency, leading to a smaller shadow price $\tilde{p}_{A_2}$. As shown in panel (b) of Figure 2.A.1, a high allocation leads to less investment in irrigation efficiency.

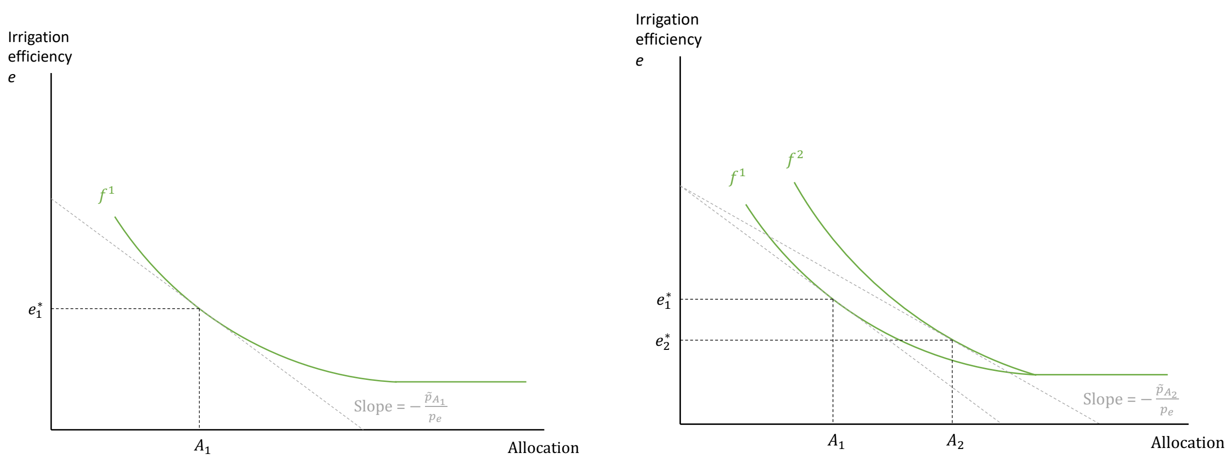

(a) Low allocation

(b) High allocation

Figure 2.A.1: Optimal investment in irrigation efficiency and output

For a particular crop option, a low allocation translates into smaller profits because (a) it induces costly investment in irrigation efficiency, and (b) production is limited. As the allocation increases, so too does profit because production increases and there is a reduction in the optimal irrigation efficiency and its associated cost. Eventually, profits reach a maximum level at the point where additional water provides no benefit to yield and there is little investment in water-saving practices thus minimizing the associated cost. Figure 2.A.2 panel (a) shows profits as a function of allocation for a low-water crop and high-water crop. The high-water crop is more profitable than the low-water crop so long as the allocation is large enough to meet the water demand of the plant. As the allocation declines, the most profitable land use changes from high-water crop to low-water crop to fallow, tracing out the envelope of profits panel (b) of Figure 2.A.2.

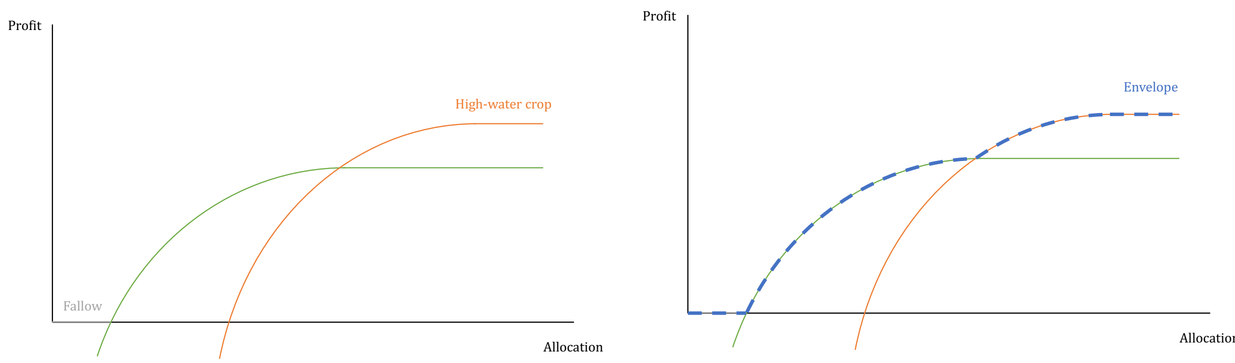

(a) Profits from different land use options

(b) Envelope

Figure 2.A.2: Profit function

In addition to extensive margin adjustments, there are also curved sections of the frontier where intensive margin adjustments occur. Conditional on growing the same crop, a decline in allocation may cause the farmer to respond along the intensive margin. This includes investing in growing practices that improve irrigation efficiency such as using drip irrigation. Growers can also reduce the quantity of effective water at the cost of a yield penalty. Responding along the intensive margin allows growers to conserve water and remain in agricultural production.

2.B Robustness checks

A threat to identification is the presence of time-varying confounders correlated with water scarcity and the outcomes. One notable difference between the north and south of California's agricultural region is the types of crops that are profitability grown given downstream processing capacity. I include a linear time trend for the north region to capture potential differences in agro-economic trends. Results are essentially unchanged, as shown in column 4 of Appendix Tables 2.B.1 to 2.B.9.

My preferred specification uses grower fixed effects. An alternative is to use field fixed effects that would further absorb unobservable differences in field quality within each grower's farm. As shown in column 2 of Appendix Tables 2.B.1 to 2.B.4, the extensive margin results are robust to replacing grower fixed with field fixed effects. The intensive margin results are similar in magnitude but less precise (see column 2 of Appendix Tables 2.B.5 to 2.B.9). Recall that the intensive margin outcomes are for a subset of field observations growing processing tomatoes. Because fields are regularly rotated, a panel of processing tomatoes field-year observations is highly unbalanced. I therefore prefer to use grower fixed effects instead of field fixed effects.

My preferred specification weights observations by field acreage to produce estimates that are representative of California-wide effects. If estimates are sensitive to the inclusion of weights, it could indicate the model is misspecified. Results are similar with or without weights as shown in column 6 of Appendix Tables 2.B.1 to 2.B.9.

The conversion of acreage in to or out of tree crops reflects a more complex, dynamic decision about land use that isn't captured by the model. I test if the extensive margin results are robust to dropping observations of tree crops (6 percent of field-year observations). As reported in column 7 of Appendix Tables 2.B.1 to 2.B.4, estimates are slightly larger in magnitude but sign and significance are unchanged.

Table 2.B.1: Fallow, robustness checks

	(1) Preferred	(2) Field fe	(3) Water district fe	(4) Grower-district fe	(5) Quadratic time trend	(6) Regional trend	(7) Basin depth	(8) No weights	(9) No tree crops
Surface water supply	-0.0008***	-0.0007**	-0.0008***	-0.0008***	-0.0008***	-0.0008***	-0.0008***	-0.0006**	-0.0009***
	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0002)
Grower fixed effects	✓				✓	✓	✓	✓	✓
Year fixed effects	✓	✓	✓	✓		✓	✓	✓	✓
Weighted	✓	✓	✓	✓	✓	✓	✓		✓
Field fixed effects		✓							
Water district fixed effects			✓						
Grower-district fixed effects				✓					
Quadratic time trend					✓				
North trend						✓			
Basin depth							✓		
No tree crops									✓
Observations	36161	36157	36161	36159	36161	36161	36113	36161	33821

Notes: Each column shows the results from a separate regression model under different modelling assumptions for the outcome variable "Fallow", using the full panel of field-year observations from 2011-2021. The reported estimate is the effect of a 1 percentage point increase in surface water supply. Additional controls in all models include: groundwater depth, temperature exposure, precipitation, and major soil type. Standard errors (in parentheses) are two-way clustered by grower and district-year. Significance: * $p<0.05$, ** $p<0.01$, *** $p<0.001$.

Table 2.B.2: High-water crop, robustness checks

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Preferred	Field fe	Water district fe	Grower-district fe	Quadratic time trend	Regional trend	Basin depth	No weights	No tree crops
Surface water supply	0.0011***	0.0007*	0.0007**	0.0012***	0.0014***	0.0012***	0.0013***	0.0008**	0.0013***
	(0.0002)	(0.0003)	(0.0002)	(0.0003)	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0002)
Grower fixed effects	✓				✓	✓	✓	✓	✓
Year fixed effects	✓	✓	✓	✓		✓	✓	✓	✓
Weighted	✓	✓	✓	✓	✓	✓	✓		✓
Field fixed effects		✓							
Water district fixed effects			✓						
Grower-district fixed effects				✓					
Quadratic time trend					✓				
North trend						✓			
Basin depth							✓		
No tree crops									✓
Observations	36161	36157	36161	36159	36161	36161	36113	36161	33821

Notes: Each column shows the results from a separate regression model under different modelling assumptions for the outcome variable "High-water crop", using the full panel of field-year observations from 2011-2021. The reported estimate is the effect of a 1 percentage point increase in surface water supply. Additional controls in all models include: groundwater depth, temperature exposure, precipitation, and major soil type. Standard errors (in parentheses) are two-way clustered by grower and district-year. Significance: * $p<0.05$, ** $p<0.01$, *** $p<0.001$.

Table 2.B.3: Low-water crop, robustness checks

	(1) Preferred	(2) Field fe	(3) Water district fe	(4) Grower-district fe	(5) Quadratic time trend	(6) Regional trend	(7) Basin depth	(8) No weights	(9) No tree crops
Surface water supply	-0.0004	-0.0000	0.0002	-0.0004	-0.0006**	-0.0004	-0.0005*	-0.0002	-0.0005
	(0.0002)	(0.0003)	(0.0003)	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0003)
Grower fixed effects	✓				✓	✓	✓	✓	✓
Year fixed effects	✓	✓	✓	✓		✓	✓	✓	✓
Weighted	✓	✓	✓	✓	✓	✓	✓		✓
Field fixed effects		✓							
Water district fixed effects			✓						
Grower-district fixed effects				✓					
Quadratic time trend					✓				
North trend						✓			
Basin depth							✓		
No tree crops									✓
Observations	36161	36157	36161	36159	36161	36161	36113	36161	33821

Notes: Each column shows the results from a separate regression model under different modelling assumptions for the outcome variable "Low-water crop", using the full panel of field-year observations from 2011-2021. The reported estimate is the effect of a 1 percentage point increase in surface water supply. Additional controls in all models include: groundwater depth, temperature exposure, precipitation, and major soil type. Standard errors (in parentheses) are two-way clustered by grower and district-year. Significance: * $p<0.05$, ** $p<0.01$, *** $p<0.001$.

Table 2.B.4: Processing tomatoes, robustness checks

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Preferred	Field fe	Water district fe	Grower-district fe	Quadratic time trend	Regional trend	Basin depth	No weights	No tree crops
Surface water supply	0.0011***	0.0007*	0.0007**	0.0012***	0.0014***	0.0012***	0.0013***	0.0008**	0.0013***
	(0.0002)	(0.0003)	(0.0002)	(0.0003)	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0002)
Grower fixed effects	✓				✓	✓	✓	✓	✓
Year fixed effects	✓	✓	✓	✓		✓	✓	✓	✓
Weighted	✓	✓	✓	✓	✓	✓	✓		✓
Field fixed effects		✓							
Water district fixed effects			✓						
Grower-district fixed effects				✓					
Quadratic time trend					✓				
North trend						✓			
Basin depth							✓		
No tree crops									✓
Observations	36161	36157	36161	36159	36161	36161	36113	36161	33821

Notes: Each column shows the results from a separate regression model under different modelling assumptions for the outcome variable "High-water crop", using the full panel of field-year observations from 2011-2021. The reported estimate is the effect of a 1 percentage point increase in surface water supply. Additional controls in all models include: groundwater depth, temperature exposure, precipitation, and major soil type. Standard errors (in parentheses) are two-way clustered by grower and district-year. Significance: * $p<0.05$, ** $p<0.01$, *** $p<0.001$.

Table 2.B.5: Planting day of year, robustness checks

	(1) Preferred	(2) Field fe	(3) Water district fe	(4) Grower-district fe	(5) Quadratic time trend	(6) Regional trend	(7) Basin depth	(8) No weights
Surface water supply	0.0439***	0.0242	0.0339*	0.1388***	0.1319***	0.0437***	0.0452***	0.0338*
	(0.0117)	(0.0152)	(0.0146)	(0.0194)	(0.0167)	(0.0117)	(0.0116)	(0.0134)
Grower fixed effects	✓				✓	✓	✓	✓
Year fixed effects	✓	✓	✓	✓		✓	✓	✓
Weighted	✓	✓	✓	✓	✓	✓	✓	
Field fixed effects		✓						
Water district fixed effects			✓					
Grower-district fixed effects				✓				
Quadratic time trend					✓			
North trend						✓		
Basin depth							✓	
Observations	8609	7511	8606	8540	8609	8609	8598	8609

Notes: Each column shows the results from a separate regression model under different modelling assumptions for the outcome variable “Planting day of year”, using the subset of observations in the processing tomatoes data from 2011-2021. The reported estimate is the effect of a 1 percentage point increase in surface water supply. Additional controls in all models include: groundwater depth, temperature exposure, precipitation, and major soil type. Standard errors (in parentheses) are two-way clustered by grower and district-year. Significance: * p<0.05, ** p<0.01, *** p<0.001.

Table 2.B.6: Season length, robustness checks

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Preferred	Field fe	Water district fe	Grower-district fe	Quadratic time trend	Regional trend	Basin depth	No weights
Surface water supply	0.0006	0.0009	-0.0001	-0.0035	-0.0027	0.0006	0.0004	-0.0002
	(0.0014)	(0.0020)	(0.0019)	(0.0018)	(0.0015)	(0.0014)	(0.0014)	(0.0013)
Grower fixed effects	✓				✓	✓	✓	✓
Year fixed effects	✓	✓	✓	✓		✓	✓	✓
Weighted	✓	✓	✓	✓	✓	✓	✓	
Field fixed effects		✓						
Water district fixed effects			✓					
Grower-district fixed effects				✓				
Quadratic time trend					✓			
North trend						✓		
Basin depth							✓	
Observations	8609	7511	8606	8540	8609	8609	8598	8609

Notes: Each column shows the results from a separate regression model under different modelling assumptions for the outcome variable “Growing days”, using the subset of observations in the processing tomatoes data from 2011-2021. Results in column (5) are for 2016-2021. The reported estimate is the effect of a 1 percentage point increase in surface water supply. Additional controls in all models include: groundwater depth, temperature exposure, precipitation, and major soil type. Standard errors (in parentheses) are two-way clustered by grower and district-year. Significance: * $p<0.05$, ** $p<0.01$, *** $p<0.001$.

Table 2.B.7: Early variety, robustness checks

	(1) Preferred	(2) Field fe	(3) Water district fe	(4) Grower-district fe	(5) Quadratic time trend	(6) Regional trend	(7) Basin depth	(8) No weights
Surface water supply	-0.0008***	-0.0007*	-0.0006*	-0.0007*	-0.0008**	-0.0008***	-0.0008***	-0.0008***
	(0.0002)	(0.0003)	(0.0003)	(0.0003)	(0.0003)	(0.0002)	(0.0002)	(0.0002)
Grower fixed effects	✓				✓	✓	✓	✓
Year fixed effects	✓	✓	✓	✓		✓	✓	✓
Weighted	✓	✓	✓	✓	✓	✓	✓	
Field fixed effects		✓						
Water district fixed effects			✓					
Grower-district fixed effects				✓				
Quadratic time trend					✓			
North trend						✓		
Basin depth							✓	
Observations	8609	7511	8606	8540	8609	8609	8598	8609

Notes: Each column shows the results from a separate regression model under different modelling assumptions for the outcome variable "Early variety", using the subset of observations in the processing tomatoes data from 2011-2021. The reported estimate is the effect of a 1 percentage point increase in surface water supply. Additional controls in all models include: groundwater depth, temperature exposure, precipitation, and major soil type. Standard errors (in parentheses) are two-way clustered by grower and district-year. Significance: * p<0.05, ** p<0.01, *** p<0.001.

Table 2.B.8: Drip irrigation, robustness checks

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Preferred	Field fe	Water district fe	Grower-district fe	Quadratic time trend	Regional trend	Basin depth	No weights
Surface water supply	-0.0005*	-0.0002	0.0001	0.0000	-0.0002	-0.0005*	-0.0006*	-0.0006*
	(0.0002)	(0.0001)	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0002)
Grower fixed effects	✓				✓	✓	✓	✓
Year fixed effects	✓	✓	✓	✓		✓	✓	✓
Weighted	✓	✓	✓	✓	✓	✓	✓	
Field fixed effects		✓						
Water district fixed effects			✓					
Grower-district fixed effects				✓				
Quadratic time trend					✓			
North trend						✓		
Basin depth							✓	
Observations	8609	7511	8606	8540	8609	8609	8598	8609

Notes: Each column shows the results from a separate regression model under different modelling assumptions for the outcome variable “Drip irrigation”, using the subset of observations in the processing tomatoes data from 2011-2021. The reported estimate is the effect of a 1 percentage point increase in surface water supply. Additional controls in all models include: groundwater depth, temperature exposure, precipitation, and major soil type. Standard errors (in parentheses) are two-way clustered by grower and district-year. Significance: * $p<0.05$, ** $p<0.01$, *** $p<0.001$.

Table 2.B.9: ln(irrigation), robustness checks

	(1) Preferred	(2) Field fe	(3) Water district fe	(4) Grower-district fe	(5) Quadratic time trend	(6) Regional trend	(7) Basin depth	(8) No weights
Surface water supply	0.0001	0.0001	-0.0001	-0.0002	-0.0002	0.0001	0.0001	0.0003
	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0002)
Grower fixed effects	✓				✓	✓	✓	✓
Year fixed effects	✓	✓	✓	✓		✓	✓	✓
Weighted	✓	✓	✓	✓	✓	✓	✓	
Field fixed effects		✓						
Water district fixed effects			✓					
Grower-district fixed effects				✓				
Quadratic time trend					✓			
North trend						✓		
Basin depth							✓	
Observations	4896	3945	4892	4848	4896	4896	4885	4896

Notes: Each column shows the results from a separate regression model under different modelling assumptions for the outcome variable "ln(irrigation)", using the subset of observations in the processing tomatoes data from 2016-2021. The reported estimate is the effect of a 1 percentage point increase in surface water supply. Additional controls in all models include: groundwater depth, temperature exposure, precipitation, and major soil type. Standard errors (in parentheses) are two-way clustered by grower and district-year. Significance: * $p<0.05$, ** $p<0.01$, *** $p<0.001$.

Essay 3

Climate, Weather, and Collective Reputation: Implications for California's Wine Prices and Quality

3.1 Introduction

Wine is the most differentiated of all farm products, with much of the differentiation based on the combination of wine grape varieties and so-called "terroir"—reflecting the soil type, topography, and climate in particular. Local climate determines the types of grapes that can be suitably grown while weather variation introduces vintage-to-vintage quality differences. Reflecting this product differentiation, prices of wine and the grapes used to produce it vary considerably. Prices of grapes of the same variety produced in the same region in the same vintage year can vary by a factor of 50 (Sambucci & Alston, 2017). Prices of wine and wine grapes also vary over time among vintage years: beneficial weather results in higher-quality wine grapes, yielding higher-quality wine that fetches a higher price. Wine produced by the same winemaker using grapes grown on the vineyard can vary from year to year by a factor of 20 or more (e.g., Ashenfelter (2010)).

This often largely uncontrolled variation in quality and prices adds to the asymmetric information or "lemons" problem that is pervasive in the market for fine wine (e.g., Livat et al. (2019)). Geographic indications (GIs) for wine were first introduced 100 years ago to address this problem in France (Mérel et al. (2021)). In the United States the counterpart American Viticultural Areas (AVAs) were introduced in 1980, enabling U.S. wine producers to label wine as coming from a specific AVA to exploit the "collective reputation" associated with that region of production. The purpose of associating a particular wine with an AVA is to create and capture price premia. Some studies have reported evidence on the value of collective reputation for wine associated with GIs (e.g., Winfree & McCluskey (2005)). However, relatively little is known about the complex relationships between prices and appellations for wine in the context of variable weather and a changing climate, and formal evidence is scant.

The objective of this study is to analyze the role of AVAs in mediating the relationship between (1) an evolving climate (the long-run expected weather in a region), (2) weather variation around the regional norm (vintage effects), and (3) the variety-specific price premia and quality (expert rating scores) for varietal wines in different parts of California. Our analysis is based on a sample of premium wines that were rated by Wine Spectator magazine between 1994 and 2022.

The primary contributions of our work pertain to quantifying the variety-specific relationships between measures of wine quality (i.e., prices and rating scores) and measures of climate and weather variation relative to climate, while controlling for appellations. We uncover complex variety-specific relationships between climate, weather, and wine prices and quality across California's diverse production regions. For Cabernet Sauvignon wines, the relationship between price and climate follows an inverted U shape where wines produced in cooler and warmer regions are discounted relative to the apparent optimum—an average daily temperature of around 18.45°C during the growing season. On the other hand, prices for Chardonnay wines are generally higher in cooler regions like Sonoma and the Central Coast, with a negative relationship between temperature and price across the range of our data. For Chardonnay, the optimum average daily temperature during the growing season may be 17°C or less, and California's growing conditions may be generally too warm. Within regions, temperatures warmer than the regional climate resulted in vintages with lower prices and rating scores for both Cabernet Sauvignon and Chardonnay wine.

A further contribution of our work is to develop an improved understanding of the role of AVAs in shaping the relationship between climate, weather, and perceived quality of wine in California. We find both Cabernet Sauvignon and Chardonnay wines are discounted when they are labeled as originating from "macro-regions" (such as California, the North Coast, or Sonoma County) rather than from smaller California AVAs that by definition lie within macro-regions in the state—as previously reported by Bombrun & Sumner (2003) and Kwon et al. (2008). We also find that prices and ratings of wines from premier AVAs are less influenced by deviations in temperature from AVA-specific climate. This is consistent with the notion that producers of high-value grapes and wines intervene in the vineyard and the winery to mitigate the effects of weather on the quality of their premium branded products–possibly by reallocation of lower-quality products to lower-quality brands–while producers of more generic grapes and wine do not.

A third contribution of this study is our measurement of key weather and climate variables. Although weather and climate can vary over relatively small distances in wine growing regions in California, previous work on wine in California uses coarsely measured weather and climate data: for example, Ramirez (2008) and Jones et al. (2005) use one weather station for the whole of the Napa Valley and all 16 AVAs within. The resulting imprecision in the measurement of weather and climate variables for the subregions adds to challenges in estimating the true relationship between weather, climate, and wine characteristics. We use spatially detailed weather data from PRISM (800m grids) to more-accurately represent the relevant concepts of weather and climate and quantify their effects on California's wine quality and price, reducing the risk of measurement error bias.

3.2 The Setting

Wine grapes are a long-lived perennial crop with a typical productive life of 25 years, and often longer. California wine grapes were valued at $3.1 billion in 2020-21, making wine grapes the second-most valuable crop in California after almonds (USDA NASS, 2021). California produces around 80 percent of total U.S. wine by volume (Wine Institute, 2023) in several distinct wine production regions (see Figure 3.1). These regions differ in terms of terrain, climate, and soil type, which drives differences in the grape varieties grown and the quality of grapes and wine produced. In the warm Southern Central Valley, wine grape production is typically high yield per acre and relatively low value per ton. The cooler areas near the coast are associated with smaller-scale production of higher-value premium wine grapes. The Napa-Sonoma region on the north coast is especially known for Cabernet Sauvignon, which is its most important variety and increasingly so, while the Central Coast region is known for cooler climate Chardonnay and Pinot Noir (Alston et al. (2015); Alston & Sambucci (2019)).

3.2.1 Geographic indications for wine

In 1980 the U.S. Government created American Viticultural Areas (AVAs) as a mechanism for producers to signal quality and better capture the benefits from collective reputation associated with the location of production (see U.S. Treasury/TTB (2022); Winfree & McCluskey (2005); and Lapsley et al. (2019)). AVAs are defined geographic areas that may be quite large and cross state or county lines, or may be quite small and lie within a county or, in some cases, another AVA. In 2021, the United States had a total of 258 established AVAs of which 142 are in California and 16 are nested within the Napa Valley AVA (see U.S. Treasury/TTB (2021)).

Wineries may label a wine as coming from an AVA if 85 percent of the grapes were grown in the AVA and the wine was fully finished in the state where the AVA is located. The use of an AVA label does not impose restrictions on production or winemaking practices,

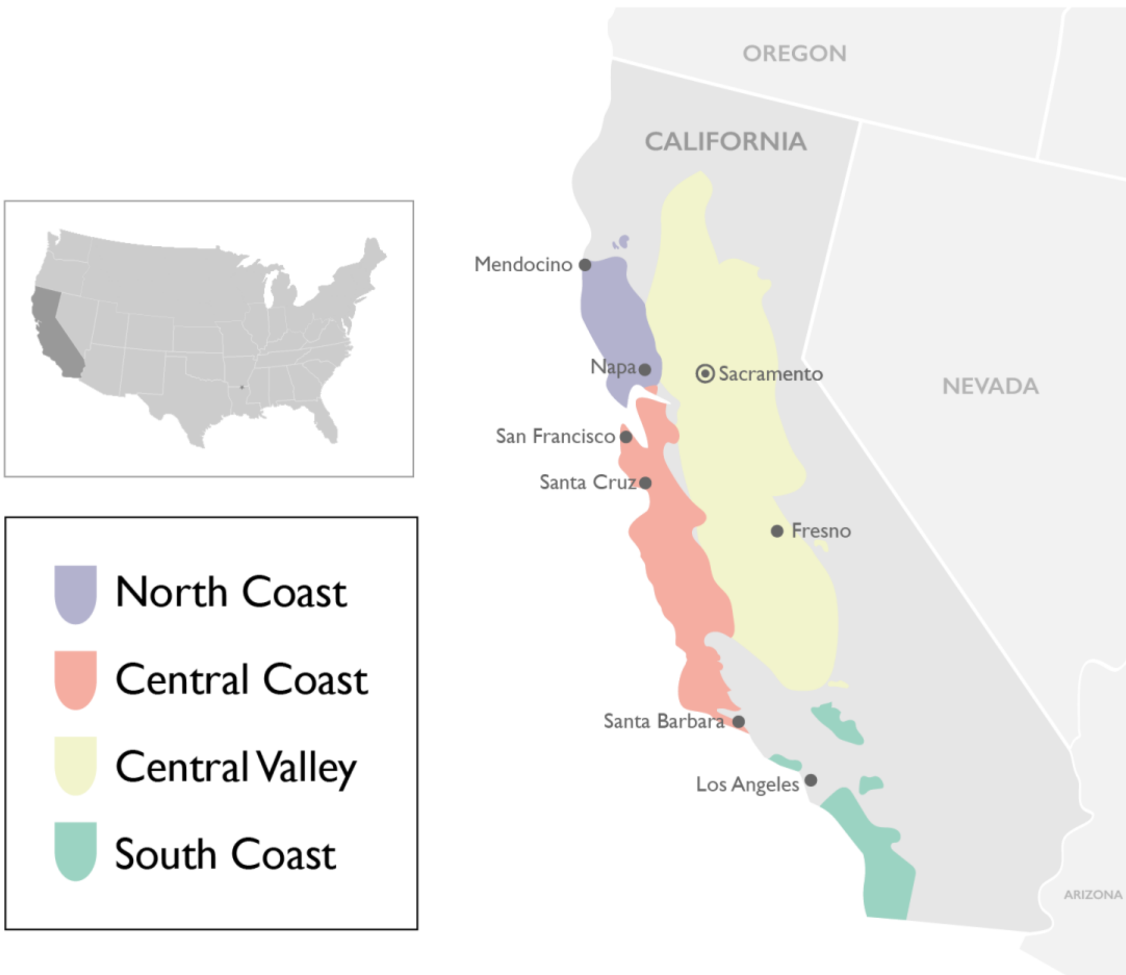

Figure 3.1: Map of main wine regions in California

Source: Getty Images.

unlike geographic indicators for wine in some other countries (for example, the Appellation d'Origine Contrôlée for French wines; see, e.g., Alston & Gaeta (2021)). California's wine can also be labelled as originating from a particular county or the state of California. For a wine to carry a county name on its label, at least 75 percent of the grapes must have been grown in that county and the wine must have been fully finished within California. A wine labeled with "California" must be made entirely using grapes from California and finished within California (U.S. Treasury/TTB, 2020).

Prices of wine grapes vary considerably among and within AVAs, even within California (see, e.g.,Alston et al. (2015); Sambucci & Alston (2017)). In the 2021 California Grape Crush Report (California Department of Food and Agriculture, 2022), prices of lots of Cabernet Sauvignon grapes from Napa County (crush district 4) ranged from as low as $600 per ton up to $62,000 per ton. Differentiation occurs along several dimensions, including wine grape varieties, terroir, vineyard management and production practices, and fruit quality (e.g., sugar content).

As described by, say, Sambucci & Alston (2017), considerable quantities of California wine grapes are vinified by the growers and not sold as such (62% of Napa-Sonoma tons crushed were sold, 38% not sold). For the wine grapes that are sold, growers often contract with wineries for the sale of grapes, particularly among growers of high-quality grapes (Goodhue et al. (2003); Franken (2014)).

Variation in wine grape prices ultimately reflects variation in the anticipated value of the wine they will be used to make, since the demand for grapes is derived from the demand for wine. The winemaking process potentially introduces additional variation in final wine quality and price. A wineries' individual reputation may also play a role in price formation. For cheaper wines, price premia are more likely a consequence of collective reputation, inferred from an AVA label, rather than firm-level reputation. For more expensive wines, the premium for an individual winery's reputation is likely to be a larger component of price (Costanigro et al., 2010).

Generally speaking, as discussed by Alston & Gaeta (2021), wine prices and ratings tend to increase as we go from broader (e.g., entire country, state, or broad region within a state), to narrower and more specific sub-regions of origin (such as north coast, or within that, Napa Valley and its sub-AVAs). For example, Bombrun & Sumner (2003) report that, after controlling for observable wine characteristics, wines using the Napa Valley AVA command a price premium over wines labeled as from "California," and some sub-AVAs like Oakville and Howell Mountain capture even larger premia; see, also, Kwon et al. (2008).

The prime purpose of creating sub-AVAs is to create and capture such premia. Hence, everything else equal, we should expect wine labeled as coming from one of the 16 sub-AVAs within the Napa Valley AVA to command a price premium over wine labeled as coming from the Napa Valley AVA (if it were not the case, winemakers might as well opt to use the broader Napa Valley AVA over the sub-AVA on the label).

3.2.2 Links between climate change and wine quality

Ashenfelter & Storchmann (2016) and Jones et al. (2022) provide summaries of main points from previous work on climate change and wine, including potential adaptation strategies.

It has long been understood that weather affects wine grape characteristics like color and acid (Winkler, 1962). Grape characteristics are inputs to wine quality and can affect the final wine's color, aroma, tannins, and other flavor attributes. Grape varieties may be characterized by an optimal temperature range–the lower limit delineates the point where grapes will ripen, and the upper limit describes the point at which grapes will be overripe or damaged. Figure 3.2 from Jones et al. (2012) illustrates the range of average growing season temperatures that are optimal for each of the world's most common winegrape cultivars.

In several papers the authors find that weather during the vintage year causes significant variation in bottled wine prices: Ashenfelter et al. (1995) and Ashenfelter (2010) for Bordeaux wines; Byron & Ashenfelter (1995) for Australian wines; Haeger & Storchmann (2006) and Ramirez (2008) for US wines. Variations in wine quality can be linked to various aspects of weather including exposure to high temperatures, number of frost days, diurnal temper-

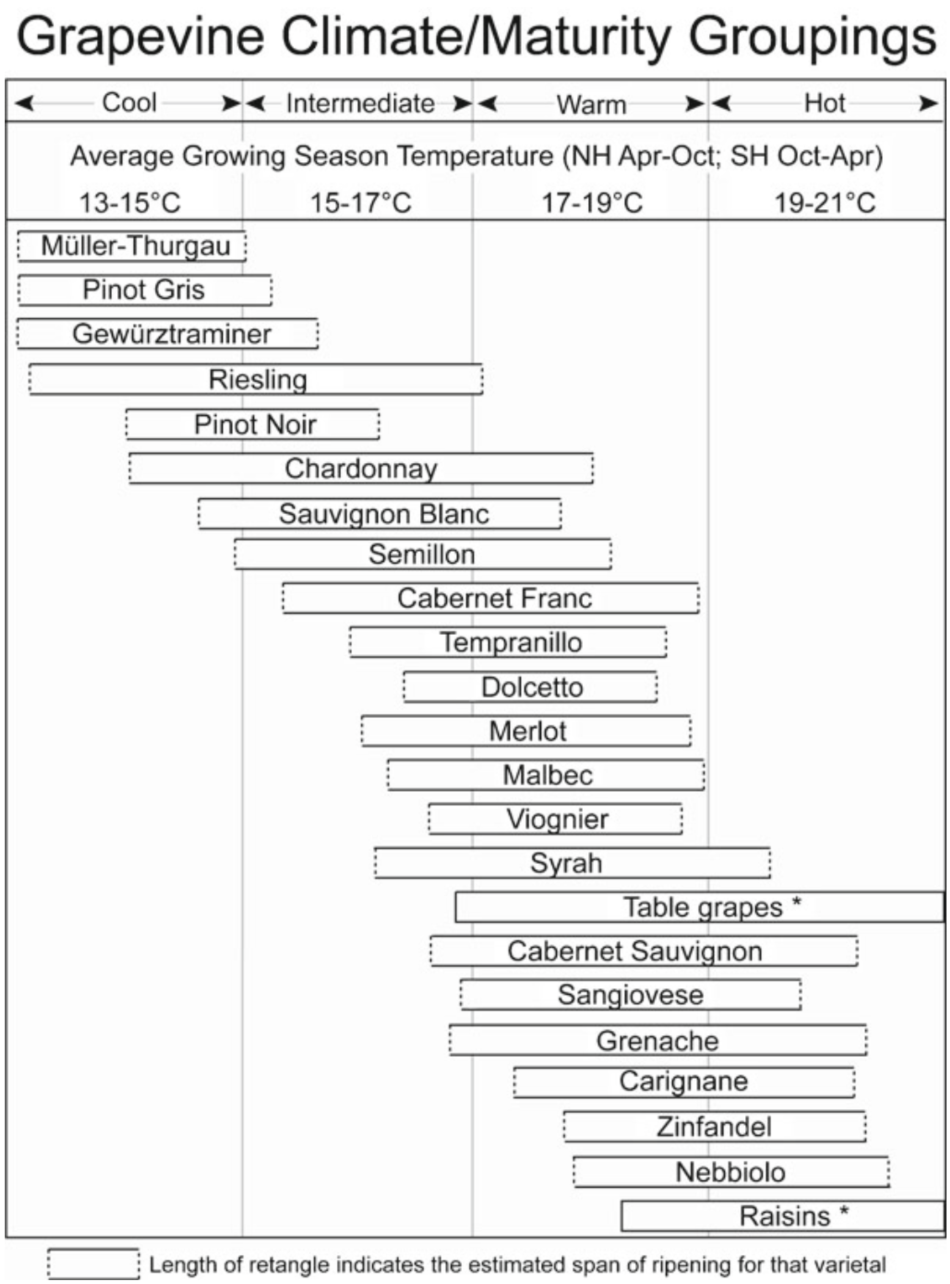

Figure 3.2: Optimal average growing season temperature range by grape variety

Source: Jones et al. (2012)
Note: The original caption in Jones et al. (2012) reads "Climate-maturity groupings based on relationships between phenological requirements and growing season average temperatures for high- to premium-quality wine production in the world's benchmark regions for many of the world's most common cultivars. The dashed line at the end of the bars indicates that some adjustments may occur as more data become available, but changes of more than $\pm 0.2 - 0.5°C$ are highly unlikely (Jones, 2006)."

ature range, rainfall, and degree day accumulation (Jones et al. (2005); Jones & Goodrich (2008); Davis et al. (2019)). Timing of weather events also matters. For example, in the case of Burgundy wines, rainfall is beneficial to quality if it occurs during the bud break period but detrimental if it occurs during the ripening phase (Davis et al., 2019).

Alston et al. (2011) found that increased heat during the growing season contributed to a statistically significant but small increase in sugar content of wine grapes grown in California over the period 1990-2008. However, rather than a climate effect, they concluded that most of the observed upward trend in sugar content (and associated increases in alcohol percentage) must be attributed to other factors including changes in vineyard management such as longer hang times. Using their estimated model parameters, even a substantial rise in average temperatures would have had only a modest effect on sugar content of wine grapes.

Since the 1950s, grape growing regions in California have experienced warmer growing seasons on average caused by a changing climate. Many regions have experienced an increase in minimum (i.e., overnight) temperatures, which has reduced the occurrence of frost days (Jones, 2004). Nemani et al. (2001) and Gambetta & Kurtural (2021) suggest that wine quality in California appears to have largely benefitted from this warming. However, warming trends have coincided with notable trends in the supply of and demand for wine that make it difficult to disentangle the effect of warming from other factors. Technological advancements, better plant material, and improved vineyard management have allowed producers to create more consistent, high-quality wine (Jones et al., 2005). Demand for wine has also grown over the past 50 years due to a larger and richer wine drinking population. Premium wine has particularly seen large increases in demand because of consumer's shifting preferences toward high-quality products (Anderson et al., 2018).

Wine grape yields can also be affected by weather. Exposure to temperatures at the extreme (i.e., frost or extreme heat) has been linked to lower yields, while moderate temperatures particularly overnight are beneficial to yields (Cahill et al. (2007); Lobell et al. (2006); White et al. (2006))

Producers can potentially mediate the effect of weather (high temperatures) and climate change (rising temperatures) in several ways. One (longer-run) response is to relocate wine grape production from warm regions to cool regions, such as towards the poles or to higher elevation areas. Several studies predict a decline in areas of vineyards acreage in key production regions (for example, southern Europe) because the regions are projected to become too hot to produce quality wine (Moriondo et al. (2013); Hannah et al. (2013); Webb et al. (2007)). However, these studies generally underestimate or ignore adaptive responses that may help preserve production in wine-growing regions that are currently culturally and economically important.

Grape varieties are diverse in their phenology and other traits related to climate and weather. As climate changes, growers may take advantage of the large amount of varietal diversity and plant a different variety that is more suited to their new climate (Wolkovich et al., 2018). However, despite the availability of more than 1,000 commercial varieties, most wine grape regions grow the same 12 varieties. In fact, the mix of wine grape varieties is becoming less differentiated in the United States, especially in California (Alston et al., 2015) and Australia (Puga et al., 2022), and instead these regions are becoming more similar to France and the rest of the world as a whole. Traditional French varieties such as Cabernet Sauvignon, Chardonnay, Merlot, Sauvignon Blanc, Pinot Noir, and Syrah (or Shiraz) are increasingly predominant in California in places that are becoming increasingly less-favored for growing those varieties. California growers have been very slow to adopt varieties from Italy or Spain that may be better suited to hotter places. Varietal adaptation in California and elsewhere is hampered by the long productive life of vineyards as well as historical association of high-quality wine from particular regions with particular varieties. This is particularly pronounced in many European regions where a wine cannot bear a geographic indication as coming from a specific PDO unless it is made using particular varieties or blend of varieties that are permitted by the PDO rules.

Changing the location of production or varieties grown can be seen as long-run, disruptive responses that essentially forsake the established identity of production that reflects the association of particular wines produced in particular places using particular varieties—at the terroir-varietal-GI nexus. Other, shorter-run responses can be undertaken seeking to preserve that identity. Specifically, producers can manage weather shocks (or trends) by adapting their growing practices, such as harvest date, canopy structure, and irrigation. Webb et al. (2021) found that smaller damages from 2009 heatwave in South-Eastern Australia were associated with (a) irrigation prior to the heatwave event, and (b) good canopy growth that protects fruit from direct radiation. Other adjustments can be made in the winery.

3.3 Data

We compiled data on prices and expert rating scores for California's wines from the Wine Spectator magazine and matched these to relevant measures of weather and climate from PRISM. Figure 3.3 summarizes the key datasets and how they were merged.

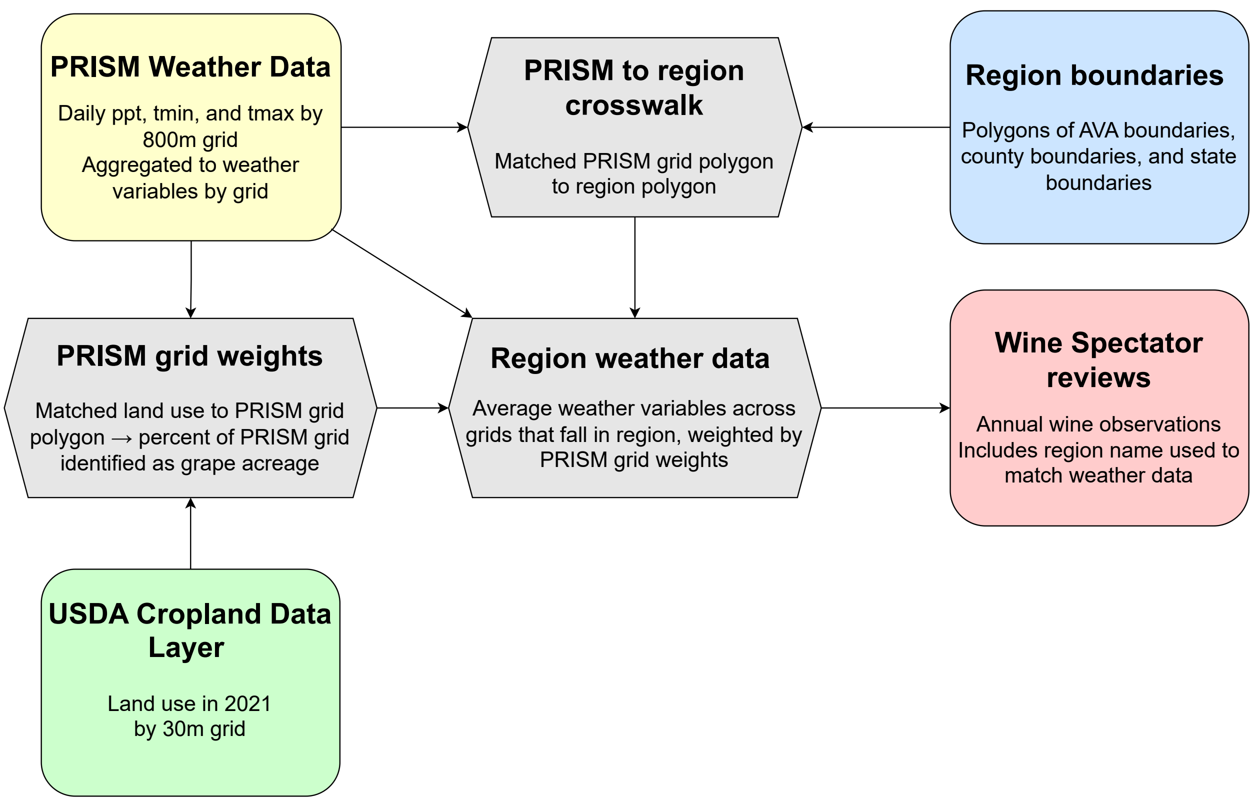

Figure 3.3: Diagram of key datasets and links between datasets

3.3.1 Wine data

The Wine Spectator (WS) magazine publishes information on recommended retail prices, expert ratings, and other information about many wines from around the world in each of its monthly issues; WS editors' blind taste and rate over 15,000 wines per year.

We collected information on wines from California published in the WS between January 1994 and December 2022. For each wine, we recorded its brand or producer, region (including AVA), vintage year, rating, suggested retail price, wine grape variety, wine type, and number of cases made. We focused on five grape varieties: Cabernet Sauvignon (27 percent of wine observations), Chardonnay (23 percent), Merlot (9 percent), Pinot Noir (26 percent), and Zinfandel (15 percent). Vintage is the year in which the grapes used to produce the wine were grown. We kept data on vintages between 1991 and 2020, with other vintages being too infrequently sampled to be included in our analysis.

Across the 28 years of WS magazines from which we collected data, some price variation reflects changes in the purchasing power of money. We converted suggested retail prices into equivalent 2022-dollar values using the Consumer Price Index (CPI) for the corresponding issue year (the year in which the wine rating was published by the WS) (U.S. Bureau of Labor Statistics (2023), specifically, the annual average of the CPI for all urban consumers, series number CUUR0000SA0). The average suggested retail price is $66 per bottle in our sample. These wines are high priced compared with California wines generally, and compared to the range of wines produced within premium brand and regions that they predominately represent.

Wine ratings reported by Wine Spectator magazine are ostensibly on a scale of 0–100 points but in practice for premium wines the typical range is 85–95 points, with exceptional wines scoring more than 95 points. Wines are rated blind, meaning information about the winery or wine (including its price) is unknown to the taster during the tasting.

In our sample, wine scores increased from an average of 85 point in 1991 to almost 90 points in 2020, and variation around the average declined over the decades shown in Figure

3.4. Wine prices also increased in real terms from \$35 per bottle in 1991 to more than \$80 per bottle by the mid-2010s (all in 2022 dollars). These trends in the complete sample are reflected in scores and prices for each variety (displayed in Appendix 3.A).

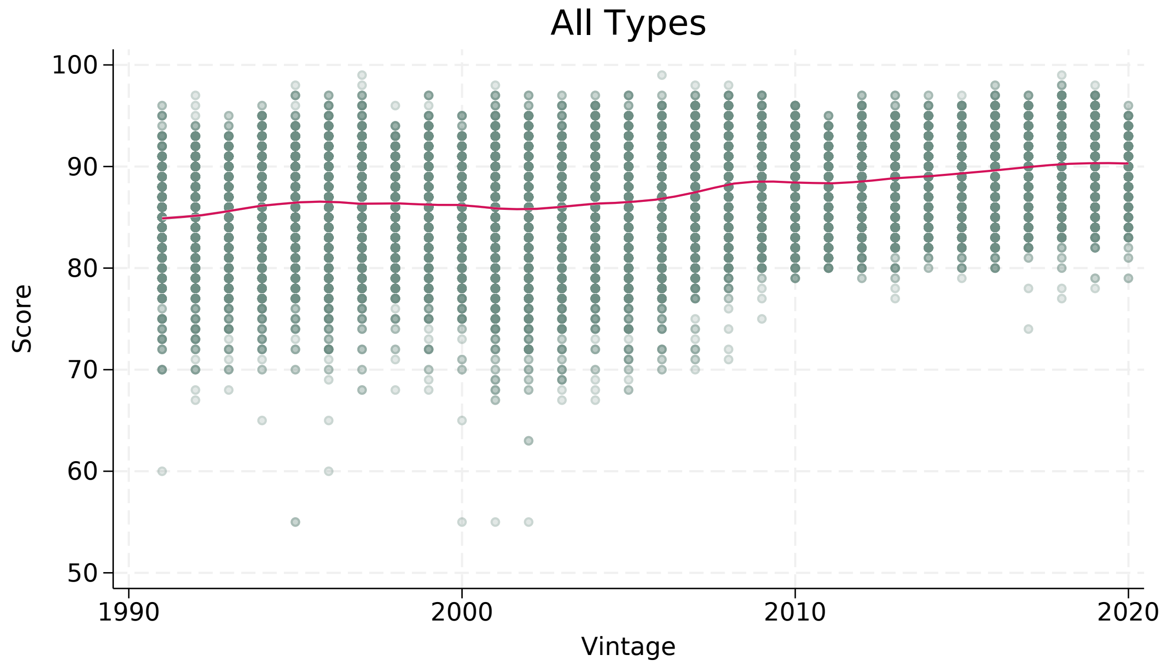

(a) Score

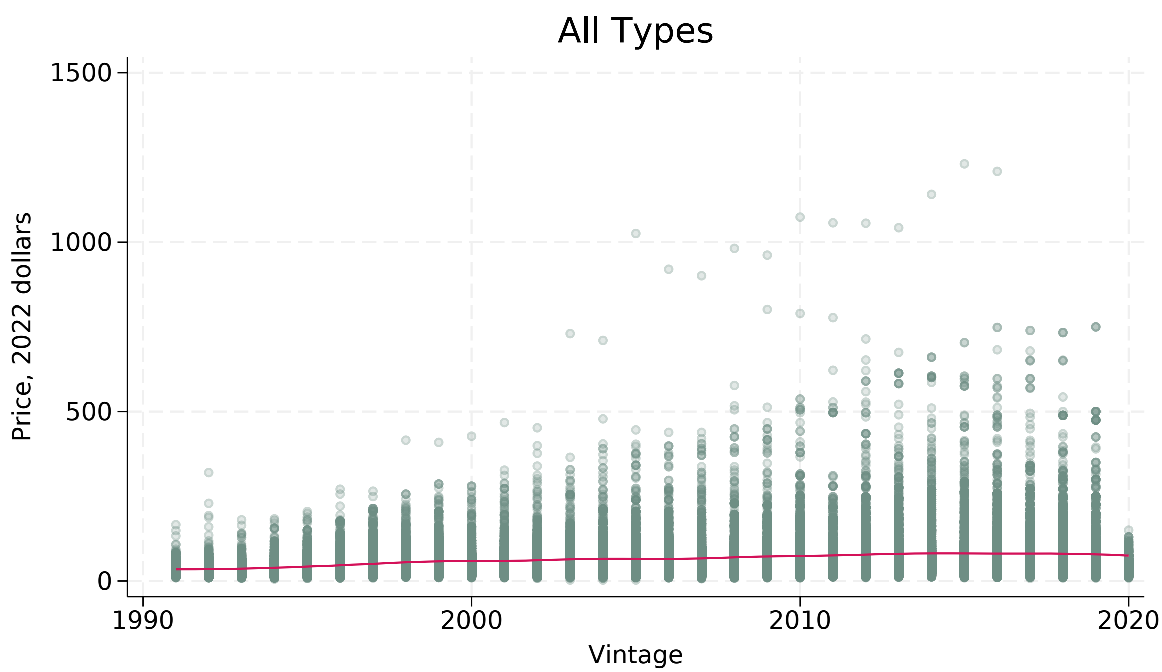

(b) Price (2022 dollars per bottle)

Figure 3.4: Wine scores and prices by vintage, all types

We defined a wine's region as the narrowest and most specific region that could be identified on the label, as reported by Wine Spectator magazine. For example, a wine from the Oakville AVA could be associated with California, Napa County, Napa Valley AVA, and Oakville AVA simultaneously, since the Oakville AVA is found within the Napa Valley AVA, which is in Napa County, California. In this example we defined the wine's region as Oakville AVA.

3.3.2 Wine regions

AVA boundaries were taken from the American Viticultural Areas Digitizing Project Team (2021) produced by UC Davis Library and UC Davis DataLab. They publish "spatial data from each of official American Viticultural Areas boundary descriptions which are accepted and published by the Alcohol and Tobacco Tax and Trade Bureau." We also used the definitions of county and state boundaries from the U.S. Census Bureau (2026).

We classify a region as a "macro-region" if it is either (a) the state of California, (b) a county in California, or (c) an AVA that contains at least one other AVA within its boundary. We classify the remaining regions as sub-AVAs: an AVA that do not contain another AVA within its boundary. See Appendix 3.B for a full list of the regions we analyzed and if they were categorized as a macro-region.

3.3.3 Weather data

We used spatially detailed weather data from PRISM (PRISM Climate Group, Oregon State University, 2020) to accurately represent regional weather and climate. PRISM interpolates daily minimum and maximum temperatures and precipitation to 800m-by-800m grids, taking into account elevation, coastal proximity, and aspect. Matching weather data (daily by 800m grid) to wine data (vintage by region) requires both temporal and spatial aggregation. Starting with temporal aggregation, we used two different approaches to aggregate daily observations into measures of weather and climate variables that may meaningfully influence wine price and quality.

The first approach is to average daily average temperatures over the key growing months

(April to October). This method is widely used in studies that link weather to agricultural outcomes because average temperature is easy to measure and interpret, however it hides potentially large differences in exposure to extreme temperatures. For example, a normal season and an abnormal season will have the same average temperature if abnormally hot and cold temperatures average each other out.

The second approach, degree days, is a concept pioneered by Snyder (1985) and popularized by Schlenker & Roberts (2009) (see Ortiz-Bobea (2021) for a summary). Degree days measure how long and by how much temperatures exceed the lower temperature bound while being below an upper temperature bound. We define growing degree days as degree days with bounds between 10°C and 35°C, summed from April to October by vintage year. The benefit of this approach is that it takes into account the distribution of temperatures throughout the day.

We calculated regional climate as the static 40-year average of regional weather variables from 1981 to 2020. An alternative definition of climate is the 10-year moving average of the weather variables, which we included in robustness checks. We chose the 10-year horizon for the practical reason of PRISM having 800m gridded data back only to 1981, 10 years before our WS data begin.

To spatially aggregate PRISM grids to wine regions, we first identified every 800m PRISM grid that intersects with each wine region using the boundaries described in Section 3.3.2. For regions that intersect with multiple PRISM grids, we calculated a single observation for the region by taking a weighted average of weather and climate variables across grids. Each grid's weight is equal to the share of the region's grape acreage within the grid in 2021, calculated using the USDA's Cropland Data Layer (USDA National Agricultural Statistics Service, 2022). Grids that were not associated with any grape acreage were assigned a weight of zero and did not contribute to the weighted average.

Figure 3.5 shows average temperatures during the growing season (April to October) over the past 40 vintages in a selected group of wine-growing regions. Average temperatures

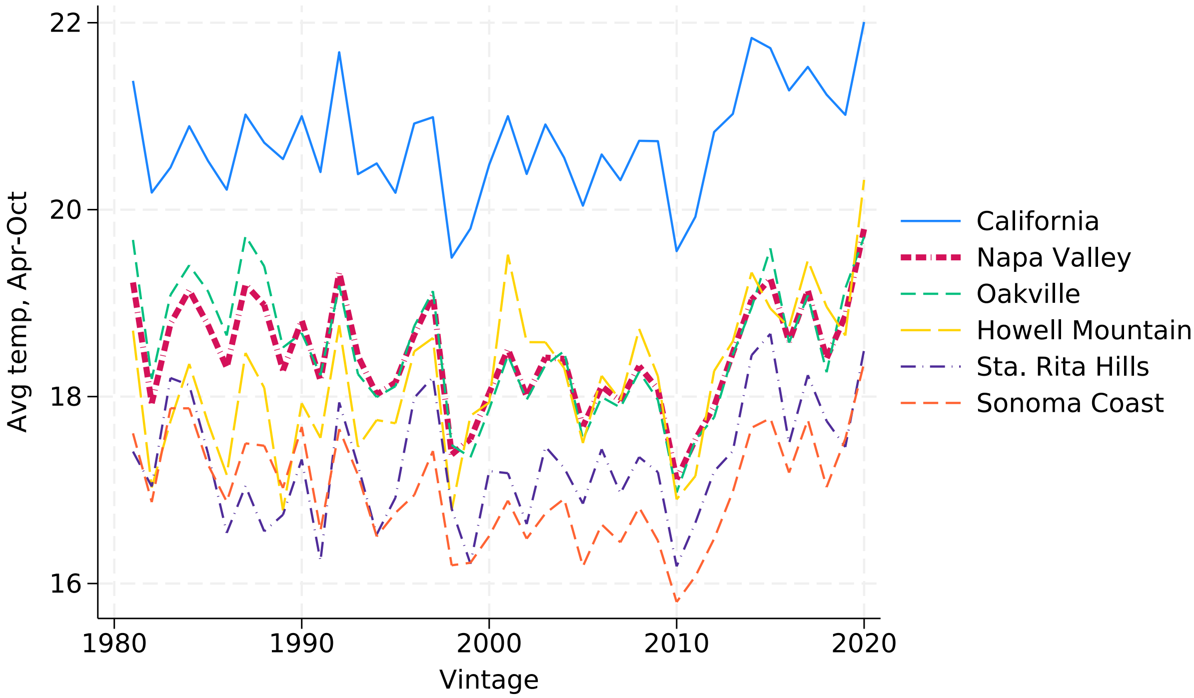

Figure 3.5: Average temperature (April to October) by vintage and region

are correlated across regions which indicates spatial dependence in weather: if one region of California is unusually warm during a growing season then it is likely that other regions in California are unusually warm too. However, the degree of correlation varies. Both Oakville and Howell Mountain lie wholly within the Napa Valley and yet Howell Mountain's weather can be quite different owing to its elevation.

Table 3.1 includes summary statistics on the price, score, cases made, weather, and proportion of the 53,930 observations associated with the various grape varieties and regions. Appendix 3.C shows average price, score, and cases made by wine grape variety and region.

Table 3.1: Summary statistics

	units	mean	sd	min	max
Price	2022 dollars	66.46	57.41	7	1231
Score	points	87.70	3.88	55	99
Cases made	no.	7939.58	40037.48	11	2600000
Vintage	year	2006.68	7.99	1991	2020
Weather					
Average temperature, annual average	° C	14.74	0.83	10	20
Minimum temperature, annual average	° C	7.28	0.98	2	14
Maximum temperature, annual average	° C	22.21	1.07	17	27
Precipitation, annual total	mm	779.93	361.50	42	2430
	count	percent of total observations			
Grape variety					
Cabernet Sauvignon	14,786	27			
Chardonnay	12,972	23			
Merlot	4,826	9			
Pinot Noir	14,459	26			
Zinfandel	8,353	15			
Grape location					
North Coast AVAs	38,195	69			
Napa Sub-AVAs	8,187	15			
Central Coast and Santa Cruz AVAs	8,275	15			
Sierra Foothills AVAs	349	0.6			
Central Valley AVAs	400	0.7			
South Coast AVAs	61	0.1			
County	5,621	10			
California State	2,495	4			
Total observations	53,930				

3.4 Climate and California's wine quality and prices

Here, we explore the associations between climate and California's wine. The causal effect of climate and climate change on wine price and score is difficult to disentangle from other factors. Climate varies in the cross-section along with other time invariant characteristics of regions, such as soil type and typography. Climate change, manifesting in warming trends, has coincided with the uptake of improved vineyard practices and technologies that enable producers to generate more consistent and high-quality product, as well as an upward trend in consumer demand for premium wine and other shifts in demand.

Figure 3.6 to Figure 3.9 show the relationship between climate and wine prices and scores for Cabernet Sauvignon and Chardonnay wines. Here, we define climate as the 40-year average of growing season average temperatures from 1981 to 2020. Figure 3.6 and Figure 3.7 show the average wine score and price by climate and region, where the size of each region's bubble corresponds to the number of observations in the WS data associated with that region. Figure 3.8 and Figure 3.9 show the distribution of wine scores and prices for different climate intervals. We chose interval cut offs so that each interval had roughly the same number of observations.

As expected, both Cabernet Sauvignon and Chardonnay wines were discounted if they were labeled as coming from a macro-region (colored blue in Figure 3.6 and Figure 3.7). Recall a macro-region is one that contains another wine region within its boundaries, such as California, Paso Robles, Santa Barbara County, Central Coast, or Napa Valley AVA (refer to Section 3.3.2 and Appendix 3.B for details).

For Cabernet Sauvignon wines in Figure 3.6, there is a distinctive cluster of wines from Sonoma including Sonoma Valley, Sonoma County, and sub-AVAs within such as Dry Creek Valley and Alexander Valley. The upper cluster is all places within the Napa Valley, hotter than Sonoma. There appears to be little relationship between wine prices and climate among the cluster of regions within the Napa Valley. This supports the idea that growers in these places have adapted to their local climate to produce high-quality Cabernet Sauvignon wine.

Leaving out the Napa Valley AVAs, with their high-quality wines, there is a generally negative relationship between price and heat across the rest of California—for middle of the road Cabernet Sauvignon. Pooling observations across all regions in Figure 3.8, the relationship between price and climate follows an inverted U shape for Cabernet Sauvignon.

In contrast, for Chardonnay wines in Figure 3.7, prices are generally higher in cooler places, especially for wines from Sonoma and parts of the Central Coast. While at any given temperature there is a broad range of prices at cooler temperatures, there is a generally negative relationship between temperature and price. This is confirmed in the distributions of wine scores and prices in Figure 3.9, which shows that Chardonnay wines from places with warmer climates tend to be priced and scored lower.

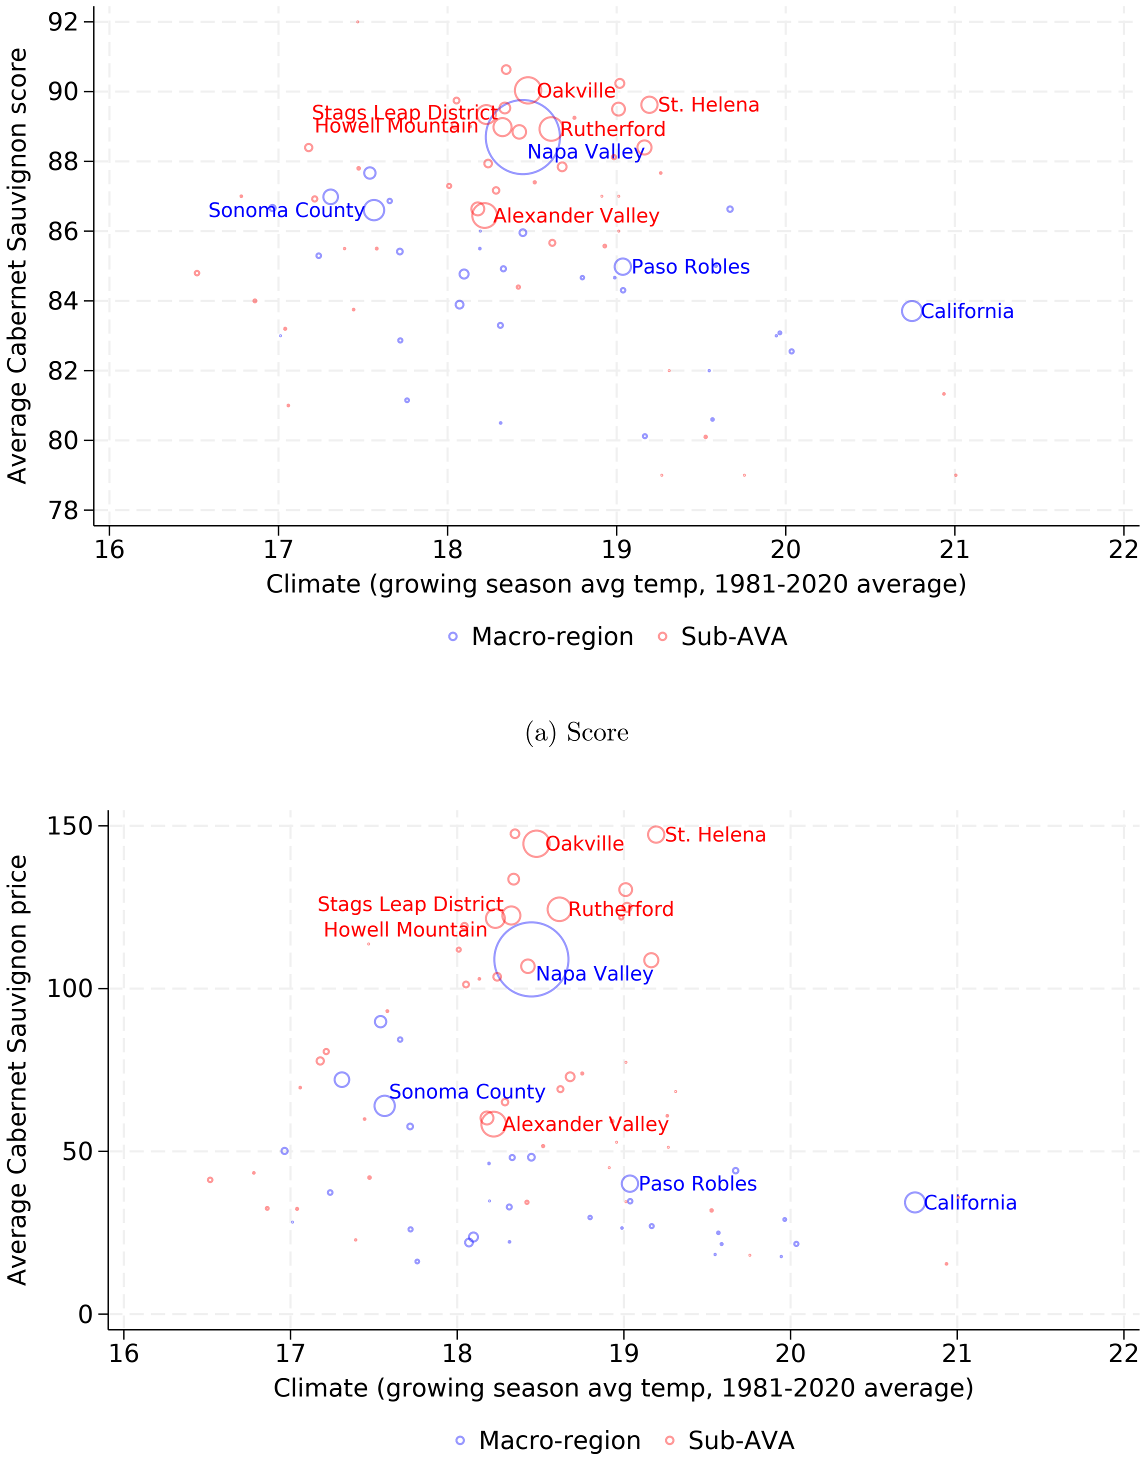

(a) Score

(b) Price (2022 dollars per bottle)

Figure 3.6: Average Cabernet Sauvignon wine scores and prices by climate and region

Note: The size of the bubble corresponds to the number of observations. We label the 10 regions with the largest number of observations.

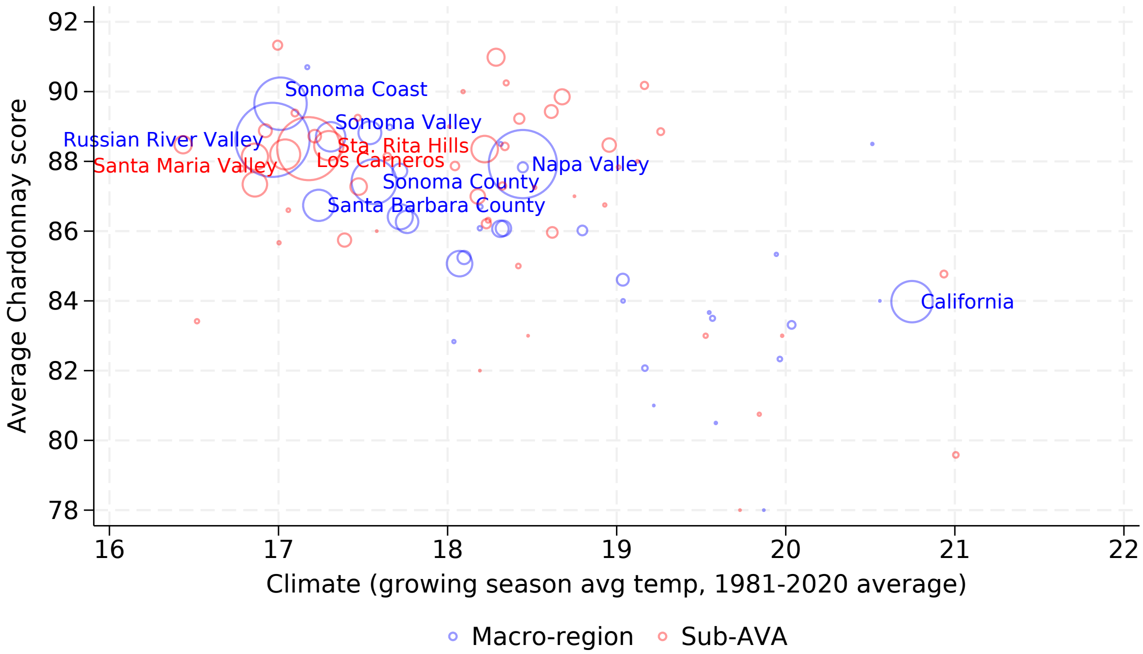

(a) Score

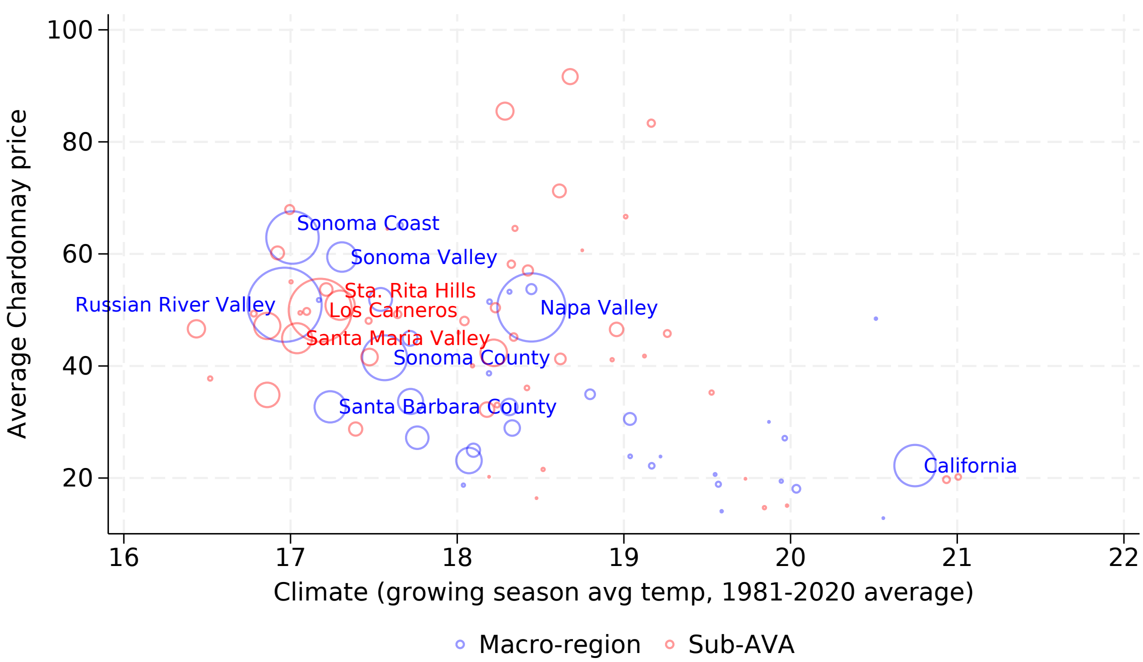

(b) Price (2022 dollars per bottle)

Figure 3.7: Average Chardonnay wine scores and prices by climate and region

Note: The size of the bubble corresponds to the number of observations. We label the 10 regions with the largest number of observations.

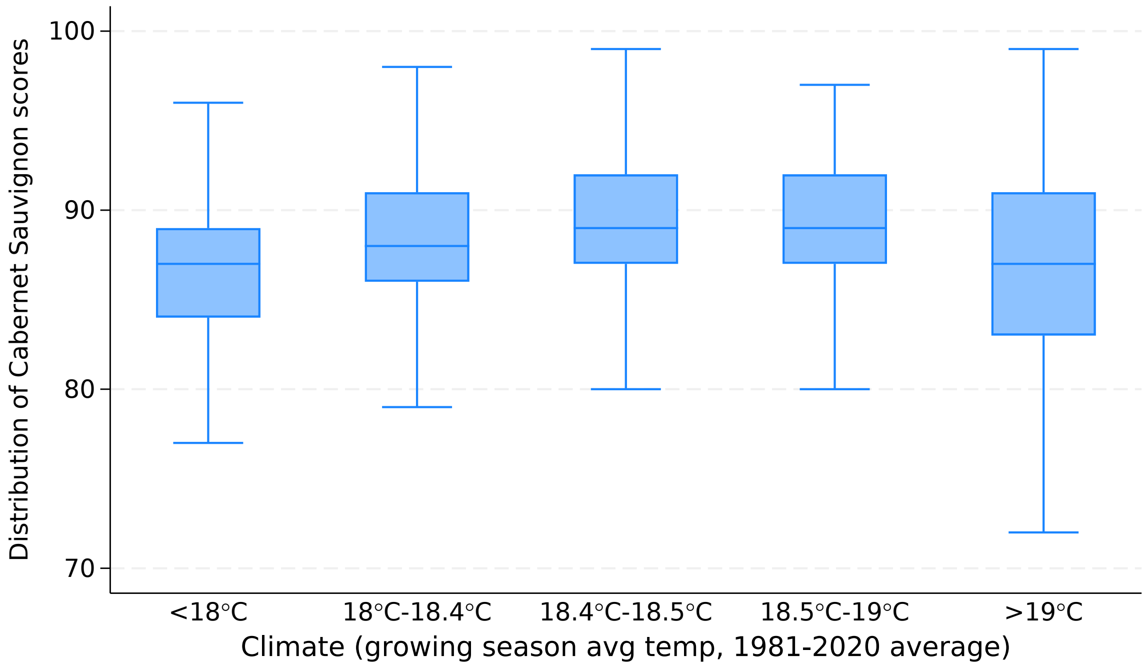

(a) Score

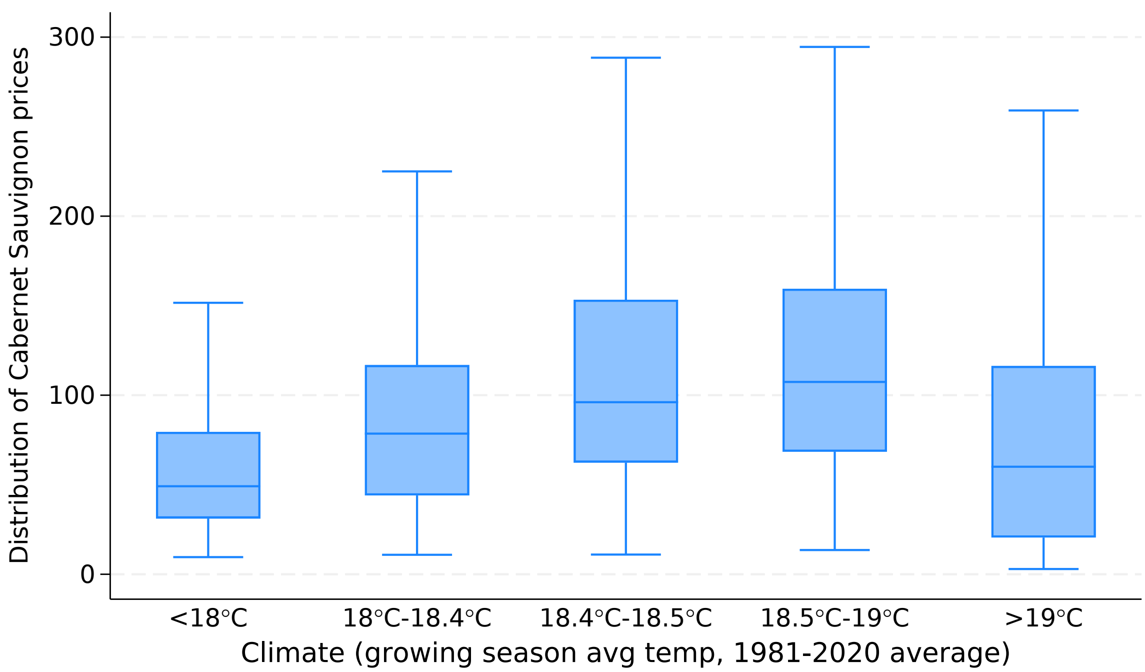

(b) Price (2022 dollars per bottle)

Figure 3.8: Distribution of Cabernet Sauvignon wine scores and prices by climate

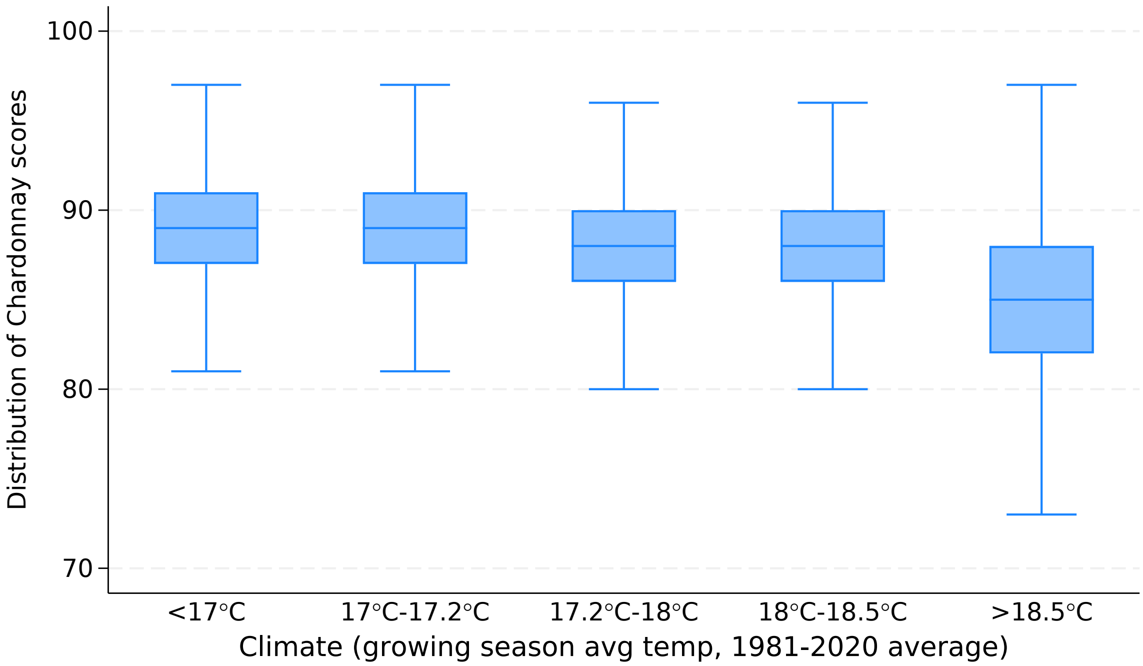

(a) Score

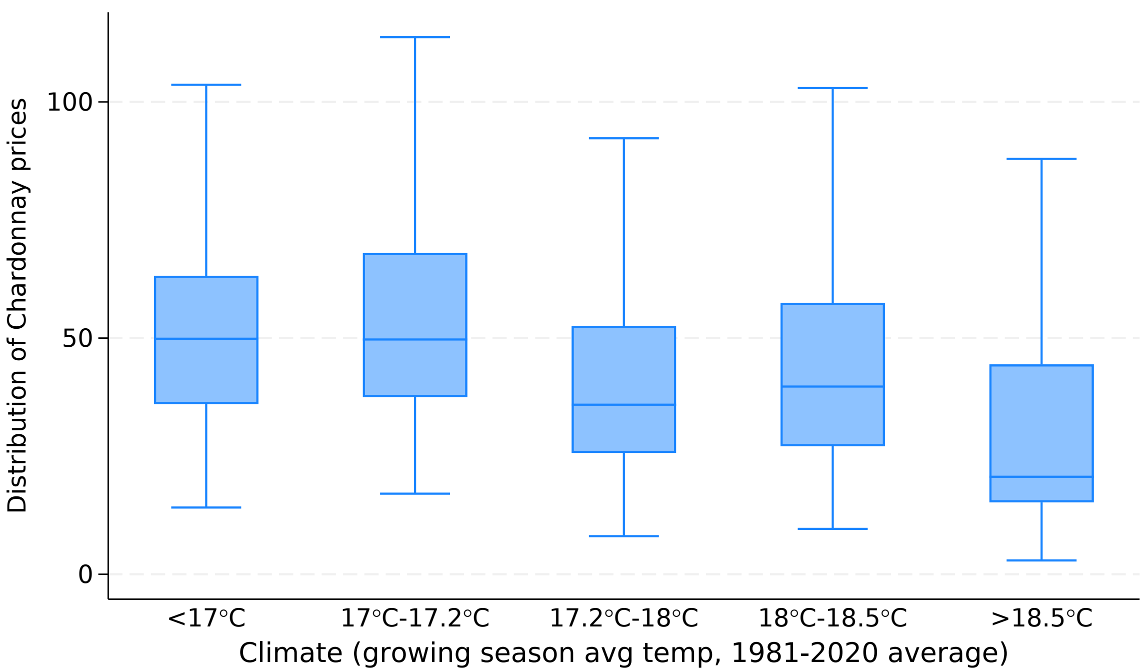

(b) Price (2022 dollars per bottle)

Figure 3.9: Distribution of Chardonnay wine scores and prices by climate

3.5 Statistical analysis of vintage weather effects

Here, we establish statistical models of varietal wine prices and scores as a function of measures of weather for each of the main varieties. The goal of this model is to estimate vintage effects arising from temperature variation around the regional norm. Since wine grape varieties (Cabernet Sauvignon, Chardonnay, Merlot, Pinot Noir, and Zinfandel) have distinct optimal climates, we estimated the model for each variety separately. We estimated:

$$y_{it} = \alpha_{r(i)} + f\left(x_{r(i)t}, \mu_{r(i)}; \beta\right) + \rho_1 ppt_{r(i)t} + \rho_2 ppt^2_{r(i)t} + \psi\left(t\right) + \epsilon_{it} \tag{3.1}$$

where y_{it} is the outcome variable (price or score) for wine observation i using grapes grown in region $r\left(i\right)$ in vintage year t. We specified the weather variable $x_{r(i)t}$ as either (a) the regional average temperature, averaged from April to October in vintage year t, or (b) the regional growing degree days (GDD) between 10°C and 35°C, summed from April to October in vintage year t. The climate variable $\mu_{r(i)}$ is the forty-year average of the weather variable (average temperature or GDD) from 1981 to 2020. Precipitation during the growing season, $ppt_{r(i)t}$ is summed from April to October. $\psi\left(t\right)$ is a quadratic time trend. We included region fixed effects $\alpha_{r(i)}$ to absorb time-invariant characteristics of the regions. The errors are heteroskedastic robust and clustered by region.

The function $f\left(x_{r(i)t}, \mu_{r(i)}; \beta\right)$ relates regional weather and climate variables to the outcome variables. We made three changes to frequently used quadratic-in-weather specification (e.g., as used by Jones et al. (2005) and Haeger & Storchmann (2006)). First, we focused on deviations in weather from local climate. Second, we fitted a flexible restricted cubic spline (also known as a natural cubic spline) that allows for asymmetry in the relationship between weather and wine outcomes. Third, we explored the role of AVAs in influencing the relationship between weather variation and wine.

We focused on deviations in weather from local climate, which we define as $d_{r(i)t} = x_{r(i)t} - \mu_{r(i)}$. The underpinning idea is that producers have adapted to their local climate, and any deviation from local climate may be sub-optimal for the quality and price of wine.

For the specification, first consider the quadratic specification $f(.) = \beta_1 d_{r(i)t} + \beta_2 d^2_{r(i)t}$. The quadratic specification assumes symmetry: the economic impacts of increases and decreases in temperature–relative to climate–are mirror images of each other.

We prefer a specification that does not impose symmetry and allows positive and negative deviations in weather from climate to have differential impacts on wine price and quality. If the relationship really is symmetric then this would be a special case of our more flexible model. However, fitting a more flexible model is not without its own potential pitfalls. For example, a high-order polynomial like a cubic can generate a curve that doesn't fit the data well at the tails i.e. when weather is especially different from climate.

We instead fitted a spline on deviations in weather from climate. We first splitted the predictor ($d_{r(i)t}$ in our case) into segments with "knots" defining the start and end of each segment. We then fitted a polynomial (such as a cubic) in each segment and force the segments to join smoothly at the knots. We opted to use a restricted cubic spline, which is a cubic spline with the added restriction that the first and last segments are linear, thus reducing the risk of overfitting the data-sparse tails. We chose four knots that split the distribution of weather deviations into quintiles.

A restricted cubic spline with four knots has a function $f(.) = \beta_1 d_{r(i)t} + \beta_2 V_{1,r(i)t} + \beta_3 V_{2,r(i)t}$. It introduces new variables $V_{1,r(i)t}$ and $V_{2,r(i)t}$ that are transformations of $d_{r(i)t}$ defined in Appendix 3.D. Together, this $f(.)$ function captures (a) the cubic polynomials between knots 1 and 2 and knots 3 and 4, (b) the restriction that individual curves join smoothly at the knots, and (c) the restriction that the segments before knot 1 and after knot 4 are linear. A downside of the restricted cubic spline is that the individual β coefficients are not intuitive. Restricted cubic spline results are best shown graphical which is the focus of the results.

Finally, we examined the role of AVAs in mediating the relationship between weather variation and price premia and quality. To examine heterogeneity in the response to weather, we interacted temperature variables with an indicator variable $I(r(i) \in \text{sub-AVA})$ where $I(.)$ takes on a 1 if the condition inside the parentheses is satisfied, and 0 otherwise. The inclusion of this indicator variable allowed us to compare the effects on wines associated with sub-AVAs and wines associated with macro-regions.

$$\begin{aligned} y_{it} =& \alpha_{r(i)} + f\left(x_{r(i)t}, \mu_{r(i)}; \beta\right) \ + \ f\left(x_{r(i)t}, \mu_{r(i)}; \beta\right)\left(I\left(r\left(i\right) \in \text{sub-AVA}\right)\right. \\ &+ \rho_1 ppt_{r(i)t} + \rho_2 ppt^2_{r(i)t} + \psi\left(t\right) + \epsilon_{it} \end{aligned} \tag{3.2}$$

Next, we show results from these statistical models that aim to uncover the effect of vintage weather on wine price and quality. Results for Cabernet Sauvignon and Chardonnay are shown in Figure 3.10 to Figure 3.13 while results for the other varieties (Merlot, Pinot Noir, and Zinfandel) are shown in Appendix 3.E. Here, we use average temperatures over the growing season (April to October) as the definition of weather. Alternative results using growing degree days are presented in Appendix 3.F. For each figure, the top panel shows the effect of a deviation in average temperature from regional climate on wine scores and prices. The x-axis is the average temperature minus regional climate: zero indicates that temperatures and climate were equal, one indicates that average temperatures were 1°C warmer than the regional climate, and negative one indicates that average temperatures were 1°C cooler than the regional climate. The bottom panel of each graph shows the distribution of average temperature minus regional climate associated with the observed wine prices and scores.

Cabernet Sauvignon wines produced with grapes grown in warmer temperatures than the regional climate have lower scores and prices on average. As shown in Figure 3.10, average temperatures 2°C warmer than the regional climate cause Cabernet Sauvignon prices to decline to an average of \$90 per bottle compared with \$106 per bottle when temperatures are at their regional norm. Across all the data, Cabernet Sauvignon scores average around 88.5 points, but this drops to 87 points for vintages that were warmer than usual by 1.5°C or more. Cabernet Sauvignon wine prices and scores benefit from cooler average temperatures.

Vintages exposed to temperatures 1.5°C cooler than the regional climate scored more than 89 points and were priced over $120 per bottle on average.

We compared the effects of deviations from climate on Cabernet Sauvignon wines from Napa Valley AVA (n=6,938) and sub-AVAs within Napa Valley (n= 3,996) including the Oakville, Rutherford, and Stags Leap District AVAs–premium producers of Cabernet Sauvignon wines. Temperatures warmer than the regional norm had significantly different effects on prices for and scores for wines from macro-regions compared with sub-AVAs as shown in Figure 3.11. We find that warmer temperatures are beneficial for the prices of Cabernet Sauvignon wine from sub-AVAs on average, although there is some imprecision in the results. By comparison, temperatures above the regional norm cause Napa Valley AVA wines to decline in price.

Chardonnay wines scores are negatively affected by both positive and negative deviations in weather from climate but because of imprecision in the estimates we can't rule out no effect (see Figure 3.12). Average temperatures above regional climate cause Chardonnay prices to decline: a 1.5°C warmer vintage has an average price of $40 per bottle compared to almost $50 per bottle when temperatures are at their regional norm.

To unpack heterogeneous effects on Chardonnay wines, we focused on Central Coast AVAs with cool coastal climates that are beneficial for Chardonnay wine grapes. We compare wines from Central Coast macro-regions (n=854) including Central Coast, Paso Robles, and Monterey, with wines from sub-AVAs within the Central Coast macro-regions (n=1,531) including Sta. Rita Hills, Santa Maria Valley, and Edna Valley AVAs. Like Cabernet Sauvignon wines, we find large negative effects of temperatures warmer than climate on the prices of Chardonnay from macro-regions shown in Figure 3.13. Wines from premier sub-AVAs are broadly unaffected by weather deviations from climate.

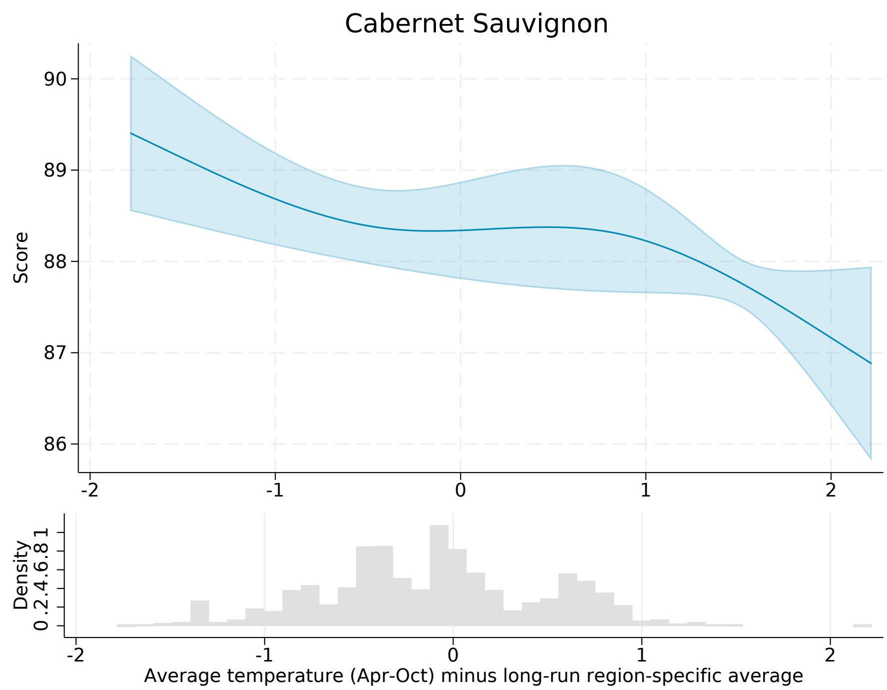

(a) Score

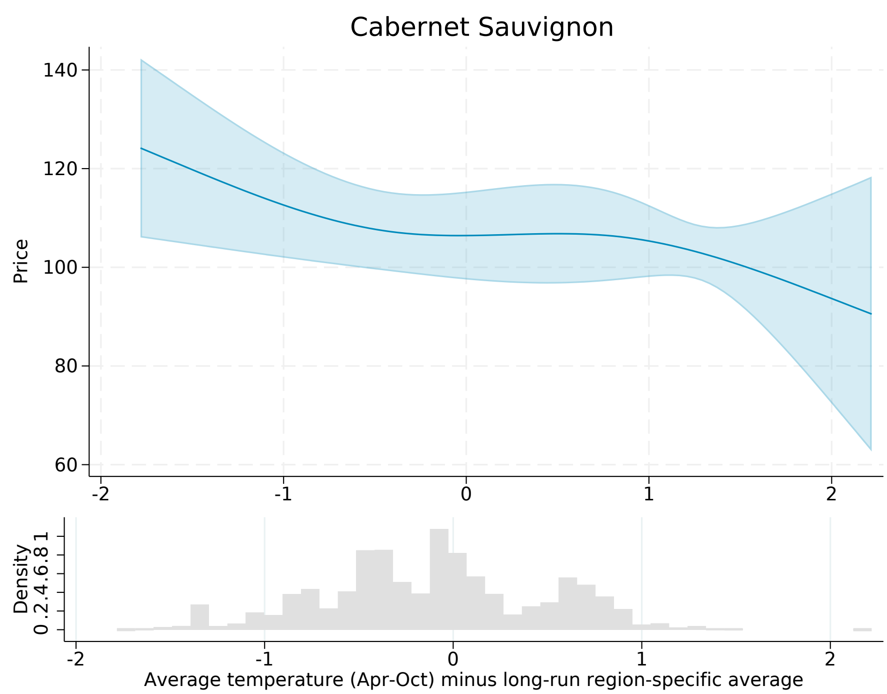

(b) Price (2022 dollars per bottle)

Figure 3.10: Effect of a deviation in temperature from regional climate on Cabernet Sauvignon wine scores and prices

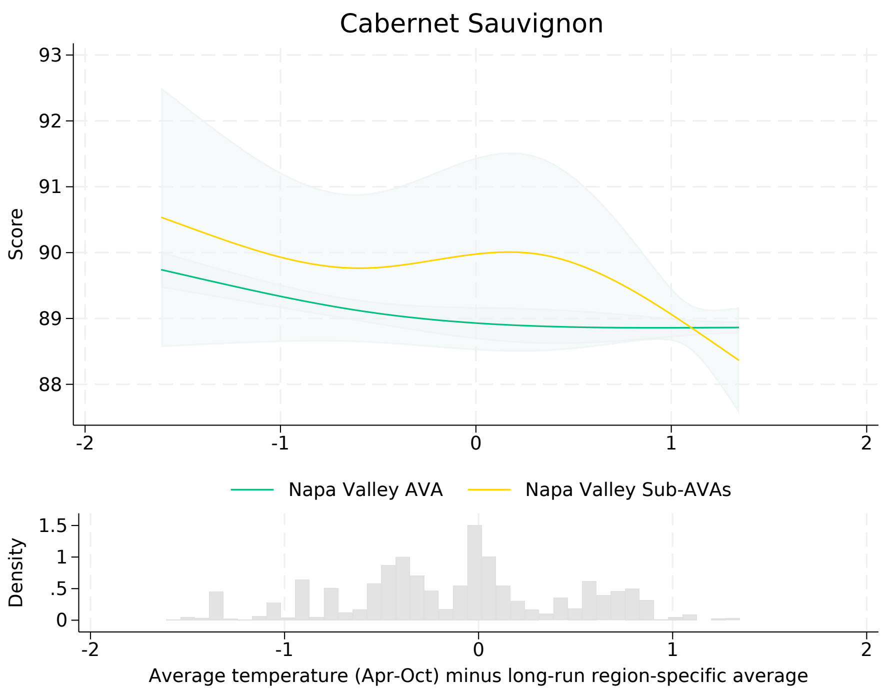

(a) Score

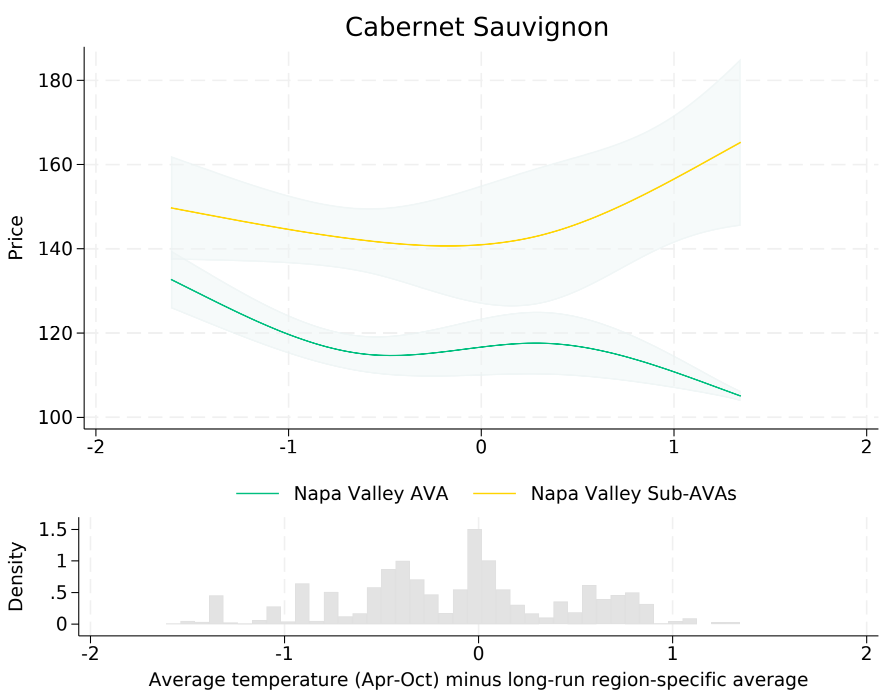

(b) Price (2022 dollars per bottle)

Figure 3.11: Effect of a deviation in temperature from regional climate on Cabernet Sauvignon wine scores and prices by Napa Valley AVA and Napa Valley Sub-AVAs

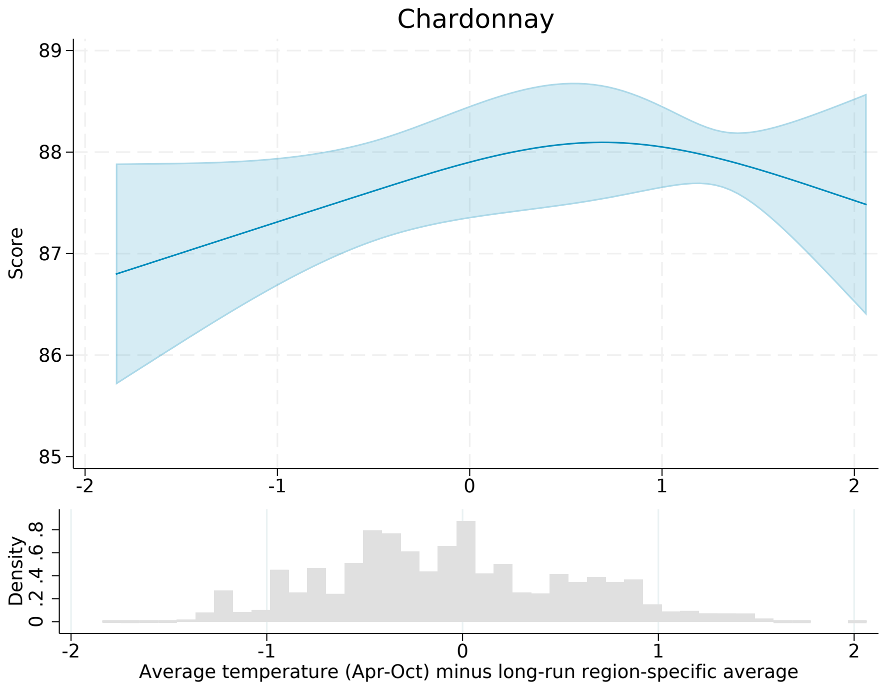

(a) Score

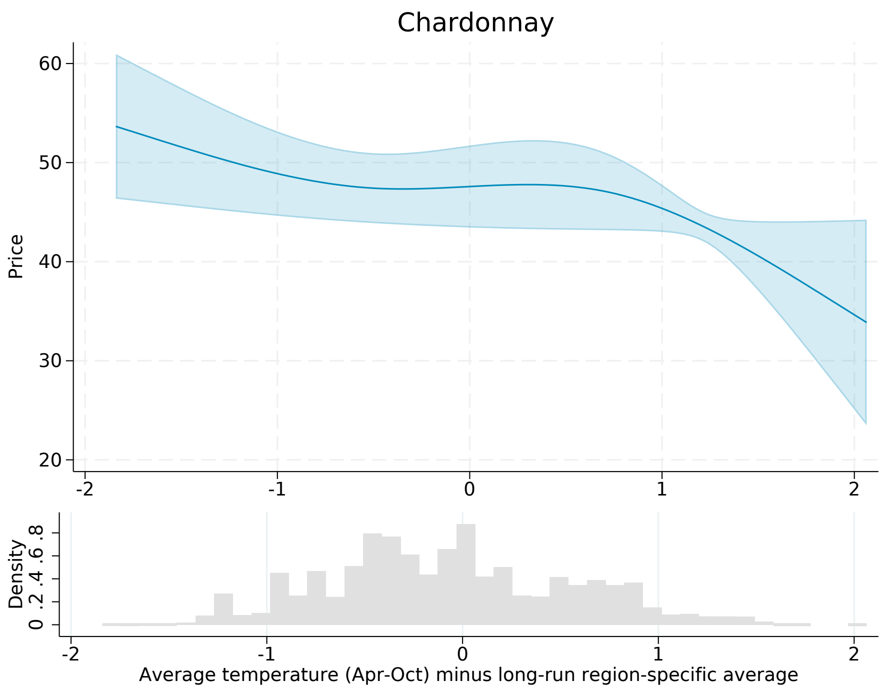

(b) Price (2022 dollars per bottle)

Figure 3.12: Effect of a deviation in temperature from regional climate on Chardonnay wine scores and prices

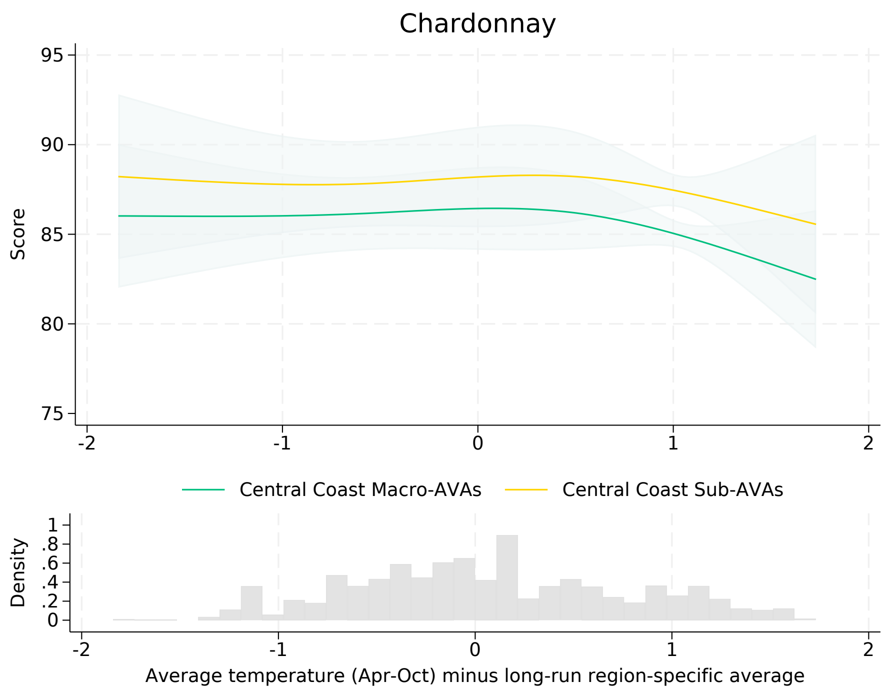

(a) Score

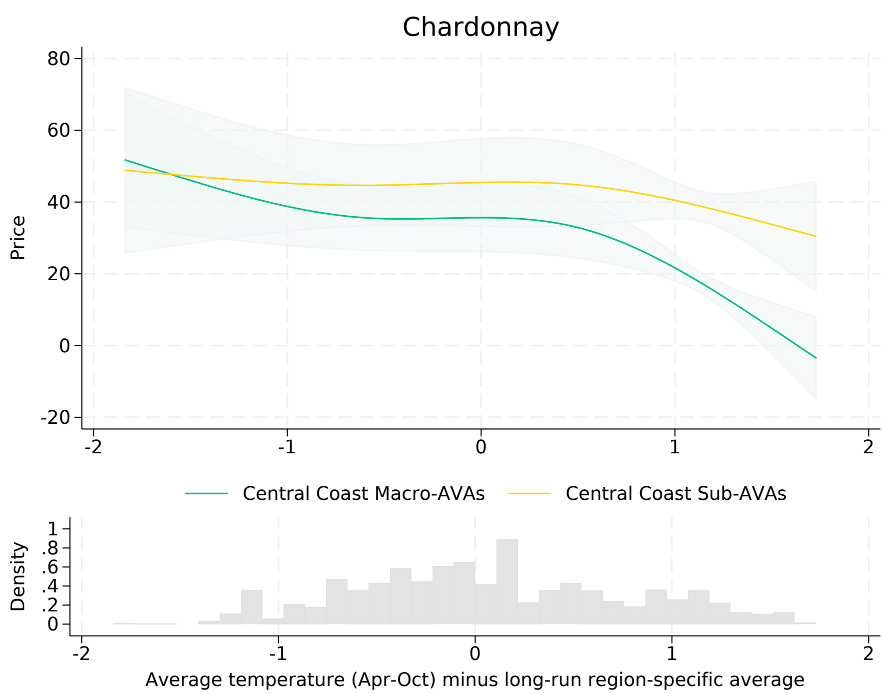

(b) Price (2022 dollars per bottle)

Figure 3.13: Effect of a deviation in temperature from regional climate on Chardonnay wine scores and prices by Central Coast Macro-Regions and Central Coast Sub-AVAs

In Appendix 3.F, we checked the robustness of our results to several modelling decisions. We show results using alternative definitions of the temperature variable (growing degree days) and climate variable (10-year moving average). We also show results for the log of the dependent variables (price and score) instead of natural units. Overall, the results in these robustness checks are broadly similar in magnitude and significance to the main results.

3.6 Discussion and Conclusion

How does weather variation around the regional norm affect wine quality and prices? Previous work finds mixed evidence of the effect of weather variation on California's wines. Haeger & Storchmann (2006) found significant evidence of a quadratic relationship between temperature and the prices of Pinot Noir wines from California and Oregon. Temperatures above the optimum caused a drop in Pinot Noir prices on average. Ramirez (2008) found that weather variability affects Napa wine quality and prices, but their results indicated that "warmer summers tend to be associated with lower, not higher, quality ratings, a result that does not coincide with expectations." Jones et al. (2005) estimated a quadratic-in-temperature specification on Sotheby's vintage ratings in wine regions across the world and found no effect of temperature on U.S. wine quality.

These three papers rely on coarse measurement of weather and climate variables. Haeger & Storchmann (2006) and Ramirez (2008) used observations from sparse weather stations to proxy for true weather (and climate) experienced at the vineyards often 20 to 30 miles away. Jones et al. (2005) used gridded weather data but took one representative grid for all coastal California regions. Our spatially detailed weather data from PRISM were interpolated to 800m grids and accounted for factors that affect the local microclimate such as elevation, coastal proximity, and aspect. We calculated a region's weather and climate variables by taking the weighted average of grids that fall within the region, where each grid's weight is equal to the share of the region's grape acreage within the grid. This allows us to represent the relevant concepts of weather and climate more accurately.

We also used a restricted cubic spline specification that does not impose symmetry for positive and negative shocks. Indeed, the effects are asymmetric: a vintage year warmer than the regional norm causes wine to be lower in quality and price on average, whereas a vintage year cooler than the regional norm causes null or positive effects on wine quality and prices depending on the grape variety.

How does being associated with an AVA influence the relationship between weather, climate, and wine? Wine prices are a function of reputation which itself is a function of historical quality performance. Costanigro et al. (2010) find that inconsistent quality, measured through an increase in the standard deviation of wine scores, causes a reduction in reputation premia for both individual wineries and AVAs. According to these authors, for cheaper wines from macro-regions, any reputation premia is more likely a consequence of collective reputation of the region. Individual investment in consistent quality is shared among all the wines that collectively associate with the region. By comparison, winemakers in premier regions benefit from investing in consistent quality through two channels: (a) the collective reputation premia, and (b) the winery's individual reputation premia.

We find that the prices of wine from macro-regions are lower if they use grapes grown during vintages that are warmer than the regional norm. By comparison, prices and scores for wines from premier sub-AVAs are comparatively resilient to deviations in weather from their regional norm, suggesting that growers and/or winemakers in sub-AVAs invest in weather-mitigating practices. This finding is consistent with the greater incentive for those making premium wines to maintain consistent quality and the preserve the reputation premia it confers on their wines.

Climate change is predicted to cause wine-growing regions in California to become warmer by the end of the century. As shown in Figure 3.14, the average temperature between April and October in the Napa Valley AVA is predicted to average almost 22°C in the last ten years of the century, more than 3°C degrees higher than the thirty-year average to 2020. Decisions by vineyard managers and winemakers will influence the extent to which climate change will affect wine quality and price.

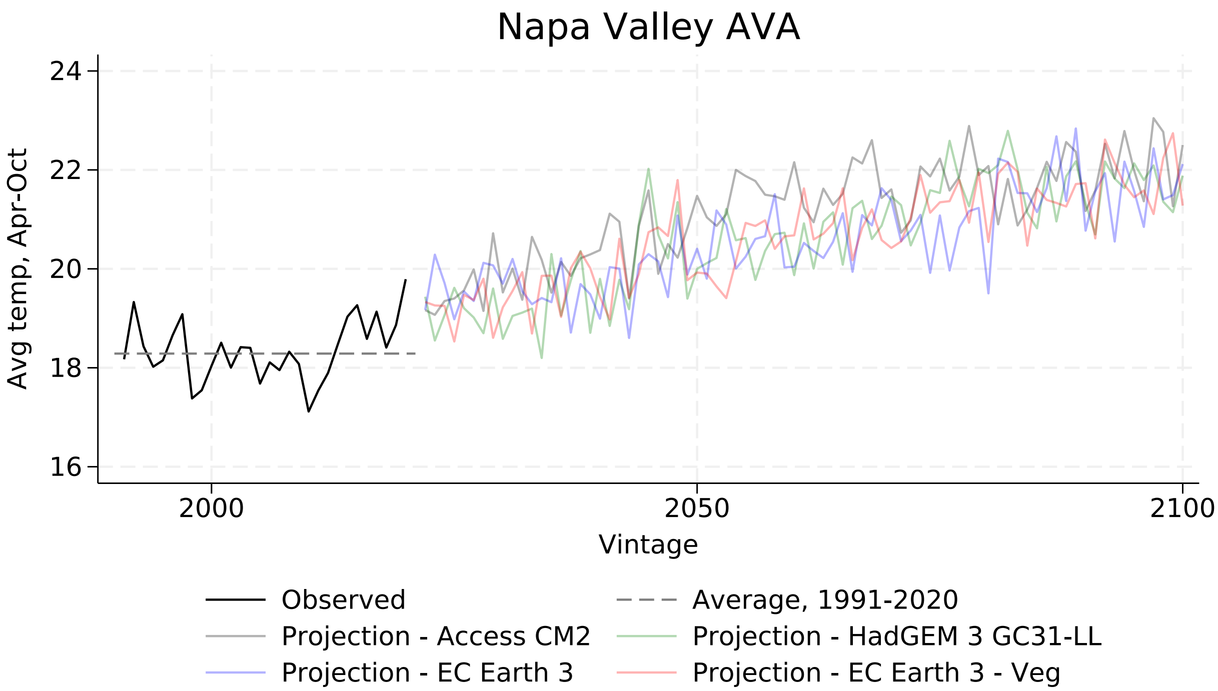

Figure 3.14: Observed and projected average temperature (April to October) in Napa Valley AVA

Note: Projected temperatures are for a middle-of-the-road global emissions scenario (SSP2-4.5), statistically downscaled to a 3km resolution in California by Pierce et al. (2023) and available from Cal-Adapt (2023).

Over the course of many decades and even centuries of wine grape growing, producers have adapted their practices to better suit their local climate. Studies on the potential to adapt to climate change predominately focus on varietal choice (e.g. Wolkovich et al. (2018)) and changing production locations (e.g. Moriondo et al. (2013)). Some work highlights potential adaptation within the production process like harvesting date, canopy management, and irrigation (e.g. Webb et al. (2021); Van Leeuwen & Darriet (2016); Santos et al. (2020)). Our results suggest that some growers and winemakers are mitigating the effects of vintage weather shocks on wine quality and prices. Further work should be done to understand whether these practices could reduce the negative consequences of climate change.

References

Alston, J. M., Anderson, K., & Sambucci, O. (2015). Drifting towards Bordeaux? The evolving varietal emphasis of US wine regions. *Journal of wine economics*, 10(3), 349–378.

Alston, J. M., Fuller, K. B., Lapsley, J. T., & Soleas, G. (2011). Too much of a good thing? Causes and consequences of increases in sugar content of California wine grapes. *Journal of Wine Economics*, 6(2), 135–159.

Alston, J. M. & Gaeta, D. (2021). Reflections on the political economy of european wine appellations. *Italian Economic Journal*, 7(2), 219–258.

Alston, J. M. & Sambucci, O. (2019). Grapes in the world economy. *The grape genome*, (pp. 1–24).

American Viticultural Areas Digitizing Project Team (2021). Current ava boundaries. Available at `https://github.com/UCDavisLibrary/ava`.

Anderson, K., Meloni, G., & Swinnen, J. (2018). Global alcohol markets: Evolving consumption patterns, regulations, and industrial organizations. *Annual Review of Resource Economics*, 10, 105–132.

Ashenfelter, O. (2010). Predicting the quality and prices of bordeaux wine. *Journal of Wine Economics*, 5(1), 40–52.

Ashenfelter, O., Ashmore, D., & Lalonde, R. (1995). Bordeaux wine vintage quality and the weather. *Chance*, 8(4), 7–14.

Ashenfelter, O. & Storchmann, K. (2016). Climate change and wine: A review of the economic implications. *Journal of Wine Economics*, 11(1), 105–138.

Bombrun, H. & Sumner, D. A. (2003). What determines the price of wine. *AIC Issues Brief*, 18, 1–6.

Byron, R. P. & Ashenfelter, O. (1995). Predicting the quality of an unborn grange. *Economic Record*, 71(1), 40–53.

Cahill, K. N., Lobell, D. B., Field, C., Bonfils, C., & Hayhoe, K. (2007). Modeling climate change impacts on wine grape yields and quality in California. *Seminare: Réchauffement climatique, quels impacts probables sur les vignobles*, (pp. 28–30).

Cal-Adapt (2023). Data derived from LOCA2 Downscaled CMIP6 Climate Projections. Cal-Adapt website developed by University of California at Berkeley's Geospatial Innovation Facility under contract with the California Energy Commission. Available at `https://cal-adapt.org/`.

California Department of Food and Agriculture (2022). 2021 grape crush report. Available at `https://www.nass.usda.gov/Statistics_by_State/California/Publications/Specialty_and_Other_Releases/Grapes/`.

Costanigro, M., McCluskey, J. J., & Goemans, C. (2010). The economics of nested names: name specificity, reputations, and price premia. *American Journal of Agricultural Economics*, 92(5), 1339–1350.

Davis, R. E., Dimon, R. A., Jones, G. V., & Bois, B. (2019). The effect of climate on burgundy vintage quality rankings. *OENO One*, 53(1).

Franken, J. R. (2014). Coordination of the California winegrape supply chain. *Journal of Wine Economics*, 9(2), 183–201.

Gambetta, G. A. & Kurtural, S. K. (2021). Global warming and wine quality: are we close to the tipping point? *Oeno One*, 55(3), 353–361.

Goodhue, R. E., Heien, D. M., Lee, H., & Sumner, D. A. (2003). Contracts and quality in the California winegrape industry. *Review of Industrial Organization*, 23, 267–282.

Haeger, J. W. & Storchmann, K. (2006). Prices of american pinot noir wines: climate, craftsmanship, critics. *Agricultural economics*, 35(1), 67–78.

Hannah, L., Roehrdanz, P. R., Ikegami, M., Shepard, A. V., Shaw, M. R., Tabor, G., Zhi, L., Marquet, P. A., & Hijmans, R. J. (2013). Climate change, wine, and conservation. *Proceedings of the National Academy of Sciences*, 110(17), 6907–6912.

Jones, G. V. (2004). Climate change in the western United States grape growing regions. In *VII International Symposium on Grapevine Physiology and Biotechnology 689* (pp. 41–60).

Jones, G. V. (2006). Climate and terroir: impacts of climate variability and change on wine. *Geoscience Canada Reprint Series*, 9, 203–217.

Jones, G. V., Edwards, E. J., Bonada, M., Sadras, V. O., Krstic, M. P., & Herderich, M. J. (2022). Climate change and its consequences for viticulture. In *Managing wine quality* (pp. 727–778). Elsevier.

Jones, G. V. & Goodrich, G. B. (2008). Influence of climate variability on wine regions in the western USA and on wine quality in the Napa Valley. *Climate Research*, 35(3), 241–254.

Jones, G. V., Reid, R., & Vilks, A. (2012). Climate, grapes, and wine: structure and suitability in a variable and changing climate. *The geography of wine: Regions, terroir and techniques*, (pp. 109–133).

Jones, G. V., White, M. A., Cooper, O. R., & Storchmann, K. (2005). Climate change and global wine quality. *Climatic change*, 73(3), 319–343.

Kwon, O. S., Lee, H., & Sumner, D. A. (2008). Appellation, variety, and the price of California wines. *ARE Update*, 11(4), 15–19. University of California Giannini Foundation of Agricultural Economics.

Lapsley, J. T., Alston, J. M., & Sambucci, O. (2019). *The U.S. Wine Industry*, chapter 5. Palgrave Macmillan.

Livat, F., Alston, J. M., & Cardebat, J.-M. (2019). Do denominations of origin provide useful quality signals? the case of bordeaux wines. *Economic Modelling*, 81, 518–532.

Lobell, D. B., Field, C. B., Cahill, K. N., & Bonfils, C. (2006). Impacts of future climate change on California perennial crop yields: Model projections with climate and crop uncertainties. *Agricultural and Forest Meteorology*, 141(2-4), 208–218.

Mérel, P., Ortiz-Bobea, A., & Paroissien, E. (2021). How big is the "lemons" problem? historical evidence from french wines. *European Economic Review*, 138, 103824.

Moriondo, M., Jones, G., Bois, B., Dibari, C., Ferrise, R., Trombi, G., & Bindi, M. (2013). Projected shifts of wine regions in response to climate change. *Climatic change*, 119, 825–839.

Nemani, R. R., White, M. A., Cayan, D. R., Jones, G. V., Running, S. W., Coughlan, J. C., & Peterson, D. L. (2001). Asymmetric warming over coastal California and its impact on the premium wine industry. *Climate research*, 19(1), 25–34.

Ortiz-Bobea, A. (2021). The empirical analysis of climate change impacts and adaptation in agriculture. In *Handbook of agricultural economics*, volume 5 (pp. 3981–4073). Elsevier.

Pierce, D., F., K. J., & Cayan, D. (2023). Localized Construction Analogs. Scripps Institution of Oceanography.

PRISM Climate Group, Oregon State University (2020). PRISM Climate Data. Available at `http://prism.oregonstate.edu`.

Puga, G., Anderson, K., Jones, G., Smart, R., et al. (2022). *Climate Change and the Evolving Mix of Grape Varieties in Australia's Wine Regions: Are They Related?* Technical report, University of Adelaide, Wine Economics Research Centre.

Ramirez, C. D. (2008). Wine quality, wine prices, and the weather: Is napa "different"? *Journal of Wine Economics*, 3(2), 114–131.

Sambucci, O. & Alston, J. M. (2017). Estimating the value of California wine grapes. *Journal of Wine Economics*, 12(2), 149–160.

Santos, J. A., Fraga, H., Malheiro, A. C., Moutinho-Pereira, J., Dinis, L.-T., Correia, C., Moriondo, M., Leolini, L., Dibari, C., Costafreda-Aumedes, S., et al. (2020). A review of the potential climate change impacts and adaptation options for european viticulture. *Applied Sciences*, 10(9), 3092.

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

Snyder, R. L. (1985). Hand calculating degree days. *Agricultural and forest meteorology*, 35(1-4), 353–358.

Stone, C. J. & Koo, C.-Y. (1985). Additive splines in statistics. In *Proceedings of the Statistical Computing Section ASA. Washington, DC* (pp. 45–48).

U.S. Bureau of Labor Statistics (2023). Consumer Price Index. Available at `https://www.bls.gov/cpi/`.

U.S. Census Bureau (2026). CA Geographic Boundaries. Available at `https://data.ca.gov/dataset/ca-geographic-boundaries`.

U.S. Treasury/TTB (2020). Wine Labeling: Appellation of Origin. Available at `https://www.ttb.gov/labeling-wine/wine-labeling-appellation-of-origin`.

U.S. Treasury/TTB (2021). Established American Viticultural Areas. Available at `https://www.ttb.gov/wine/established-avas`.

U.S. Treasury/TTB (2022). American Viticultural Area (AVA). Available at `https://www.ttb.gov/wine/ava.shtml`.

USDA NASS (2021). California County Agricultural Commissioners' Reports Crop Year 2020-2021. Available at https://www.nass.usda.gov/Statistics_by_State/California/Publications/AgComm/2021/CAC_2021_errata.pdf.

USDA National Agricultural Statistics Service (2022). Cropland Data Layer. Available at https://nassgeodata.gmu.edu/CropScape/.

Van Leeuwen, C. & Darriet, P. (2016). The impact of climate change on viticulture and wine quality. *Journal of Wine Economics*, 11(1), 150–167.

Webb, L., Whetton, P., & Barlow, E. (2007). Modelled impact of future climate change on the phenology of winegrapes in australia. *Australian Journal of Grape and Wine Research*, 13(3), 165–175.

Webb, L. B., Watt, A., Hill, T. C. J., Whiting, J., Wigg, F., Dunn, G., Needs, S., & Barlow, S. (2021). Extreme Heat: Managing Grapevine Response. Documenting Regional and Inter-regional Variation of Viticultural Impact and Management Input Relating to the 2009 Heatwave in South-Eastern Australia.

White, M. A., Diffenbaugh, N., Jones, G. V., Pal, J., & Giorgi, F. (2006). Extreme heat reduces and shifts United States premium wine production in the 21st century. *Proceedings of the National Academy of Sciences*, 103(30), 11217–11222.

Wine Institute (2023). California & US Wine Production. Available at https://wineinstitute.org/our-industry/statistics/california-us-wine-production.

Winfree, J. A. & McCluskey, J. J. (2005). Collective reputation and quality. *american Journal of agricultural Economics*, 87(1), 206–213.

Winkler, A. J. (1962). *General Viticulture, By AJ Winkler*. University of California Press.

Wolkovich, E., García de Cortázar-Atauri, I., Morales-Castilla, I., Nicholas, K., & Lacombe,

T. (2018). From pinot to xinomavro in the world's future wine-growing regions. *Nature Climate Change*, 8(1), 29–37.

3.A Scores and prices by vintage and grape variety

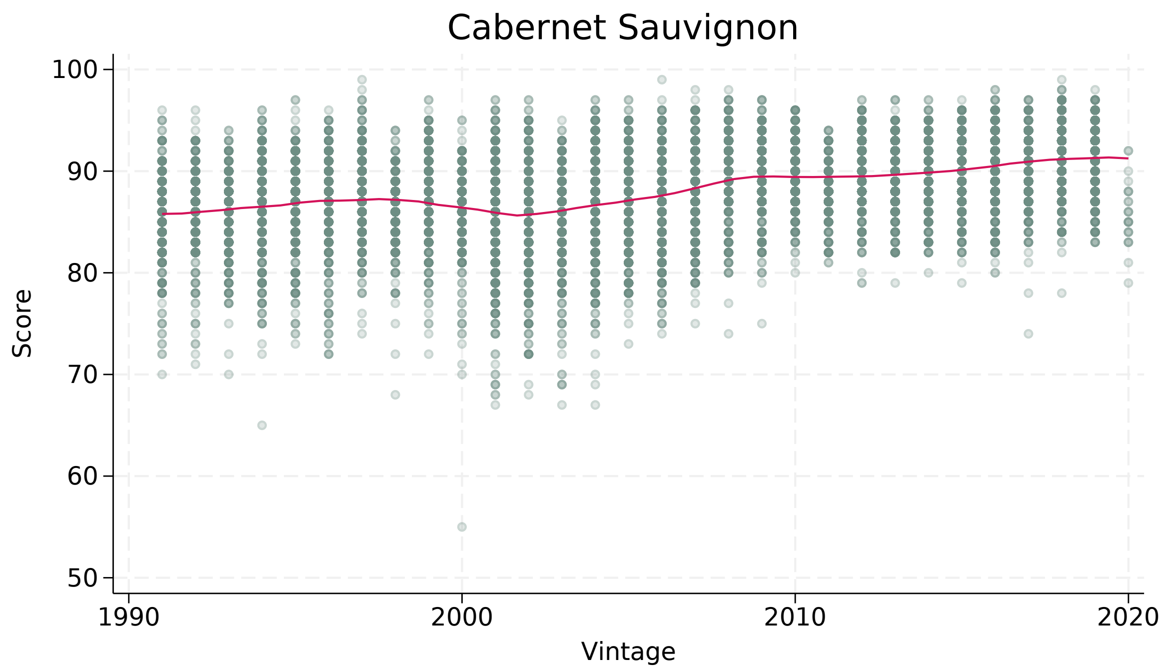

(a) Score

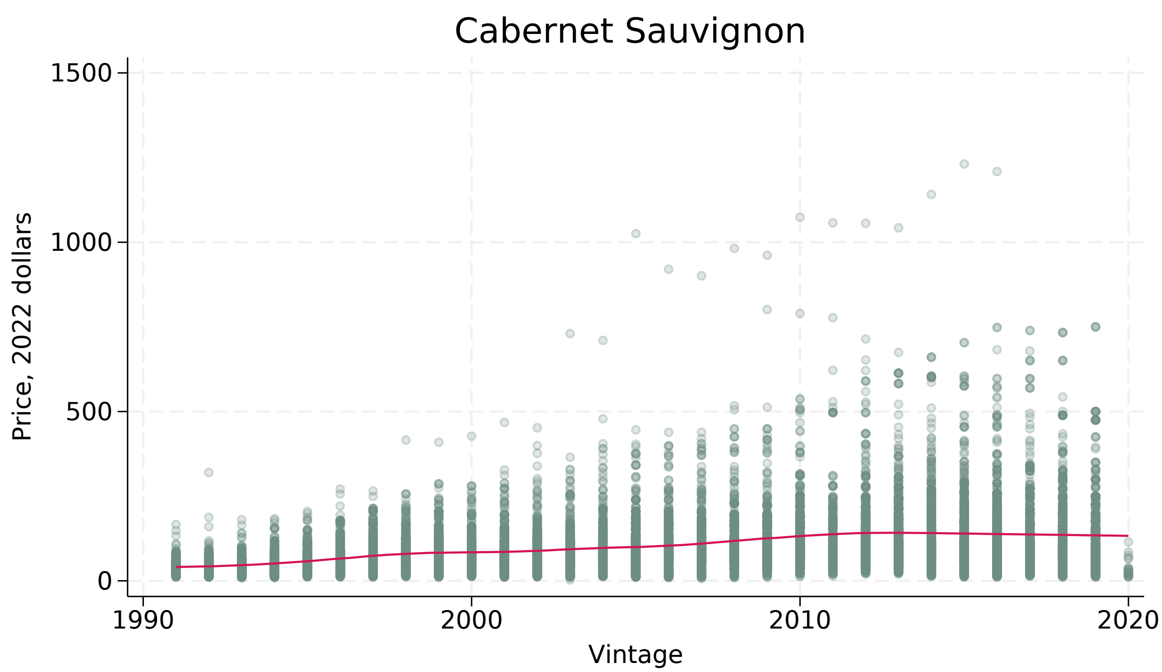

(b) Price (2022 dollars per bottle)

Figure 3.A.1: Wine scores and prices by vintage, Cabernet Sauvignon

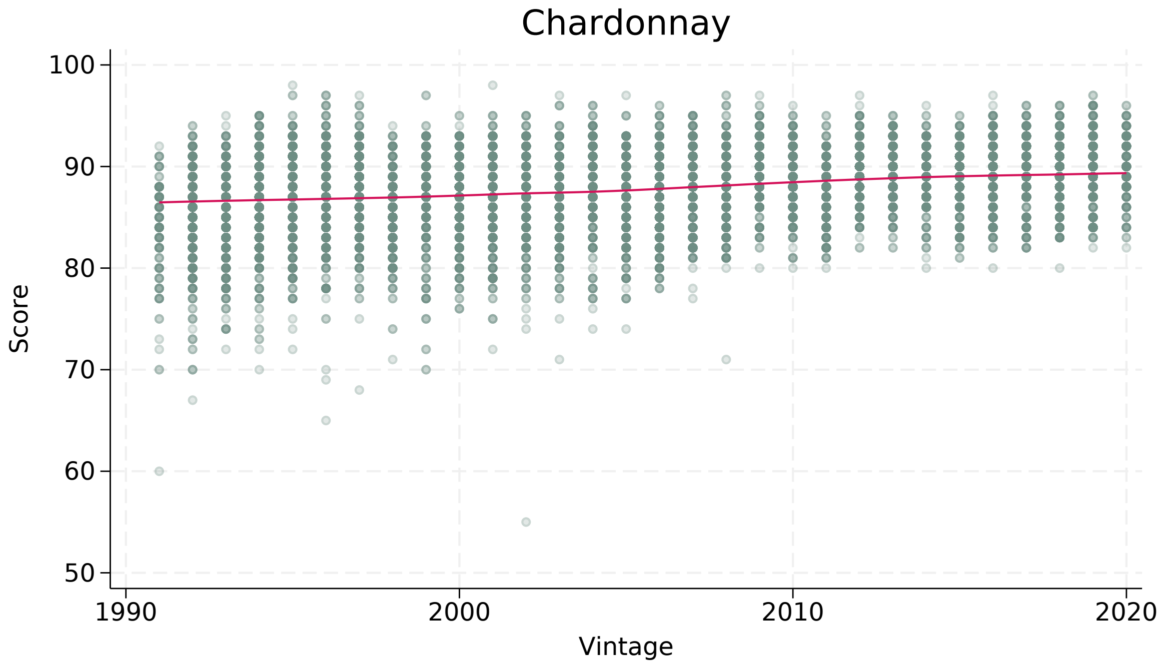

(a) Score

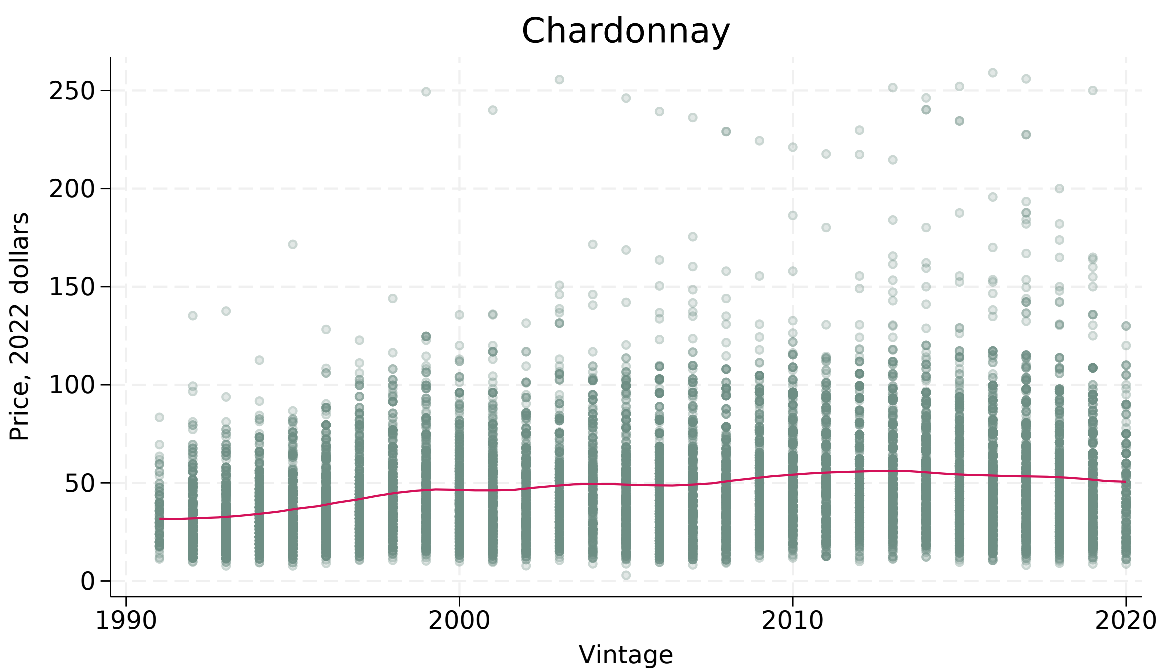

(b) Price (2022 dollars per bottle)

Figure 3.A.2: Wine scores and prices by vintage, Chardonnay

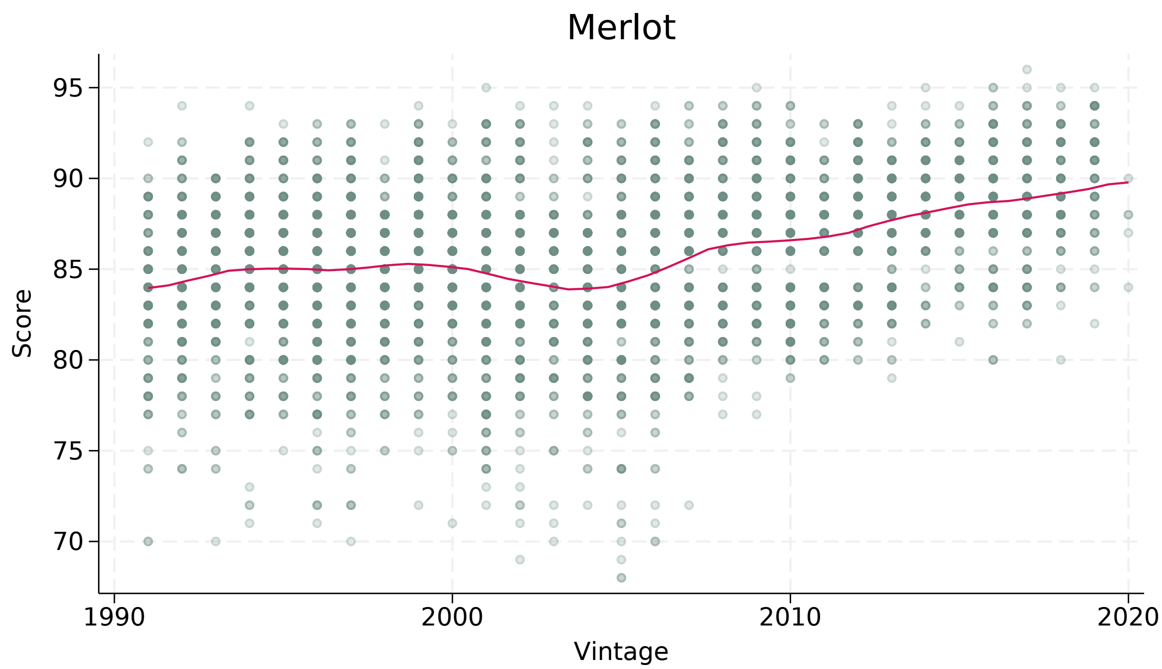

(a) Score

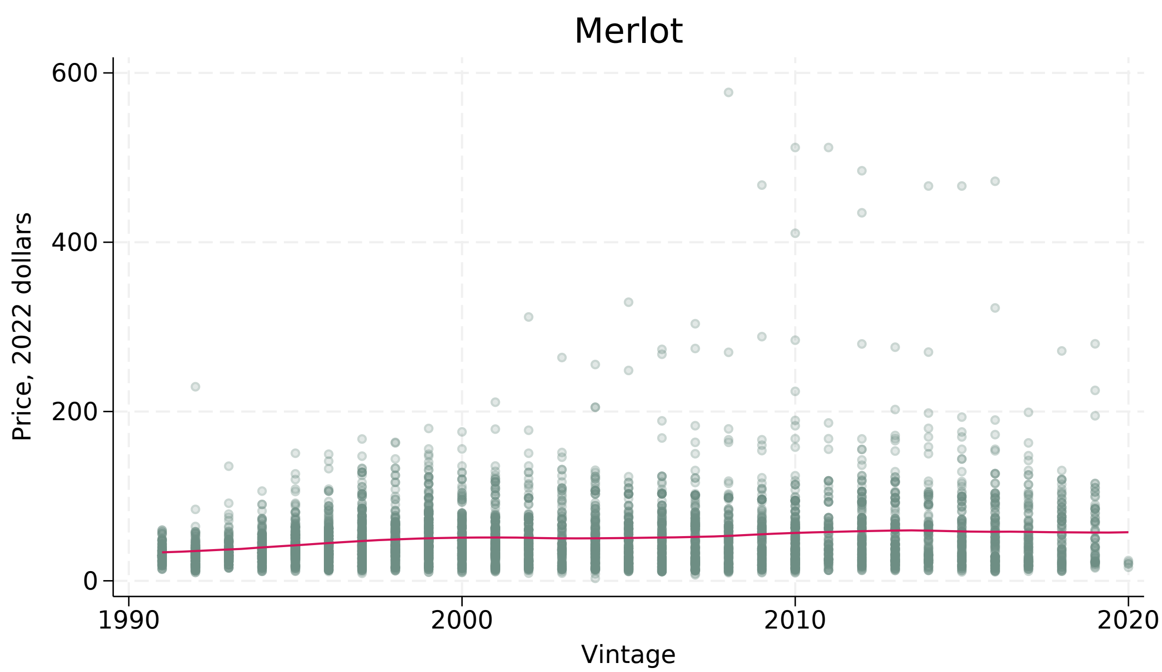

(b) Price (2022 dollars per bottle)

Figure 3.A.3: Wine scores and prices by vintage, Merlot

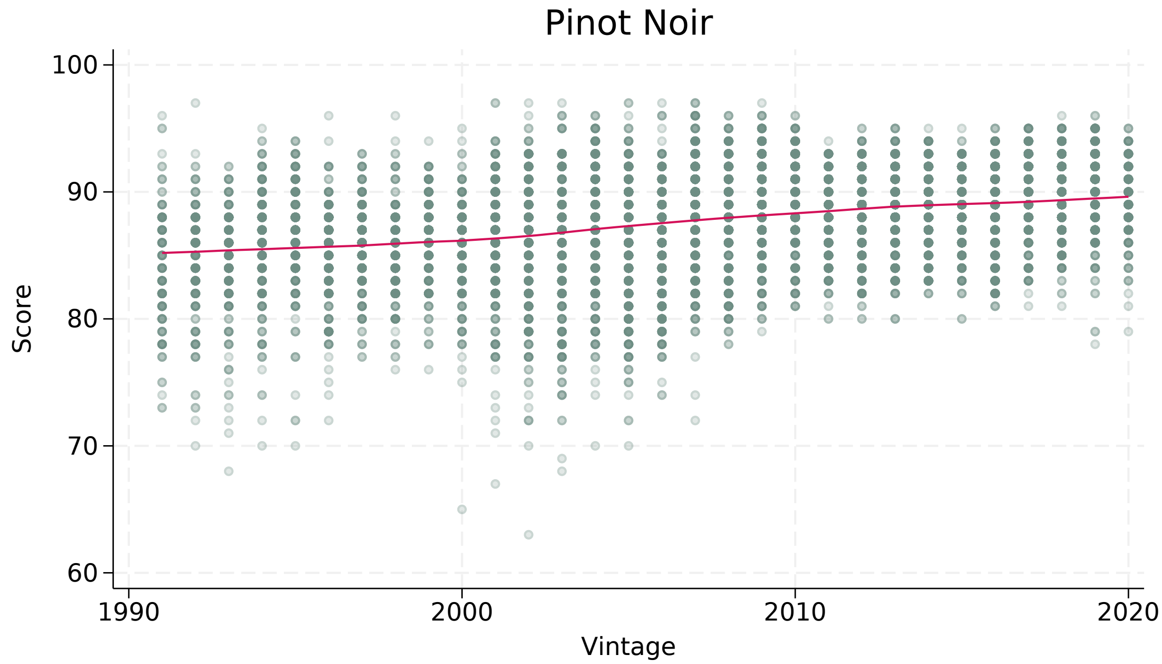

(a) Score

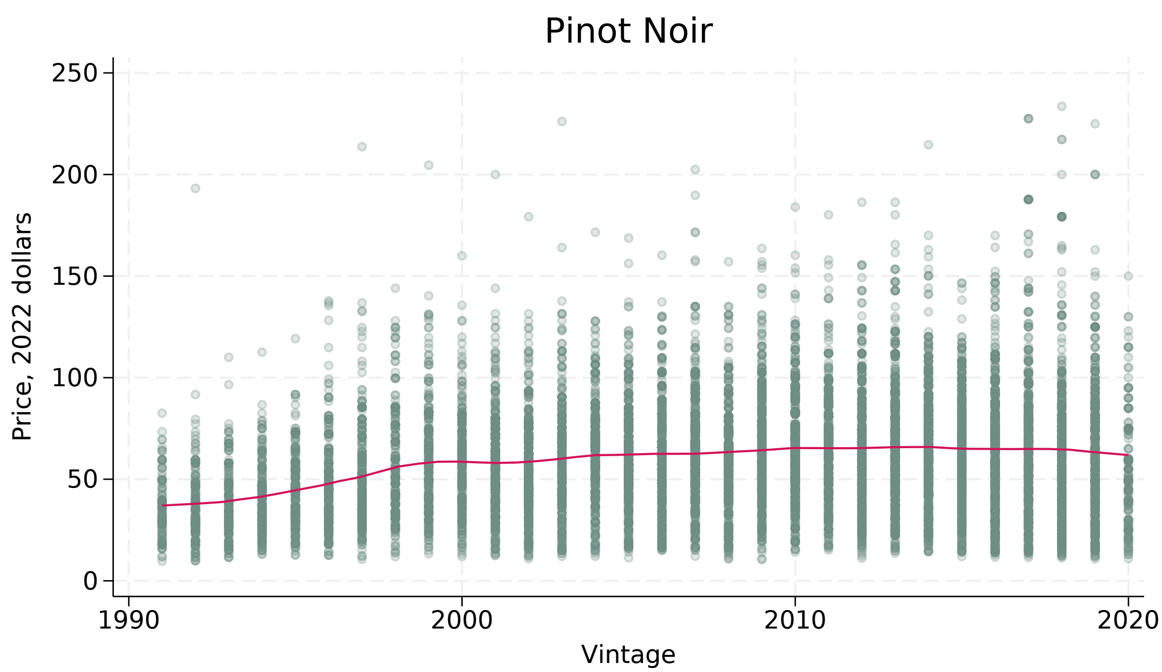

(b) Price (2022 dollars per bottle)

Figure 3.A.4: Wine scores and prices by vintage, Pinot Noir

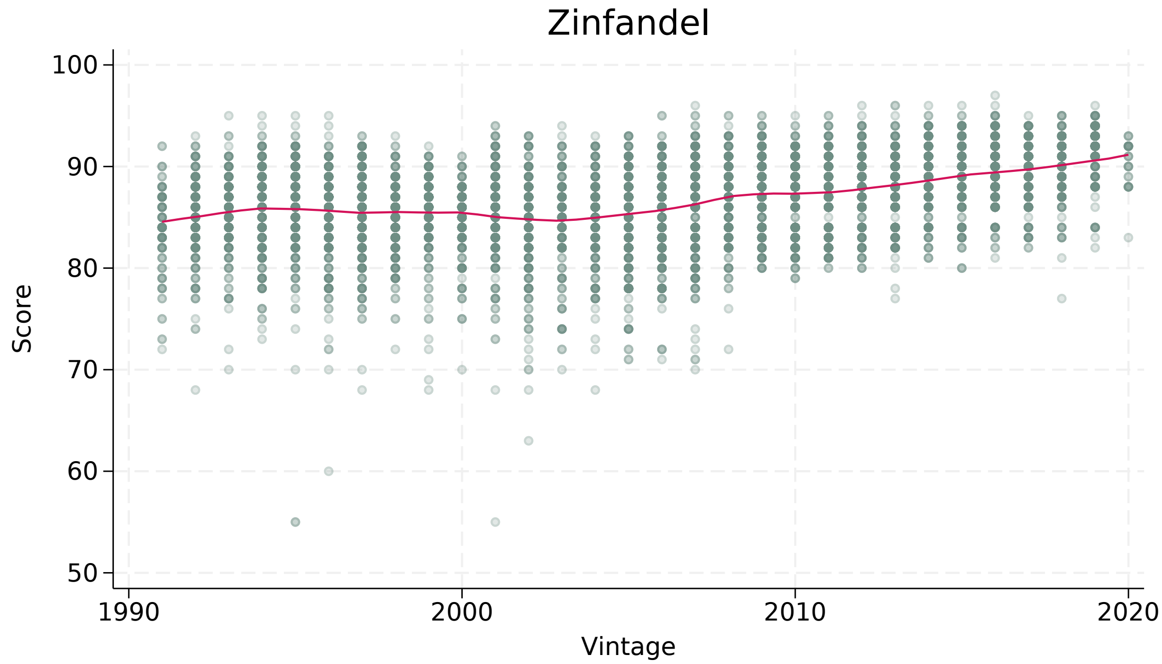

(a) Score

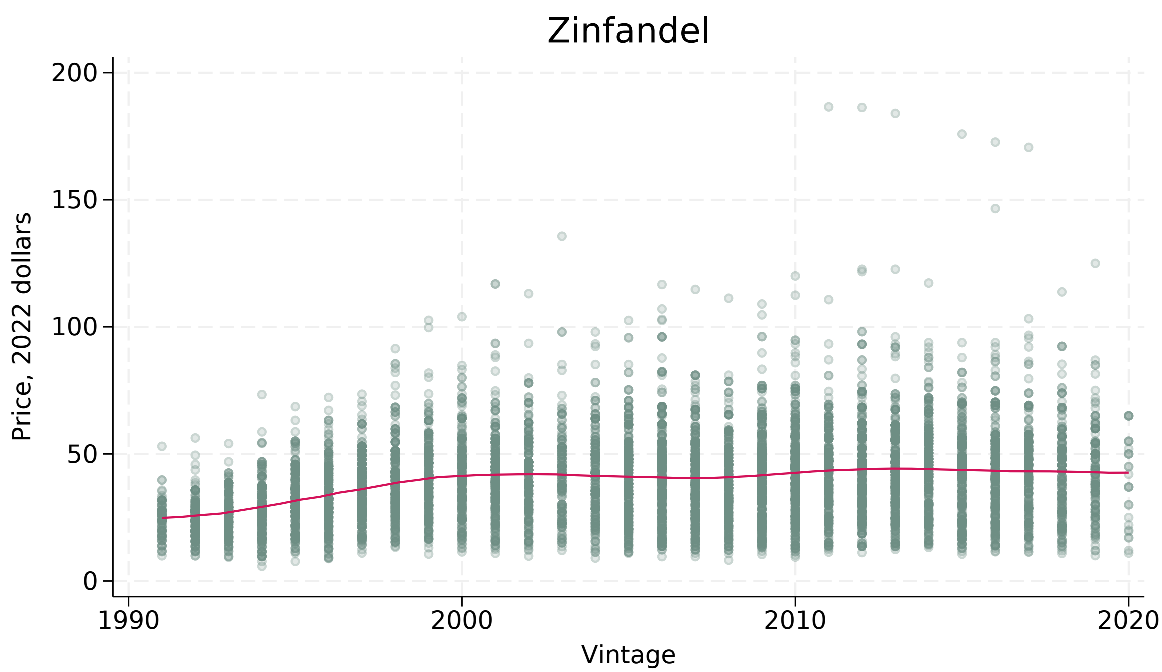

(b) Price (2022 dollars per bottle)

Figure 3.A.5: Wine scores and prices by vintage, Zinfandel

3.B Categorization of regions into macro-regions and sub-AVAs

Table 3.B.1: Categorization of regions into Macro-Region or Sub-AVA

Macro-Region?	Region	count
North Coast AVAs		
✓	Napa Valley	11,814
✓	Russian River Valley	6,420
✓	Sonoma Coast	3,501
	Los Carneros	3,127
	Dry Creek Valley	1,771
	Alexander Valley	1,625
✓	Sonoma Valley	1,528
	Anderson Valley	1,070
	Oakville	973
	Rutherford	857
	Howell Mountain	721
	Stags Leap District	498
	St. Helena	404
	Mt. Veeder	351
✓	Mendocino	345
	Spring Mountain District	329
✓	North Coast	316
	Sonoma Mountain	278
	Diamond Mountain District	227
	Knights Valley	216
	Yountville	194
	Green Valley of Russian River Valley	192
	Calistoga	132
	Fort Ross-Seaview	132
	Coombsville	126
	Oak Knoll District of Napa Valley	123
	Redwood Valley	117
	Chalk Hill	116
	Mendocino Ridge	114
	Rockpile	107
	Atlas Peak	102
	Moon Mountain District Sonoma County	55
	Petaluma Gap	52

continued . . .

. . . continued

Macro-Region?	Region	count
	Bennett Valley	50
✓	Northern Sonoma	44
✓	Clear Lake	37
	Yorkville Highlands	36
	Guenoc Valley	21
	Eagle Peak Mendocino County	16
	Wild Horse Valley	14
	Fountaingrove District	10
	Chiles Valley	9
	Suisun Valley	8
	Potter Valley	7
	Pine Mountain-Cloverdale Peak	6
	High Valley	2
	Solano County Green Valley	2
Central Coast AVAs		
	Sta. Rita Hills	1,508
	Santa Lucia Highlands	1,350
✓	Paso Robles	1,104
	Santa Maria Valley	932
✓	Central Coast	659
✓	Santa Cruz Mountains	554
✓	Monterey	440
	Edna Valley	412
	Arroyo Grande Valley	279
✓	Santa Ynez Valley	228
	Chalone	154
	Mt. Harlan	147
	Livermore Valley	129
	Arroyo Seco	113
	Carmel Valley	45
✓	San Francisco Bay	44
	Adelaida District	35
	Happy Canyon of Santa Barbara	32
✓	Santa Clara Valley	30
	Santa Margarita Ranch	19
	Cienega Valley	18
	Paso Robles Willow Creek District	15
	Ben Lomond Mountain	8
	York Mountain	7
	San Ysidro District	5
	Ballard Canyon	2
	Paso Robles Highlands District	2

continued . . .

... continued

Macro-Region?	Region	count
	El Pomar District	1
	Hames Valley	1
	Paicines	1
✓	San Benito	1
Central Valley AVAs		
✓	Lodi	355
	Dunnigan Hills	24
✓	Clarksburg	11
	Borden Ranch	5
	River Junction	4
	Mokelumne River	1
Sierra Foothills AVAs		
✓	El Dorado	133
✓	Sierra Foothills	70
	California Shenandoah Valley	66
	Fiddletown	62
	North Yuba	16
	Fair Play	2
South Coast AVAs		
	Cucamonga Valley	22
	Temecula Valley	20
	Malibu-Newton Canyon	18
✓	South Coast	1
County and state		
✓	California	2,495
✓	Sonoma County	2,793
✓	Santa Barbara County	856
✓	Mendocino County	465
✓	Monterey County	387
✓	Amador County	325
✓	Contra Costa County	196
✓	Napa County	156
✓	San Luis Obispo County	122
✓	Lake County	91
✓	Marin County	61
✓	Red Hills Lake County	46
✓	Calaveras County	33
✓	Santa Clara County	16
✓	San Benito County	13
✓	El Dorado County	9
✓	San Mateo County	8
✓	Nevada County	7

continued ...

. . . continued

Macro-Region?	Region	count
✓	San Joaquin County	7
✓	Trinity County	5
✓	Yolo County	4
✓	Kelsey Bench-Lake County	3
✓	San Diego County	3
✓	Santa Cruz County	3
✓	Solano County	3
✓	Los Angeles County	2
✓	Stanislaus County	2
✓	Tehama County	2
✓	Alameda County	1
✓	Siskiyou County	1
✓	Ventura County	1

3.C Summary statistics by grape variety and region

Table 3.C.1: Average price, average score and average number of cases made per wine by location, all grape varieties

	Observations	Price	Score	Cases made
	no.	2022 dollars	points	no.
North Coast AVAs	38,195	76.44	88.30	3,820
Central Coast & Santa Cruz AVAs	8,275	51.64	87.20	6,153
Sierra Foothills AVAs	349	31.03	84.90	2,036
Central Valley AVAs	400	25.77	84.40	24,509
South Coast AVAs	61	37.51	83.67	6,315
County	5,621	40.27	86.20	14,827
California State	2,495	23.01	83.39	64,743
Total	55,396	65.97	87.65	7,940

Table 3.C.2: Average price, average score and average number of cases made per wine by location, Cabernet Sauvignon

	Observations	Price	Score	Cases made
	no.	2022 dollars	points	no.
North Coast AVAs	12,663	116.98	88.94	4,430
Central Coast & Santa Cruz AVAs	797	53.70	85.44	15,372
Sierra Foothills AVAs	36	28.81	80.92	2,507
Central Valley AVAs	27	22.00	83.11	32,630
South Coast AVAs	11	58.95	84.55	1,119
County	747	52.29	85.94	23,419
California State	505	24.93	82.71	72,628
Total	14,706	106.72	88.35	8,116

Table 3.C.3: Average price, average score and average number of cases made per wine by location, Chardonnay

	Observations	Price	Score	Cases made
	no.	2022 dollars	points	no.
North Coast AVAs	8,242	53.79	88.58	5,724
Central Coast & Santa Cruz AVAs	2,385	41.64	87.51	9,854
Sierra Foothills AVAs	29	24.91	82.69	3,971
Central Valley AVAs	49	18.97	84.08	47,475
South Coast AVAs	16	25.83	80.31	21,283
County	1,584	36.79	86.86	22,668
California State	667	21.01	83.69	84,199
Total	12,972	47.56	87.88	12,362

Table 3.C.4: Average price, average score and average number of cases made per wine by location, Merlot

	Observations	Price	Score	Cases made
	no.	2022 dollars	points	no.
North Coast AVAs	3,305	56.70	86.60	4,955
Central Coast & Santa Cruz AVAs	418	35.11	83.71	9,777
Sierra Foothills AVAs	22	27.22	80.50	1,381
Central Valley AVAs	21	16.97	80.05	24,159
South Coast AVAs	9	50.29	82.11	814
County	615	46.92	85.03	13,931
California State	436	18.47	81.93	79,479
Total	4,826	49.81	85.66	12,924

Table 3.C.5: Average price, average score and average number of cases made per wine by location, Pinot Noir

	Observations	Price	Score	Cases made
	no.	2022 dollars	points	no.
North Coast AVAs	8,931	64.90	88.26	1,966
Central Coast & Santa Cruz AVAs	3,993	61.20	87.89	2,442
Central Valley AVAs	4	16.51	82.50	9,090
County	1,183	43.12	86.06	10,279
California State	348	31.07	84.47	31,146
Total	14,459	61.27	87.89	3,422

Table 3.C.6: Average price, average score and average number of cases made per wine by location, Zinfandel

	Observations	Price	Score	Cases made
	no.	2022 dollars	points	no.
North Coast AVAs	5,054	45.15	87.41	1,750
Central Coast & Santa Cruz AVAs	682	38.31	86.25	2,213
Sierra Foothills AVAs	262	32.33	86.06	1,827
Central Valley AVAs	299	27.96	84.90	20,259
South Coast AVAs	25	30.95	86.00	1,381
County	1,492	32.94	86.21	6,450
California State	539	22.16	84.16	44,084
Total	8,353	39.87	86.76	5,832

3.D Details on restricted cubic spline specification

We chose $K = 4$ knot locations — $\kappa_1 < \kappa_2 < \kappa_3 < \kappa_4$ — that split the distribution of weather deviations d into quintiles.

We defined:

$$V_i = \frac{(d-\kappa_i)_+^3 - (\kappa_K - \kappa_{K-1})^{-1}\left\{(d-\kappa_{K-1})_+^3(\kappa_K - \kappa_i) - (d-\kappa_K)_+^3(\kappa_{K-1} - \kappa_i)\right\}}{(\kappa_K - \kappa_1)^2} \tag{3.3}$$

for $i = 1, 2$

See Stone & Koo (1985) for details on how the linearity constraints on a restricted cubic spline are used to derive V_i.

3.E Results for other varieties

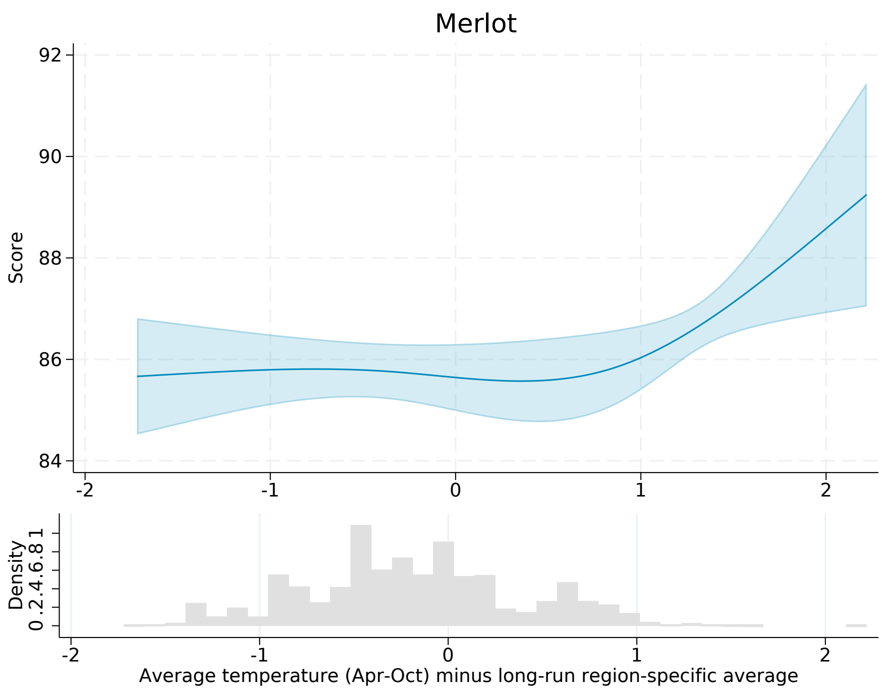

(a) Score

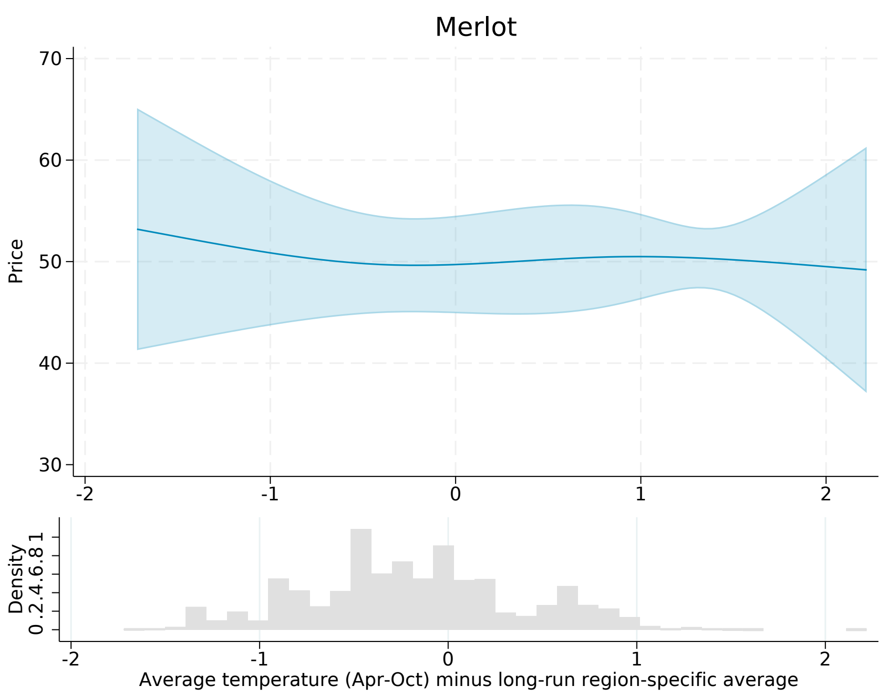

(b) Price (2022 dollars per bottle)

Figure 3.E.1: Effect of a deviation in temperature from local climate on Merlot wine scores and prices

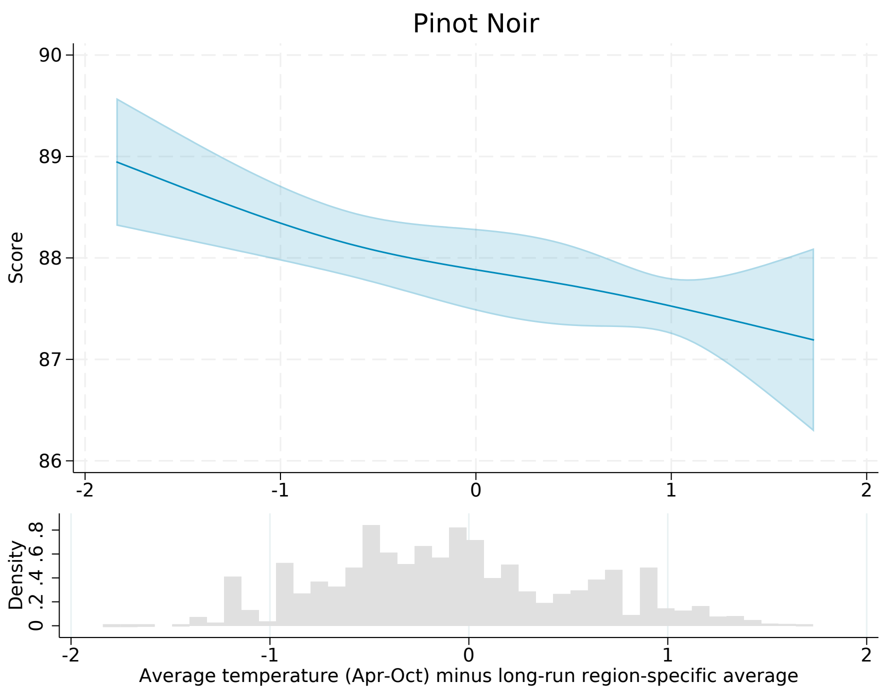

(a) Score

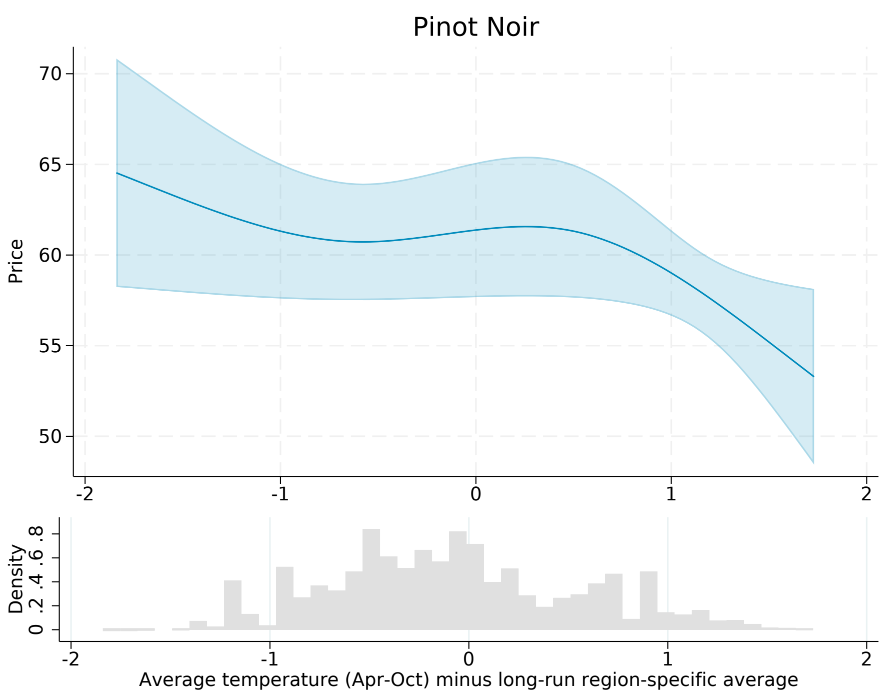

(b) Price (2022 dollars per bottle)

Figure 3.E.2: Effect of a deviation in temperature from local climate on Pinot Noir wine scores and prices

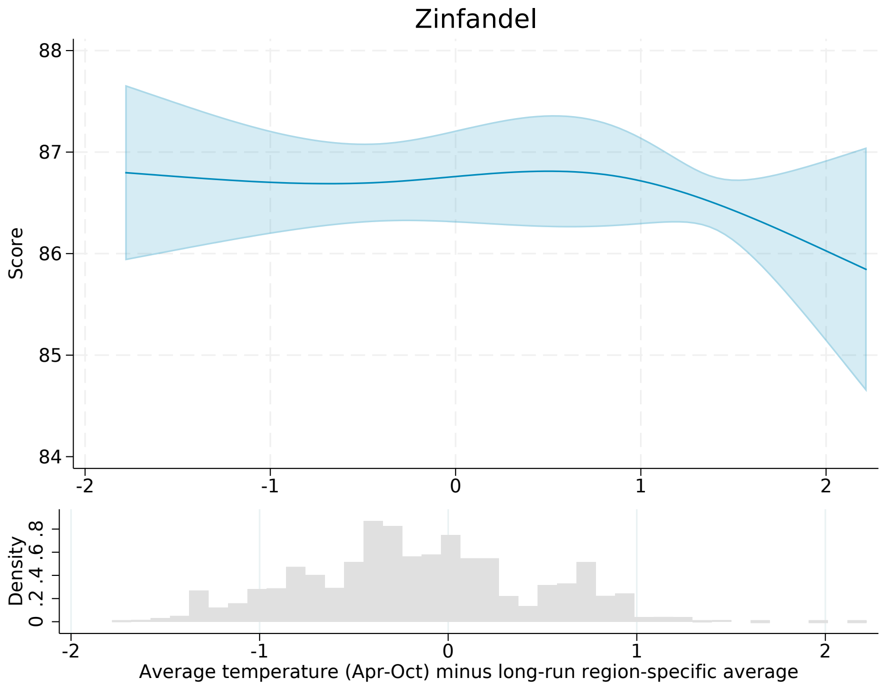

(a) Score

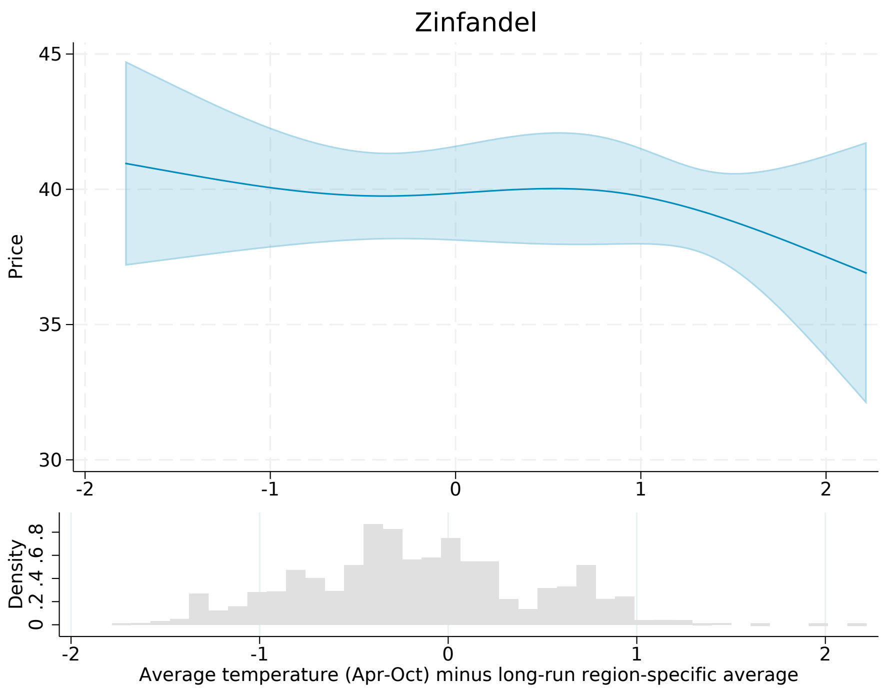

(b) Price (2022 dollars per bottle)

Figure 3.E.3: Effect of a deviation in temperature from local climate on Zinfandel wine scores and prices

3.F Robustness checks

We considered several approaches to aggregating daily observations of temperature. For most of our analysis we use growing season average temperature, which we calculate by taking the average of daily observations of average daily temperature from April to October. In this robustness check, we use an alternative measure called degree days. Following Ortiz-Bobea (2021), we calculate growing degree days (GDD) as the time spent between 10°C and 35°C, summed from April to October. Overall, results using GDD (Figure 3.F.1 and Figure 3.F.2) are similar to results using average temperature. This is expected because average temperatures and GDD are highly correlated.

There are also different ways to define climate. Here, we show results where climate is defined as the 10-year moving average of regional average temperature. Overall results are virtually unchanged (Figure 3.F.3 and Figure 3.F.4).

Finally, we show results where the dependent variables (price and score) are in log form instead of levels. Prices are skewed so one may argue that taking the log of price before estimation is preferred. Overall, we find that the shape of the estimates is very similar (Figure 3.F.5 and Figure 3.F.6). We opt to show the main results in levels for ease of interpretation.

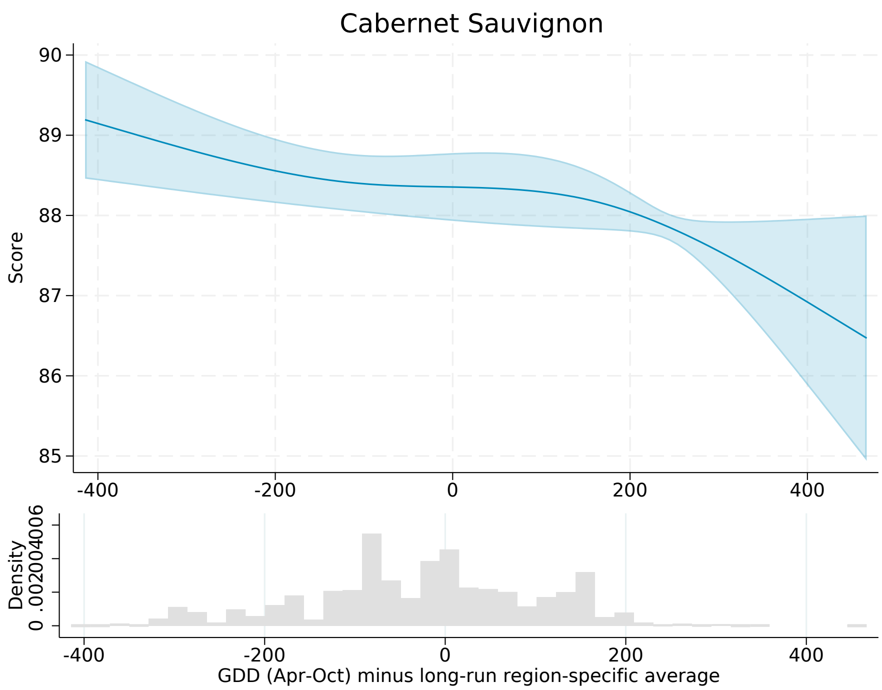

(a) Score

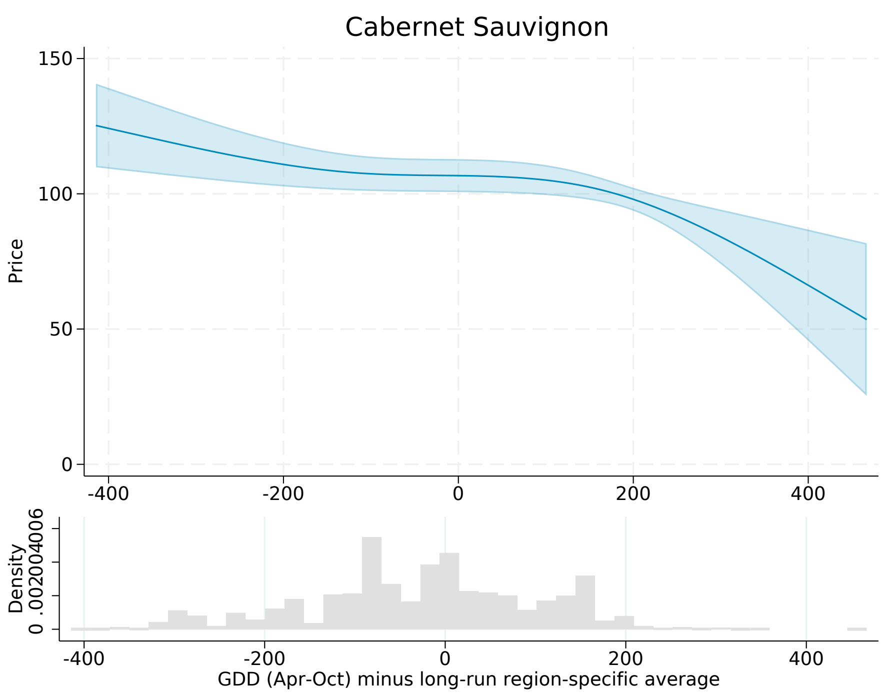

(b) Price (2022 dollars per bottle)

Figure 3.F.1: Effect of a deviation in GDD from local climate on Cabernet Sauvignon wine scores and prices

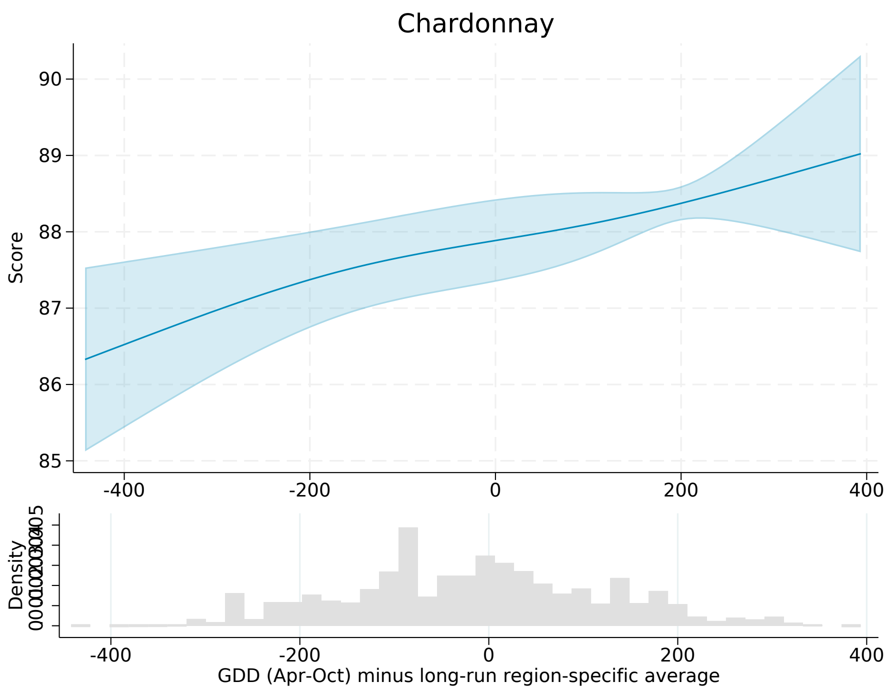

(a) Score

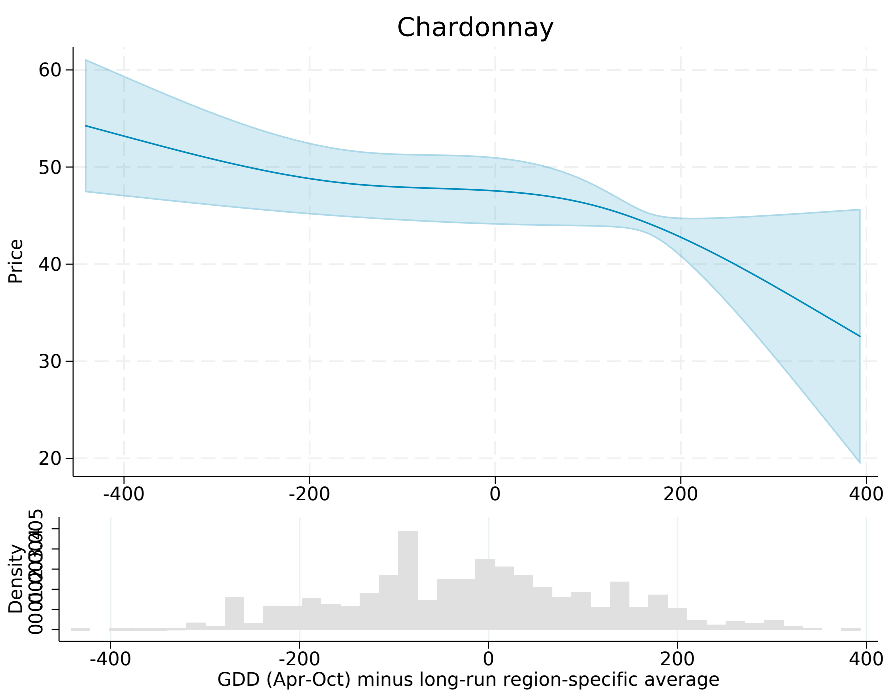

(b) Price (2022 dollars per bottle)

Figure 3.F.2: Effect of a deviation in GDD from local climate on Chardonnay wine scores and prices

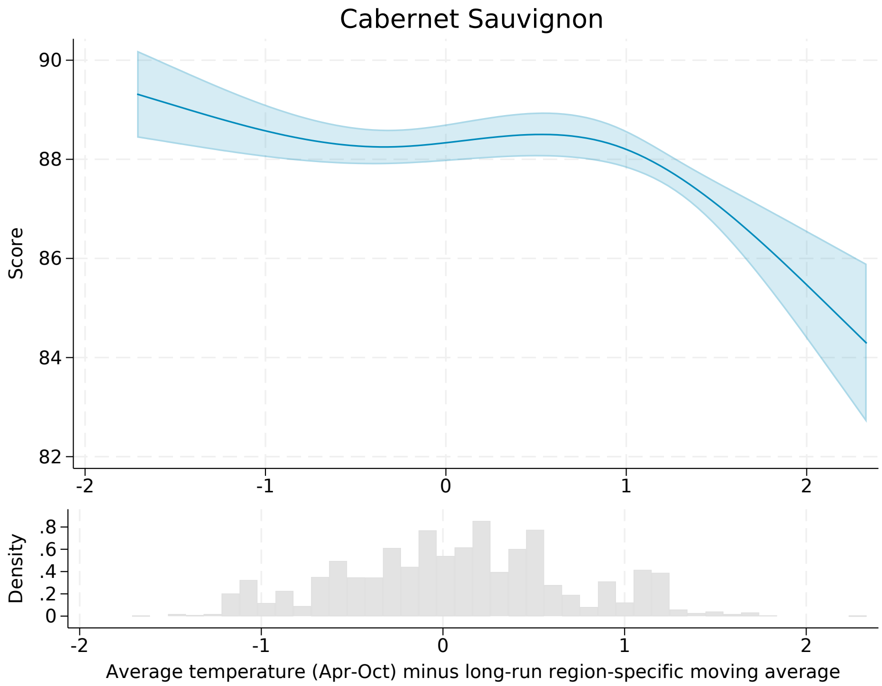

(a) Score

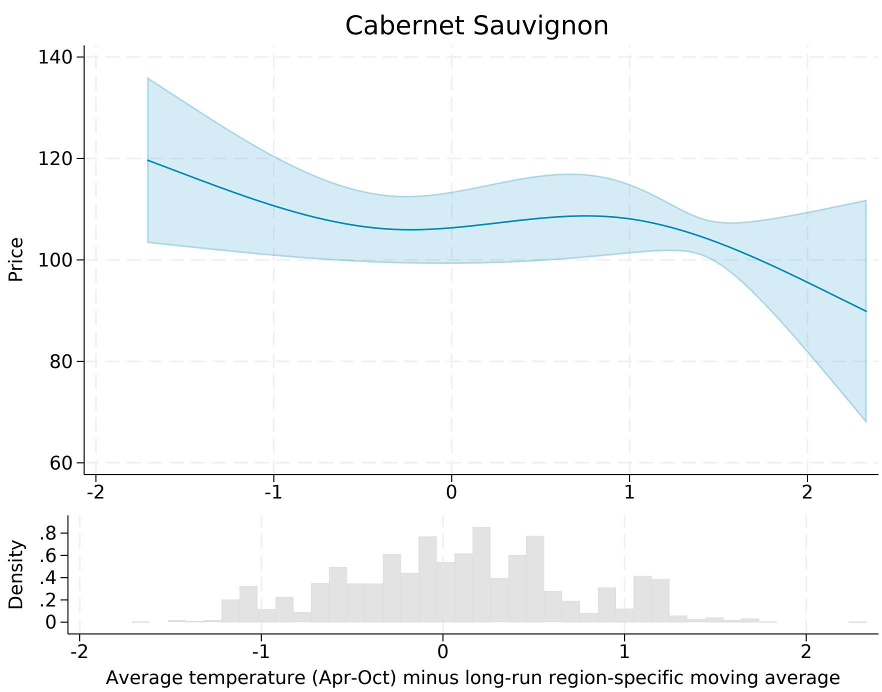

(b) Price (2022 dollars per bottle)

Figure 3.F.3: Effect of a deviation in temperature from local climate (10-year moving average) on Cabernet Sauvignon wine scores and prices

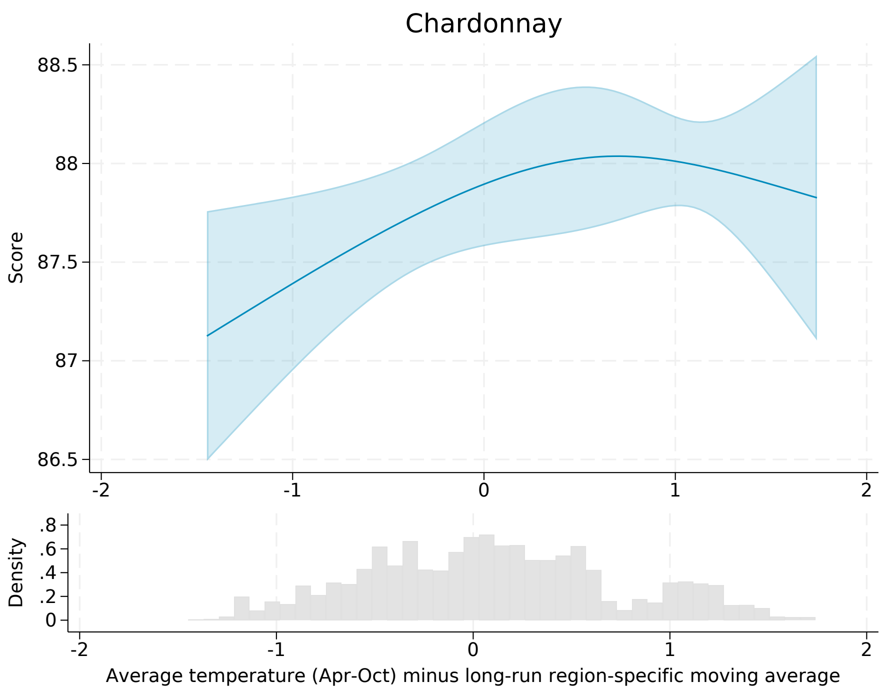

(a) Score

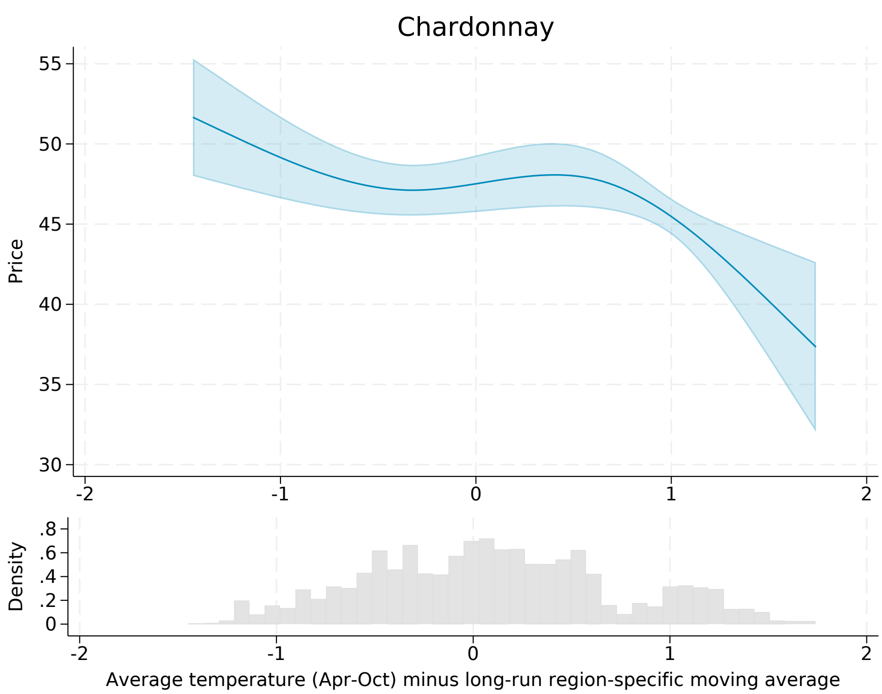

(b) Price (2022 dollars per bottle)

Figure 3.F.4: Effect of a deviation in temperature from local climate (10-year moving average) on Chardonnay wine scores and prices

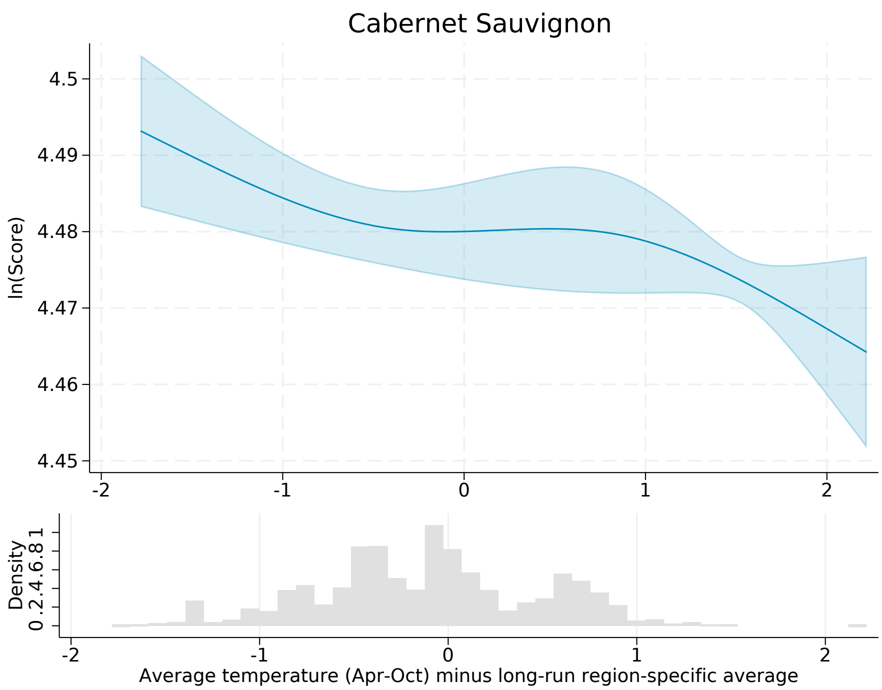

(a) log of Score

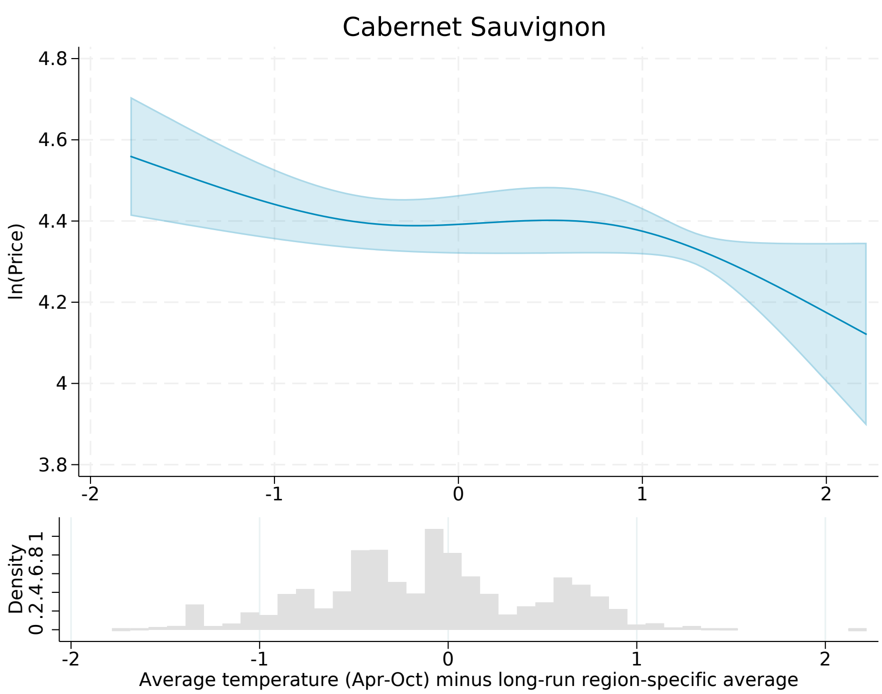

(b) log of Price (2022 dollars per bottle)

Figure 3.F.5: Effect of a deviation in temperature from local climate on Cabernet Sauvignon wine scores and prices (in logs)

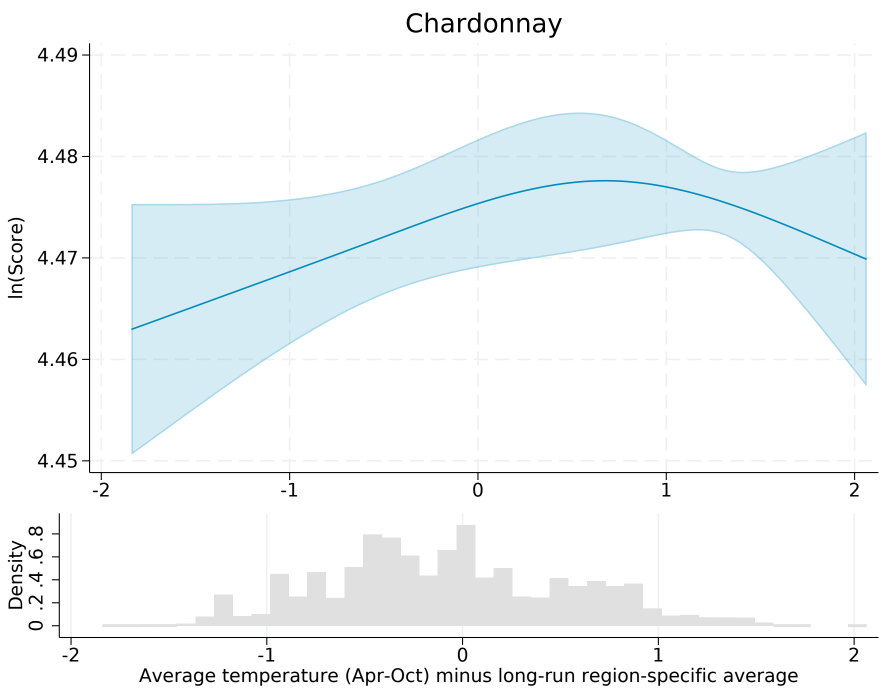

(a) log of Score

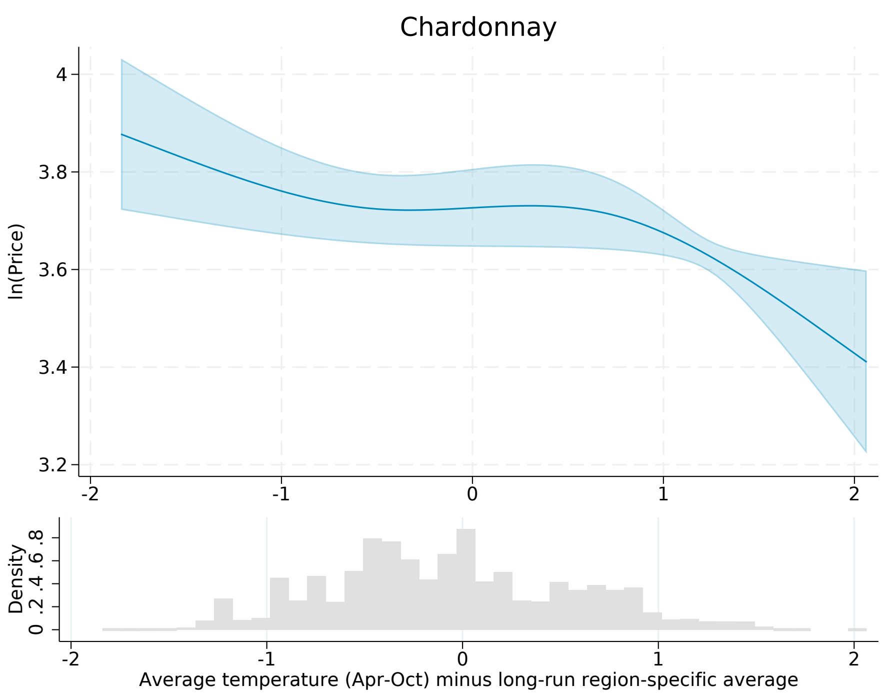

(b) log of Price (2022 dollars per bottle)

Figure 3.F.6: Effect of a deviation in temperature from local climate on Chardonnay wine scores and prices (in logs)