Costs of Climate Adaptation: Evidence From French Agriculture*

Tristan Du Puy and Jeffrey G. Shrader

[1]Columbia University

November 27, 2024

Abstract: Adaptation costs are one of the main missing elements from the existing literature on the effects of climate change. Policy to address climate change depends on how costly it is for people to adapt, but a lack of cost-related data means that such estimates are rare. In this paper, we use uniquely rich data on agriculture in France to provide novel, direct estimates of the marginal cost of adapting to changing temperatures. The dataset is a farm-level panel with measures of outputs, inputs, and prices from 1994–2018. We merge the farm data with measures of realized and forecasted weather. Controlling for realized weather, we use forecasts as information shocks to estimate costs of ex ante adaptation. For the year that heat shocks arrive, we find that for the average farm in France the cost of adaptation is low. In contrast, the benefits of adaptation are large. This difference is driven by the behavior that farms engage in when responding to forecasts. They mainly use the forecasts to change the timing of planting and harvesting decisions, as well as their crop mix, rather than to change costly inputs. The large, observed difference between the benefits and costs of adaptation contradicts a widely used sufficient statistic approach to estimating climate damages. We further show that while costs of adaptation are low initially, future periods are marked by increases in costs and decreases in profit, potentially capturing the dynamic nature of farms' adaptation strategies. Finally, we show that forms of adaptation currently implemented by farmers are less likely to remain relevant in a warmer world.

*We are grateful for comments from Andrew Hultgren, Derek Lemoine, and Wolfram Schlenker as well as numerous participants in seminars at Columbia University, the MIT/NOVA School of Business Lisbon Workshop on Climate Change Adaptation, NBER EEE, and the Transatlantic Kellen Conference Series. We thank the French Ministry of Agriculture and INSEE for providing access to the data, as well as the CASD for secured data access. We gratefully acknowledge funding from the Columbia Center for Environmental Economics and Policy, the Columbia Climate School, and the Columbia Center for Political Economy. All remaining errors are ours.

1 Introduction

The potential consequences of climate change are a central concern for our century. A large and growing body of research seeks to understand the effects of climate change on the economy. One of the longstanding challenges to this work is estimating the benefits and costs of adaptation—the action people take to prepare for or adjust to a changing climate. As argued by Carleton et al. (2022) and Kelly et al. (2005), we should be interested in the sum composed of the damages of climate change and the costs of adaptation, in part because actions taken to mitigate climate change would not only prevent damage but also save people from paying adaptation costs. In this paper, we provide novel estimates of the costs of adaptation in the agricultural sector.

The challenge to estimating costs of adaptation is primarily a lack of data. Even in agriculture, one of the sectors of the economy most directly affected by climate change, and therefore widely studied in the climate economics literature, data on costs are typically unavailable. Crop yields, profits and land values are commonly observed and have been used to estimate the damages from changing weather (e.g. Mendelsohn et al. (1994), Deschênes and Greenstone (2007), Schlenker and Roberts (2009)). But the input levels and prices needed to recover adaptation costs are rarely observed. This lack of cost data is not limited to the agricultural setting. For example, recent work on temperature-related mortality estimates adaptation costs using a revealed preference approach, where they are bounded by estimates of adaptation benefits (Carleton et al. (2022)).

We overcome this challenge by using a uniquely rich panel dataset of French agriculture for 1994-2018. We combine it with realized and forecasted weather data, in order to isolate the farm-level variation related of ex-ante adaptation. The administrative agricultural data contains accounting measures for both costs and revenue directly at the farm level, and we can hence recover independent estimates of the marginal costs and benefits of adaptation. Observing data at the farm level, we can also include farm fixed effects, a more granular control than previous studies and one that helps address the challenges set by recent work highlighting the potential endogeneity of local climate (Braun and Schlenker (2023)). We further observe store level input price data, and can disentangle responses in input prices from responses in levels. Our farm sample is representative of the French market, and farms are observed for multiple seasons in a row. These repeated observations further allow us to recover the consequences of adaptation over time, both in the season when adaptation action is undertaken, and in the ones that follow.

We find that farms undertake costly adaptation strategies in response to information shocks about incoming heat shocks. Controlling for realized weather, we use forecasts to isolate the farmers' responses to weather related news shocks. These clearly identify the cost, revenue and profit consequences of ex-ante adaptation, holding the direct effect of weather constant. We further identify the main adaptation strategies undertaken by farms, paying attention to timing and crop choices, as well as input usage. Within the period of the shock, we show adaptation translates into positive revenue gains, with no visible effect on production costs. These are however unevenly distributed over time. In the longer run, adaptation proves costly and the net present discounted consequence of adaptation on profit is negative. Conditional on farmer rationality, these provide a lower bound for the counterfactual no-adaptation effect of similar variation in realized heat.

The first part of this paper looks at farm behavior within the period of the shock. Within period, farmers are able to leverage heat forecasts to generate relatively large revenue gains at small costs. A one standard deviation increase in forecasted heating degree days leads to a .8% increase in revenue, while costs have no statistically significant response. In comparison, farms respond to more moderate temperature forecasts by increasing the scale of their production, at an equivalent marginal cost. Both these results hold across alternative specifications and robustness tests. Making use of field level data, we show that adaptation to heat shocks in part takes the form of ex-ante timing changes affecting different steps of the growing season. We further provide evidence for crop switching. We identify two potential patterns of switching from high heat-sensitive crops towards less sensitive ones: away from sunflower and towards colza, away from corn and towards wheat. We further observe increases in the area allocated to peas and beans in line with the results of Aragón et al. (2021).

The differences in the cost and revenue consequences of ex ante farm adaptation naturally translate into net profit gains of ex ante adaptation. In a second section, we investigate the farm-level profit responses to forecasted and realized weather. We further discuss the implications of these net positive profit responses for the applicability of the static envelope theorem to our context. In a climate change damages setting, the envelope theorem states that the cost and benefit consequences of re-optimizing production in response to marginal weather variation should offset each other. In our empirical context, this correspondence does not hold. The envelope theorem has been invoked by recent studies of climate change to provide estimates of adaptation costs (Deschênes and Greenstone (2007), Hsiang (2016), Deryugina and Hsiang (2017), Merel et al. (2022), Carleton et al. (2022)).[1] In particular, the envelope theorem allows one to equate marginal benefits and marginal costs of adaptation, and hence to infer a measure of the costs of adaptation.

We discuss the potential reasons for the breach of the static envelope theorem. A first element is that growing and heating degree days temperature aggregations bring together marginal and non-marginal variations in weather. Over our period of study, growing degree days are mostly composed of marginal variations,[2] while heating degree days are mostly composed of non-marginal variation. This source of variation tracks with our observed, statistically indistinguishable cost and benefit responses to forecasted growing degree days (GDDs), and net-positive profit responses to forecasted heating degree days (HDDs). A second element of answer is the difference in nature of the response to forecasted growing and heating degree days. Responses to forecasted GDDs imply some timing changes, as well as input responses, while responses to forecasted HDDs are uniquely driven by significant timing and cropping changes. These changes are likely to drive non-continuous changes in farm-level profits, and are hence less suited to the assumptions of usual envelope theorems.

Focusing on the within-period consequences of ex ante adaptation does not, however, provide a complete picture of the consequences of adaptation. We further investigate the profit and cost response to ex ante adaptation over time. While the period of the forecast shock is marked by net profit gains, the following two periods show significant profit losses. The cost movement parallels the evolution of profit, with increases in costs in future periods. Overall, it seems that farmers' ex ante adaptation leads to a net decrease in their present discounted flow of profit. Because we identify crop switching as one of the main adaptation strategies used by farmers, the presence of these future costs suggests that changes to one's crop mix entail switching costs. The presence of such costs has been central to the growing dynamic land use literature (Livingston et al. (2008), Scott (2013), Burlig et al. (2024)) but we are one of the first papers to recover evidence for them in a reduced form setting.

In a last step, we investigate whether the farm responses to forecasted HDDs are likely to be good indicators of future adaptation strategies to climate change in France. Specifically, we highlight the heterogeneity of responses to the scale of heat shocks experienced by farms. To do so, we select the top decile of the distribution of department-level HDD realizations, and re-run our analysis on that subset of bad years. We find that costs and benefits of adaptation evolve significantly as we move to these hotter years. We stress the appearance of costly on-the-spot adaptation in the form of input increases and storage decreases, which, jointly, appear to be successful at mitigating the negative impacts of heat realizations. ex ante adaptation shifts as well. From relying mostly on timing changes to generate positive profit gains, ex ante adaptation now happens earlier, reacting to two-month ahead forecasts, and implies a down-scaling of production and an associated decrease in production costs.

[1] See reviews from Carleton and Hsiang (2016); Auffhammer (2018)

[2] We code marginal variations as variations more than 1.96 standard deviations away from the usual hour-day-month realization over the last thirty years.

Related Literature: This paper is first and foremost related to the literature on climate adaptation and climate damages. This literature has long debated how to provide comprehensive estimates of the consequences of climate change on human systems. This in turn has implied accounting for the costs of adaptation, or isolating measures of direct weather damages (hence indirectly of climate damages).

The literature on climate damages has grown rapidly over the last decades, crossing from Ricardian approaches into global dynamic general equilibrium models and fine scale panel studies of context-specific dose-response functions.[3] Our paper stands out in this literature by the unusual precision of the data directly measured at the agent-level, and within a repeated panel format. While studies of adaptation to climate change have previously relied on farm-level data, for example Aragón et al. (2021) for smallholder farms in Peru, our repeated structure and the depth of the data collected is to the best of our knowledge unprecedented. In the context of developed economies, most research has relied on aggregated data for which yields or profits are available (Schlenker and Roberts (2009) for example). In France, previous research has been conducted at the department-level and over a longer period of time by Gammans et al. (2017), but has focused on yields specifically, while research that has been done at the establishment-level has focused on a small subset of farms (Bareille and Chakir (2023)). We build on these papers by distinguishing costs and revenue adaptive responses.

Our paper also builds on the subset of the literature which has approached adaptation with a keen focus on the definition of farmers' climate beliefs. Such papers have highlighted how the nature of agents' beliefs about the climate process structures what can be interpreted as adaptive behaviour to changes in climate. Burke and Emerick (2016) illustrate clearly how the definition of beliefs drives our understanding of adaptation. Prior to Burke and Emerick (2016), Kelly et al. (2005) defined two sets of costs related to environmental changes: transition costs and equilibrium costs. The first ones being incurred as the economy moves along the transition path towards a new equilibrium where the new environmental context is fully taken into account. Kala (2019) goes further by comparing different learning models for the timing of the monsoon in India. She makes the point that recovering farmers' learning behavior can depend on our modeling of the objective that they maximize. What we measure as the extent of their adaption to changing climate patterns relies on the behavioral model and the objective function assigned to them. Finally, Shrader (2023) offers a way to use weather forecasts to disentangle adaptation effects from direct climate damages. Our current work builds on this. Once controlling for weather realizations, we use forecasts as shocks to the agents' beliefs that in turn drive the agents' adaptive actions. Related to this literature, we show that here, not controlling for weather forecasts induces an upward bias in the measurement of the profit impact of extreme heat—extreme heat seems less damaging that it really is. We run a distributed lag model that includes all the forecasts available during the growing season (at all the possible lead values), and hence attempts to capture as much as possible of the agents' beliefs formation and adaptive behavior, and their consequences on farm profit. Running such a regression significantly increases the negative impact of extreme heat on French farm profits. This is suggestive evidence that as we better control for the indirect consequences of adaptation on profit, we can better isolate the negative consequences of marginal variation in extreme heat.

Finally, we note that our work relates very closely to two papers which have approached climate damages without relying on the envelope theorem. The first of them is Guo and Costello (2013), which looks at the impact of climate change in a forestry context with both discrete and continuous decision making. In their context, foresters can implement continuous adaptation actions for which envelope conditions apply, and discrete actions for which they do not. They derive a model where marginal variation in weather can lead to arbitrarily large impacts on profit by inducing discrete changes in the type of tree species planted by foresters. Empirically, they find that extensive margin (discrete) adaptive actions can have large profit responses, but typically induce a net-zero profit response. Similarly, Bareille and Chakir (2023) also rely on

[3]See Auffhammer (2018) for a review.

structural methods to disentangle adaptation effects from direct weather effects, but their method requires a complete enumeration and modeling of adaptation methods, as well as instruments separately identifying adaptation effects from the weather impacts. Mirroring the discussions in the sufficient statistic literature, by-passing the envelope theorem is a difficult task which requires a structural setting and implies trading some assumptions for others. We see our work as shedding additional light on the trade-offs between structural and envelope-free, and reduced form and envelope-dependent methods.

The rest of the paper proceeds as follows. In section 2, we describe our data sources, and in section 3 we link our production model to the estimating equations used in our empirical strategy. In section 4, we present the main results on the marginal benefits and costs of adaptation to weather shocks. section 5 examines the effects on profit, laying out a test for sufficient statistics assumptions, and **??** examines limits to the adaptation found in the prior section. Finally, section 6 concludes the paper.

2 Empirical Context and Data

2.1 Agriculture in France

We focus on French cereal, oil, and protein crop production for our analysis. These sectors of the agricultural industry are of interest in a climate change context, as they are known to be clearly sensitive to variations in weather. They are also less likely to be shielded from variations in weather through the use of greenhouses—as might be the case with horticulture—and less heavily irrigated than other crops. In France in 2015, only 5.8% of the total used agricultural area was irrigated (Colas-Belcour et al., 2015). This low exposure to irrigation in the data also reduces potential concerns of endogeneity linking irrigation and local weather, as discussed in Braun and Schlenker (2023), and also matches recommendations to avoid studying the impact of weather shocks in areas where irrigation is likely to be subsidized (Schlenker et al., 2005).

French agriculture is mechanized, and farms have seen decades-long increases in their average size (to reach 69ha in the 2020 agricultural census), and in their yields as discussed in Schauberger et al. (2022). France also possesses the largest agricultural area within the EU, and is one of the largest cereal exporter in the world, exporting to other EU member states, Morocco, Algeria, and other countries. As such, exposure to climate change and the measurement of its potential damages on agricultural production is a key policy question.

French farms benefit from a moderate climate, but farm production is still sensitive to weather shocks. There is significant variation in weather due to the frequent arrival of low-pressure areas from the Atlantic (Canal (2015)). The Mediterranean regions of France also suffer from high temperatures, water stress over the summer, and heavy localized rainfall.

Seasonal forecasts have been developed for France since the 1990s, but high-quality coupled atmosphere-ocean models started to be used for seasonal forecasts only in 1999, allowing to move beyond a four-month lead for the predictions (Canal (2015)). In this paper, we will use hindcasts produced by the European Center for Medium-Range Weather Forecasts (ECMWF), up to 1994, and will use leads up to five months. Most of the analysis however will focus on one-month-ahead forecasts. These should hence have been in the farms' information set from 1994 onward. Today, different companies offer an access to month-ahead forecasts throughout the season, and the EU's Joint Research Center has been providing real-time cereal yield predictions based on the Crop Yield Forecasting System with its MARS bulletin since 2007.

Below we show the allocation of land in our sample for the four main crops, accounting together for about 85% of our sample's agricultural area.[4]

[4] We note that the thresholds used for crop-specific tolerance are only indicative. Tolerance to heat varies across the growing stages of each crop, and also relate among else to drought conditions. We take these as only indicative that wheat is more heat tolerant than other crops, and is likely to be less responsive to extreme heat events observed for France in our sample. References are Gammans et al. (2017), Schlenker and Roberts (2009), Elferjani and Soolanayakanahally (2018).

Table 1: Crop mix composition

	Share	Mean Temperature	Threshold	Reference
Wheat	47.89	10.2	33C	Gammans et al, 2017
Corn	15.29	10.5	29C	Schlenker and Roberts, 2009
Barley (Winter)	11.38	10.0	33C	Gammans et al, 2017
Rapeseed	11.26	10.2	29C	Elferjani et al, 2018

Notes. We compute the shares of cropland allocated to each crop in our dataset, and show the four largest ones. None of the remaining crops account for more than 5% each of the total land considered. We show the associated average temperature, among farms growing the crop, and a threshold for heat damages taken from the literature.

This table is indicative of the fact that wheat plays the largest role in the French agricultural system. Heat thresholds taken from the economic and agronomic literature are shown in order to indicate that wheat is relatively more heat tolerant than other crops. As such, we expect farms in our sample to be relatively heat tolerant.

We run regressions similar to the ones of Schlenker and Roberts (2009) to describe the general impact of weather shocks in French agricultural production, and as visible in our sample. We use a restricted cubic spline to model the non-linear relation between exposure to temperature and outcomes, include farm fixed effects, region-specific quadratic time trends. We show the results in subsubsection A.1.4. Output shows a non-linear relation with temperature, with total output remaining flat and decreasing sharply after 30C°. Crop yields show heterogeneous relations with temperature, both wheat and colza show no clear relations with temperature, while corn and sunflower yields decrease sharply at high temperatures.

We further compute unconditional and conditional growing season mean temperature realizations at the department level, for our period of study. Figure A20 shows splines describing the evolution over time of these realizations. These are helpful to characterize the average climate in France. On average, temperatures remain around the 10°C line, with little change over our period of study. On average, below 0°C temperature are not very negative, and on average extreme heat temperature remain around the 30°C level. As we have seen above, these extreme heat events for the French context are likely to hurt some crops, but will have a more ambiguous impact on wheat production. Only extreme heat shocks for the French context would significantly impact wheat production.

Figure A6 in the annex further shows that while there is more dispersion across French department in extreme heat events, the general dispersion of average conditional temperatures remains moderate. We also provide maps showing the geographic dispersion in growing and heating degree days in France over our time period in the annex.

We end this section with a discussion of two institutional features of interest for our study. The first one is that of agricultural subsidies. Agricultural policy has been centralized at the European Union since the late 1960s under the umbrella of the Common Agricultural Policy (CAP), although more recent iterations of the CAP have allowed for some flexibility in its country-level implementation. For our period of interest, the main types of subsidies accessible through the CAP to row crop farmers will be in the form of subsidies distributed per hectares of sown and harvested land. These should hence be relatively insensitive to the heat forecasts on which we focus. The second institutional element is that of agricultural insurance. In France, the main insurance scheme is one of harvest coverage—where extreme weather events leading to substantial losses in harvest lead to compensation. We note however that in 2023, the French government indicated that only 17% of the total agricultural area in France was insured.[5] This is a stark difference with

[5]See https://agriculture.gouv.fr/animation-la-reforme-de-lassurance-recolte

the US context, where insurance subsidies constitute one of the main schemes for agricultural intervention. To account for both these features, we will use two profit variables throughout the paper. The first one will be a measure of value added, roughly accounting for the value of production minus its costs. The second one is a more general measure of establishment profits which accounts for expenses in insurance, insurance compensation, as well as subsidies received during the season. The comparison of these two results should help confirm that our results are not driven by any of these two institutional features. We give more precise definitions of all our variables in subsection B.1.

2.2 Agricultural Data

2.2.1 Farm Data

We use the European Union's Farm Accountancy Data Network (FADN) for France, which is a subset of an EU-wide annual panel of farms. Started in 1968 in France, it currently corresponds to a sample of around 7,000 farms meant to be statistically representative of French commercial farming. [6]

The FADN contains precise accounting data on the farm's activities. For our analysis, we focus on farms mainly producing cereal and oil crops. Specifically, we use a sub-sample of farms whose sales in cereals, sunflower and rapeseed account for at least 50% of their total yearly sales. Soybean and protein crops correspond to a small fraction of large field crops in France, which justifies our focus on cereals and oil crops. This has the main consequence of excluding animal, as well as vegetable, fruit and wine farming.

For these selected farms, the FADN indicates the total and crop-specific sales, total farm-level spending on inputs per type, and different values for profit. It also indicates crop-specific production values in quantity, the amount of production sold (in quantity and value), the total amount sold per crop (accounting for potential storage), and the amount of land allocated to each crop. We describe how the different accounting variables are constructed in the appendix section B.1.

We geocode the farms at the department level, and match them with department-level weather data. While farms locations at the commune-level are available from 2000 onward, the forecast data is too coarse for such granularity, and we choose to measure realized and forecasted weather at a similar level. We also note that there are 101 departments in France, which make for slightly larger entities than US counties. We perform a robustness analysis where realized weather is computed at the commune-level and forecasts kept at the department-one over the 2000-2018 period.

2.2.2 Input Price Data

Our input price data comes from two sources. From the Land Market Value survey (Valeur Venale des Terres) for land prices, and from the Observational Survey of Intermediate Consumption Costs Necessary to Farmers (EPCIA) for seeds, fertilizers and pesticides. EPCIA is mandated by the European Commission, and is used to build Laspeyres price indices for agricultural intermediary goods in France. The EPCIA records price series for a representative amount of goods within each input category.[7] Each series is defined according to the nature of the product sold, to its sales condition and the firm that sells it.[8] We localise the sale points at the department level, and match them with department-level weather data. We also use

[6]The FADN is not representative of all agriculture, but only of commercial farming. The definition of a commercial farm changed in 2010, but this only led to the changing of the the rules for choosing replacements for the farms leaving the sample, and not to an overhaul of the sample population itself. Before 2010, a commercial farm was defined as a farm with a unique manager, which sells more than half of their production, and whose manager's working hours correspond to at least 75% of their total annual work hours. Finally, farms with less than 5ha of land were removed from the targeted population if they were not specialized. In 2000, there were 380,000 such farms recorded in the Agricultural Census out of 663,800, but together they accounted for 95% of the country's total agricultural production. From 2010 onward, the working hours requirement was removed, and the 5ha threshold was replaced by a requirement that farms have a production capacity of at least €25,000.

[7]The relative number of goods sampled within a category is proportional to that category's relative sales.

[8]The spread of one good's series across firms is proportional to the firms' market share for the sale of that specific good.

the Laspeyres price indices derived by the INSEE from the EPCIA in order to deflate the FADN farm-level input bills for seeds, fertilizers and pesticides. Specifically we use the Ipampa price index series, from 1994 to 2020. Figure A14 compares input price indices from Ipampa with a time series for nitrogen-based fertilizers producers prices from Fred, and shows the strong correlations in fluctuations. This should confirm that the indices are efficient at tracking variations in prices. Irrigation expenses is only deflated using a regular CPI index, as we do not have a water-specific price index.

2.2.3 Land Price Data

Land price data comes from the Land Market Value yearly and department-level survey (Valeur Venale des Terres), which is fielded every year by the statistical services of the French departmental administration for agriculture and forestry. These are based on data provided by the public company in charge of land management (SAFER), which authorizes agricultural land purchases and consolidations when transactions surpass a given threshold. This data is then complemented by data provided by local notaries, and several local administrations. We use this data for 1994-2015. The data was digitized from scanned data catalogues for the first years of the series.

2.2.4 Plot Level Data

Plot level data comes from the survey on Agricultural Practices for field crops (Pratiques Cultures sur les Grandes Cultures). We use the surveys fielded in years 1994, 2001 and 2006. Plots surveyed are selected among the farms that benefit from the European Union's Common Agricultural Policy. The survey focuses on land plots defined as the set of contiguous land for which the same crop is cultivated, with homogeneous agricultural practices (fertilizer and pesticide use for example). For each crop, the survey selects the minimum number of regions covering at least 95% of that crop's production, and within each region the minimum set of departments accounting for at least 90% of the region's production. Within departments, the survey selects farms with at least .1 hectare cultivated, and less than 200ha. A unique plot is selected within each farm. For the waves that we study, around 20,000 plots are sampled each time. Importantly, this dataset is a repeated cross-section and does not allow for within plot or farm analysis.

2.3 Climate Data

2.3.1 Weather Data

Our realized climate data comes from the European Centre for Medium-Range Weather Forecasts (ECMWF)'s ERA5 reanalysis product. ERA5 gives hourly estimates of climate data, out of which we use precipitation and temperature (temperature 2m above the surface of the Earth). ERA5 combines observational weather data with model-based data into a 0.25×0.25 gridded dataset.

We extract that data at the French department level, cropping the grid with department shapes, and averaging the data using simple area weights. Using the time separability assumption common in the literature on climate impacts on agriculture, we aggregate the hourly data into growing-season observations. Temperature exposure is measured through growing degree days (GDD) and heating degree days (HDD). GDDs are computed over the $[4°, 30°]$ degree interval, and HDDs sum the realized temperature above 30°C. Thus, GDDs measure moderate temperature exposure while HDDs measure extreme heat exposure. In order to match the forecast dataset, we only use 4 daily measurements of temperature to compute the GDDs and HDDs, specifically at midnight, 6h, 12h and 18h. The GDD and HDD integrals are computed over these sets of four observations a day, and later divided by 4 in order to obtain variables which are equivalent in scale to actual growing and heating degree days computed using 24 observations per day. Growing-season precipitation is measured by summing precipitation over each day.[9]

Our relevant unit of observation being the farm, and farms being mostly multi-product production units—we aim at taking into account this heterogeneity when building our climate variables. In Table A1, we aggregate crops into three categories, and show the conditional probability that farms grow different crop categories. On average, most farms growing industrial crops like sugar beets, flax or hops will also grow cereals, and the same stands for oil and protein crops like sunflower, rapeseed or soy. In Figure A13, we show the pair-by-pair probabilities of joint production. This shows that French farms are inherently multi-product economic entities, likely maximizing their profits across product lines, rather than product-by-product. As such, our weather variables have to account for the entire set of weather shocks that might impact their production schedule. To account for this, we take an expansive definition of the growing season spanning October of the previous year, to July of the current. This also matches the winter wheat growing season, which is our main crop here.[10] We show in Figure A9 and Figure A10 the cross-sectional and cross-temporal variations in growing seasons for resp. wheat and corn in France.

The same argument justifies adopting a more expansive definition of heating degree days. Gammans et al. (2017) study the impacts of climate change on wheat and barley yields at the French department level over 1950-2014, and find that heat shocks happening over the spring and summer seasons can cause significant yield losses when they pass the 33°C threshold. Their dose-response function remains however relatively flat before the 33°C mark. Results by Schlenker and Roberts (2009) however show that corn, which accounts for 15% of the surface we study, suffers from heat from the 29°C mark onwards. We adopt 30°C as a threshold to account for this heterogeneity in crop responses, and later show robustness tests using a 28°C threshold for growing and heating degree days.

2.3.2 Seasonal Weather Forecasts

The forecast data is taken from ECMWF's SEAS5 seasonal forecasting system. Specifically, we use SEAS5 system 8 data originating from Météo France. Forecast are produced on the first of each month for the following 5,160 hours (215 days or roughly 7 months).[11] For temperature, forecasts are produced for four moments of the day located at 6h intervals from each other. They hence give an instantaneous point prediction of temperature. For rainfall, forecasts give the accumulation of rainfall every 24h. As such, the rainfall forecasts available for the second day of January will be the following: a 24h ahead forecast provided on the 1st of January, a 32 days ahead forecast produced on the first day of December, and so on until the lead value exce eds 5,160 hours. Temperature forecasts work in a similar way, but are just produced with more granular time steps.

The main issue for our purpose is that given that forecasts are produced every first of the month, different days within a month will not be provided a forecast with the same lead—here the 10th and 15th forecasts of a month will always be produced on the same day, and the 15th is likely to accumulate more uncertainty. We would ideally like to build the forecast-equivalents of our growing-season aggregates for realized weather, just for different forecast lead times. For example, the rainfall forecast for the growing season, with a constant one-month lead throughout the season (or the equivalent of the farmer's knowledge about rainfall one month in advance throughout the growing season).

We approximate this by bundling together forecasts produced one calendar month ago (what we refer to as a one-month lead for the rest of the paper), produced two calendar months ago, up to five months ago. As such, the forecasts that we aggregate into growing season observations are not homogeneous in terms of lead value, but are the closest equivalent of it that we can observe. Practically, for each year-day-hour (or

[9] We perform robustness tests where we use realized growing and heating degree day measures computed using the full set of EAR5 hourly data, and the set of four hours a day for the forecasts.

[10] Gammans et al. (2017) also define the winter wheat growing season for France as spanning October to July.

[11] ECMWF provides an ensemble of 25 forecasts, which we average.

year-day for precipitation), we rank the forecasts by their lead time and use the first-ranked forecasts for our one-month ahead measure, our second-ranked ones for the two-months ahead, and so on. A one month ahead aggregate for a given month will then aggregate over forecasts for the 2nd day of the month which amount to one-day ahead forecasts, and for the last day of the month made on average 30 days in advance. A better description for them might be that one-month ahead forecasts give forecasts issued between zero and one month ahead in advance, two-months ahead between one and two months ahead, and so on.

Similarly to the weather realization data, we aggregate the gridded forecast data into department-level observations, using area weights.

2.3.3 Forecast Accuracy

We show in Figure A1 and Figure A2 the distribution of the difference between realized and forecast weather. We note that for growing degree days, the distribution is centered near zero, but with a small, negative bias. There is a long left tail to the GDD errors. We trim the estimation sample to drop errors less than -500 (removing 120 observations in total) to avoid the effect of outliers.[12] The distribution of HDD forecasts is also tightly centered around 0, with some evidence of positive skewness, indicating that the forecast was more likely to come in too low rather than too high. Rainfall also shows an upward bias in the forecasts, but with distributions centered close to zero and largely symmetric.

The relatively narrow distribution indicates that forecasts do contain meaningful information about future weather. This provides statistical corroboration to the discussion in subsection 2.1 that forecasts contain information that should shift farmers' beliefs about weather.[13] This conclusion is further reinforced by Figure A4 and Figure A5, which show calibration plots for the forecasts of GDDs and HDDs respectively. For GDDs, the figures show that the forecasts are well calibrated at both a 1 and 2-month horizon (further horizons show similar patterns). Forecasts and realizations are highly correlated, and the scatterplot shows that forecast and realizations fall close to the 45°line. Compared to GDDs, forecasted HDDs are more likely to under-estimate the realized HDD.

We construct heating and growing degree days from 4-times-daily observations of realized and forecasted temperature. The highest and lowest temperatures during the day might occur between observations, and, therefore, our calculation could understate the realized extreme temperatures actually faced by farms. The forecasts, however, are only available at this temporal frequency, so using them to calculate degree days accurately captures the granularity of information available to farmers.

3 Empirical Strategy

To identify the cost and revenue effects of farms' responses to expected weather shocks, we use an estimating equation that regresses total costs, or total revenue on both realizations and forecasts of weather. In particular, we use the following estimating equation.

$$\begin{aligned} y_{jt} = &\beta_1^w \mathrm{GDD}_{d(j)t} + \beta_2^w \mathrm{HDD}_{d(j)t} + \beta_1^f \mathrm{FGDD}_{d(j)t} + \beta_2^f \mathrm{FHDD}_{d(j)t} + \\ &g(P_{d(j)t}) + g_2(FP_{d(j)t}) + \gamma_j + \eta_t + \zeta^1_{r(j)} t + \zeta^2_{r(j)} t^2 + \varepsilon_{jt} \end{aligned} \tag{1}$$

The outcome variable is either costs or revenues for farm j in growing season t. In later results, we also explore effects with different outcomes including farm inputs, planting decisions, and profits. The main right-hand-side variables are realizations of temperature (GDD and HDD) and forecasts of temperature (FGDD

[12]Our results are robust to removing observations with an absolute error value over 500. Results are not shown in the paper, but available upon request.

[13]As discussed in subsection 2.1, seasonal forecasts, and forecasts up to four months ahead, have been available in France since the 1990s. See Canal (2015) for a discussion of the evolution of forecasting methods in France.

and FHDD) experienced by farms located in department d during the growing season. Given that weather and forecasts vary at the department and year level, we use two-way standard errors at the department and year level.

We focus on the marginal effects of forecasts, and do so for two reasons. First, conditional on weather realizations, variation in forecasts should cleanly identify changes in information available to farmers. This information change, in turn, should affect production decisions. In that sense, forecasts are a way to avoid the attenuation bias caused by recovering cost and revenue effects off weather realizations–which, to a minimum, will contain both predicted and unpredicted weather. When looking at outcomes such as revenue, rather than farm decisions, coefficients recovered off weather realizations will further capture the joint effect of farmer action and weather realization. Controlling for realizations, and using forecast variation isolates farmer decisions. Second, if the farmer faces adjustment costs when choosing actions, then they have an incentive to choose actions prior to the arrival of weather. In such a case, forecasts provide more powerful identification of the effect of temperature on farmer actions than looking at realizations of temperature. In a farm setting, adjustment costs are likely high given that many actions need to be taken prior to the growing season (e.g., the choice of which crops to plant, total cropped area) or prior to weather arrival during the growing season (e.g., fertilizer application, defense of crops against freezing). We focus, in the initial results, on one-month-ahead forecasts. In cases with convex adjustment costs, marginal value of information falls as forecast horizon increases. Thus, short-horizon forecasts should again improve power to detect effects. In additional results, we examine forecasts with longer horizons.

The estimating equation also includes controls for the level and square of realized and forecasted precipitation over the growing season to account for effects of precipitation on farm outcomes. We write these as $g(P_{d(j)t})+g_2(FP_{d(j)t})$. We include region specific time trends in the form of $\zeta^1_{r(j)}t$ and $\zeta^2_{r(j)}t^2$. These account for potential sub-national trends that would correlate with weather and our outcomes of interest.[14] Finally, farm fixed effects, γ_j, and year fixed effects, η_t, mean that effects are identified from within-farm variation in weather over time, while accounting for national time series patterns in both weather and agricultural costs or revenues, as well as region-level quadratic trends. The identification assumption is that the remaining error term, ε_{jt} is uncorrelated with the temperature forecast variables. The control set is similar to prior work on the effects of climate on agriculture, with one important difference: we are able to use farm fixed effects rather than geographic area fixed effects (e.g., many studies in the U.S. include county fixed effects). This more granular cross-sectional control should alleviate concerns about confounding farm-level characteristics like geographic features that determine crop suitability and weather patterns.

4 Results: Marginal Costs and Benefits of Adaptation

4.1 Cost and Revenue Effects

We first show the effect of forecasted moderate and extreme temperature on farm costs and revenues, controlling for realizations of temperature. Table 2 displays the results. The outcomes are revenues (measured by total sales) and a measure of the costs of production. The measure of costs is broad and includes the cost for intermediate inputs, social contributions to workers, personnel expenses, taxes, and insurance.[15] This

[14]Recent work has highlighted the endogeneity of agricultural technical change to the heterogeneous exposure of crops to heat, with crops more exposed being a relatively higher focus of innovation. See Moscona and Sastry (2022). We also run regressions without the department-specific trends, and find very similar results. This structure of time trends follows from Schlenker and Roberts (2009), albeit we also include year fixed effects to account for potentially non-linear France-level shocks. Year fixed effects seem particularly relevant in the European context, with an integrated agricultural market and likely spatially correlated heat shocks across countries which will impact overall demand.

[15]We perform a test and run the regressions using only expenses for intermediate inputs, as a check for potential mismeasurement, and find similar results. Results are available upon request. See subsection B.1 for the definition of intermediary inputs.

broad measure is less likely to miss potential adaptation costs.

GDDs, as the name implies, help crops grow and can thus be interpreted as positive productivity shocks. HDDs, in contrast, are temperatures so extreme that they cause crop losses. We expect that forecasts of GDDs are useful for taking advantage of better growing conditions, while forecasts of extreme heat are useful to either cut production costs, to increase input usage to compensate for adverse conditions, to better target input usage, or to modify decisions such as the timing of harvest or one's crop mix composition. In a simple model where farms optimize profit, and local adverse weather (either lower GDD or higher HDD) corresponds to a negative TFP shock—but where output prices are not responsive to local conditions—we would expect forecasts of worse conditions to lead to a reduction in the scale of production, leading in turn to reduced revenues and costs. The expected negative TFP shock shifts the marginal cost curve upwards, and under decreasing returns to scale, induces farms to reduce their scale of production until it matches the market price again. On the other hand, changes in the timing of harvest could allow for a positive revenue response without substantial changes in costs.

Table 2 shows that revenues respond positively to forecasted GDDs and HDDs. In contrast, costs respond positively to forecasted GDDs but exhibit an insignificant response to forecasted HDDs. And, in comparison to the forecasted HDD effect on sales, the response is quantitatively small. The point estimates of the effects of realized GDDs and HDDs are generally in keeping with our assumptions that higher GDDs are productivity improving for firms while higher HDDs are generally productivity reducing. The interpretation of these coefficients, however, is not straightforward. They mix direct effects of realized weather with the effect of ex-post adaptation actions. In our sample, these effects are also not statistically significant. Given that the central goal of the paper is to identify and quantify costs of adaptation, we do not give further attention to these coefficients.

Table 2: Cost and Revenue Reactions to Forecasts (1 month lead)

Dependent Variables:	Revenue	Costs
Model:	(1)	(2)
Variables		
GDD	5.103	1.213
	(11.84)	(7.809)
GDD (F)	37.39	30.79**
	(34.66)	(14.48)
HDD	-158.3	-191.9
	(332.2)	(156.9)
HDD (F)	1,192.2***	-56.71
	(402.7)	(178.5)
Mean	155,386.2	123,249.1
Unique Farms	2,603	2,603
Fixed-effects		
Farm	Yes	Yes
Year	Yes	Yes
Fit statistics		
Observations	18,917	18,917
R^2	0.88715	0.93666

Notes: Estimates are based on Equation 1, using the baseline sample. Realized and forecasted rainfall in levels and squares, as well as quadratic region-specific time trends, are included as controls in addition to the indicated fixed effects. Observations are weighted using the sample weights provided in the FADN. Two-way department-by-year standard-errors in parentheses. Stars indicate estimate is significantly different from zero: $^*p < .10$, $^{**}p < .05$, $^{***}p < .01$.

Looking first at the effect of forecasted GDDs, one can see that the effects on costs and sales are of comparable magnitude. This result suggests that for moderate temperatures, farmers are taking costly actions to arrive at an increase in revenue. This is in line with the hypothesis laid out above that forecasts of better conditions (higher GDDs) lead farmers to expand the scale of their production. The effects on both revenues and costs are small relative to the mean for a one-unit increase in GDDs. The standard deviation of forecasted GDDs is about 240, so a typical change in GDDs will lead to an about 6% change in sales and costs.

The effect of forecasted HDDs shows a different pattern. Sales increase while costs do not show substantial or significant responses. An increase by one degree-day in HDD forecasted one month in advance leads to an increase in sales of €1,192. Given typical sales per farm of about €$150,000$ and a standard deviation of forecasted HDD of just over 1 (see Table A2), this coefficient indicates that a typical change in forecasted HDD causes sales to change by about .8%, on average. In contrast, forecasts of extreme heat have no statistically significant effect on production costs. And from comparing the revenue and cost effects, one can see that the effect on costs is also practically small in magnitude. The results indicate that there are large benefits to adaptation and near-zero costs of adaptation, on average, for these farmers.[16]

[16] Table A7 compares results to one-month and two-months ahead forecasts. Figure A15 compares the results for all available leads.

Finally, we look at a distributed lag model which includes both weather realizations and the entire set of forecasts with a lead from one through five months ahead. That is, we run the following regression on our outcome measures:

$$
\begin{aligned}
y_{jt} =& \beta_1^w \mathrm{GDD}_{d(j)t} + \beta_2^w \mathrm{HDD}_{d(j)t} + \sum_{\ell=1}^{5} \left(\beta_{1,\ell}^f \mathrm{FGDD}_{d(j)t}^{\ell} + \beta_{2,\ell}^f \mathrm{FHDD}_{d(j)t}^{\ell} + g_{2,\ell}(FP_{d(j)t}^{\ell}) \right) + \\
& g(P_{d(j)t}) + \zeta_{d(j)}^1 t + \zeta_{d(j)}^2 t^2 + \gamma_j + \eta_t + \varepsilon_{jt}
\end{aligned}
\tag{2}
$$

where all variables are the same as in Equation 1 except we have added forecasts for each horizon, as indicated by the variables FGDD^{ℓ}, FHDD^{ℓ} and $FP_{d(j)t}^{\ell}$. This regression identifies the precise timing of information arrival. When running regressions with a unique set of forecasts, these will capture the information received at other leads, via their non-controlled-for correlation with these other forecasts. Essentially, not including all the leads available creates a form of omitted variable bias where the included forecast captures a composite of the effects of all forecast horizons, with the composite effect being determined by the autocorrelation of forecasts across horizons. As long as we control for realized weather, this omitted variable will not be an issue for identifying benefits or costs of adaptation, because the interpretation of the forecast coefficients is still that it causes changes in the agent's action. However, it is useful to include all the possible forecasts leads in order to understand which forecast lead is most useful to farmers—in the sense that it generates the largest response.[17] It also puts the forecast and realization effects on similar footing in the sense that both are then identified by shocks: surprising realizations in the case of the realized temperature and news shocks in the case of all forecasts for horizons less than five months ahead.

Figure 1 plots the coefficients associated with heating degree days from estimating two different versions of Equation 2. In particular, the figure shows β_{ℓ}^f for $\ell \in [1, 4]$ and shows β^w for the lead value of 0 for regressions with costs and sales on the left hand side. We exclude the coefficients for lead 5, given that these are less cleanly identified, accounting for all the information received more than four months in advance.[18]

[17] In a model with risk averse farmers, a specific forecast lead might generate a larger response, either because forecasts produced at that lead value are more precise (in general forecasts become more precise as they come closer to the predicted event), or because the timing of the favored adaptation responses matches its lead value the most. Figure A1 shows that forecast errors do not change significantly across lead values; at least when aggregated into our growing season variables. As such, we can expect that here, differences across lead values are mainly driven by questions of timing.

[18] As discussed in Canal (2015), forecasts were only available up to four months in advance in the 1990s in France, and our five-month ahead hindcast is hence less likely to correspond to information available to farmers at the time.

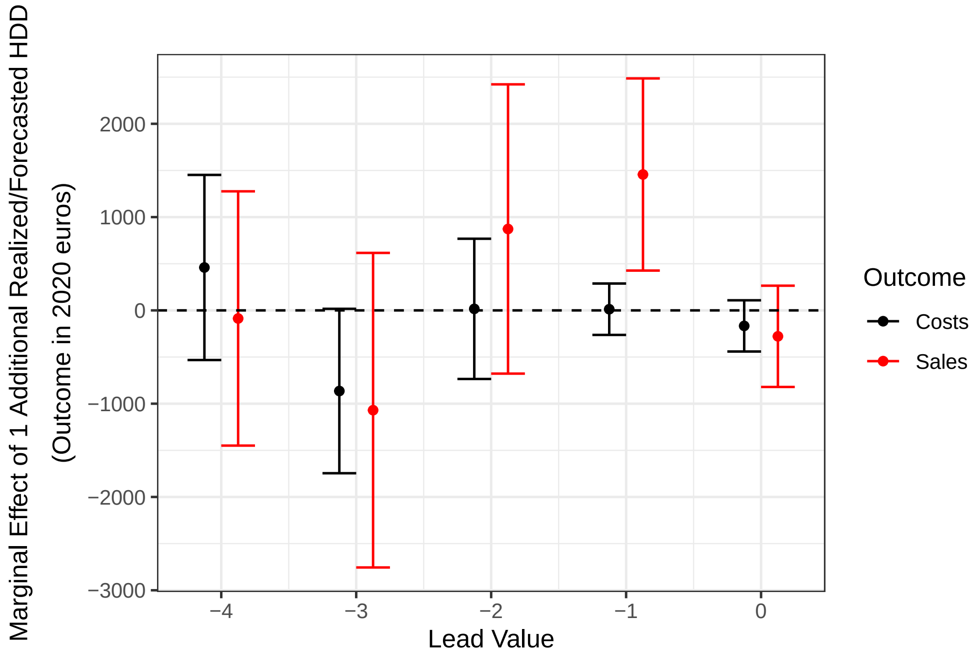

Figure 1: Costs and Revenue Reactions

We see that the the effect on sales associated with a one-month-ahead forecast is statistically significant and of a similar magnitude to the effect found in Table 2. The cost effect of a one-month-ahead forecast of HDDs remains centered around zero. Revenue responses are not statistically significant at the 5% level for horizons longer than one month, but the point estimates indicate that information close to the shock allows for revenue-enhancing adaptation while information farther from the shock leads to a decrease in revenues. Costs are near zero for all horizons aside from three months ahead. The fact that costs do not move more strongly at short horizons than at longer horizons is inconsistent with adjustment costs driving the results that we find (Downey et al., 2023). The significant three-month-ahead forecast could indicate that there are adjustments that are uniquely available at a quarterly frequency that allow farms to reduce the scale of their operation in anticipation of adverse weather conditions.

We also observe that the coefficients associated with realizations become more negative and almost statistically significant at the 5% level. This might mean that including more leads allows for more powerful estimation of the direct effect of extreme heat on agricultural revenues and costs.

The large effect on revenues and small effect on costs naturally raises a question: how are farmers achieving an improvement in revenues with little to no change in costs? Below, we explore farm-level behavior that does and does not respond to forecasts to shed light on this question. We also test for—and rule out—a variety of measurement and identification arguments that could explain the results.

4.2 Ruling Out Explanations Due to Measurement Error, Offsetting Effects, Specification Choices, Identification of News, or Adjustment Costs

Before investigating the mechanisms by which firms might be adapting to temperature, we first rule out some alternative explanations of the results.

A first issue is that input quantities and prices responses might offset each other. On this, we show separate regressions: one for the quantity of inputs and one for their prices.[19] The table corresponding to the figure below is in Table A10. We use input prices as recorded in an agricultural input price survey used to build nationally representative input-specific Laspeyres indices. The precision of this survey is useful, in that we can run the regressions including store and product fixed effects in addition to year fixed effects. We show the results pooling together all categories of input prices here, and outline the price responses per input category in Table A15.[20] We also show results for three of the main inputs, and the ones we consider the easiest to adapt to one-month ahead information shocks: irrigation, fertilizers and pesticides. We show in Table A13 results for a larger set of inputs, and in Figure A16 and Figure A17.

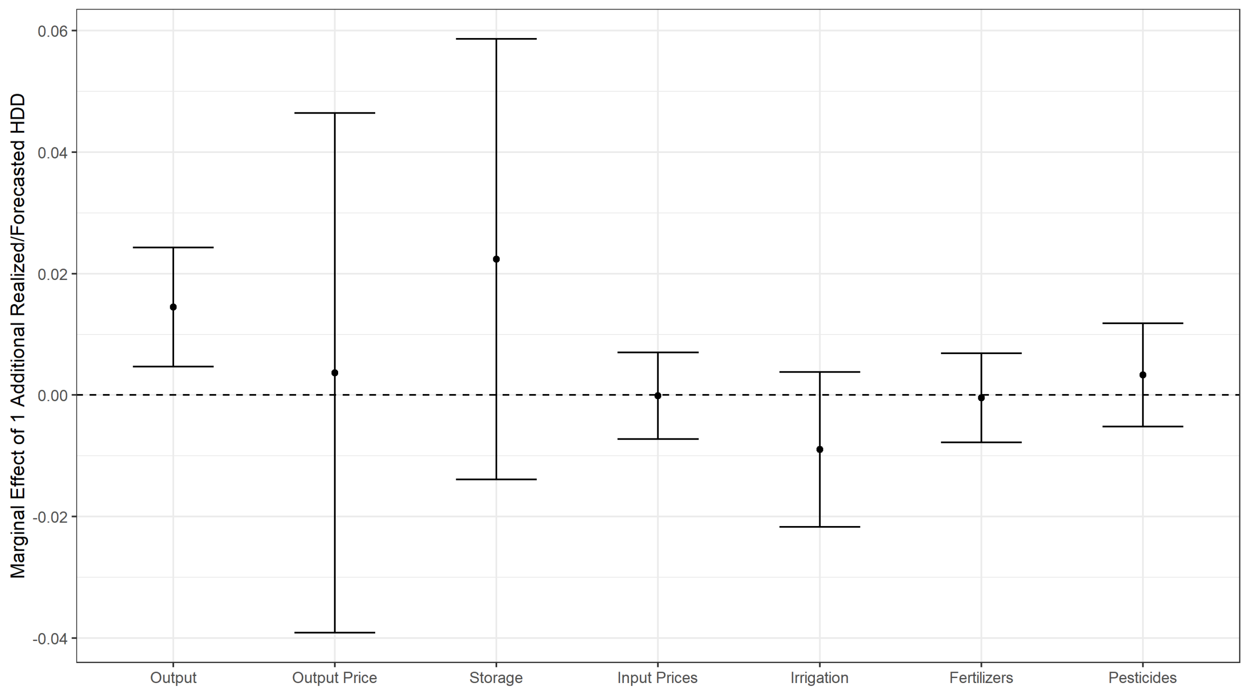

Figure 2: Prices and Quantities for Outputs and Inputs

We further decompose land and labor responses per sub-categories of land ownership and labor contracts.[21] Store prices are unresponsive to HDD forecasts, as are inputs. In Table A13, we see that none of land, labor, fertilizers, phytosanitary products, seeds and irrigation respond to forecasted HDDs. Labor does respond to forecasted GDDs, matching the increase in total costs of production see in Table 2. Decomposing land responses in Figure A16, we see that this does not seem to be driven by an aggregation of different land categories with heterogeneous adjustment costs. Both the response of rented land with short rental contracts, and of owned and used land are non-statistically significant. Additionally, fallow land does not respond as well, whereas we could have considered it as a potential source of extensive margin adaptation. When decomposing labor responses in Figure A17, we note that the two most flexible types of labor respond

[19] Fertilizers and phytosanitary products are measured as deflated input bills, using France-level input-specific Laspeyres price indices which translate the deflated input bills into 2020 euros. See subsection B.1 for detailed definitions of the different variables. Because we do not observe price indices at the farm-level, we are unable to account for variation across farms in input qualities, and variations in input qualities potentially induced by weather shocks. Input upgrading in response to weather shocks is to our knowledge a question which has never been studied, although research in firm dynamics has shown that firms can respond to local shocks by changing the quality of inputs they use. See Verhoogen (2023) for a discussion.

[20] Table A16 gives the two-months ahead results.

[21] Table A14 shows similar results for two-month ahead forecasts.

negatively to expected hot weather, with non-regular salaried workers dropping in a statistically significant way by about 2.83 hours per additional forecasted HDD. This response remains small, given that labor is measured in total hours of work per year, and that the standard deviation in forecasted HDD is of one.

Overall, these results confirm that the non-responsiveness of farm-level total costs to forecasted HDDs does not stem from countervailing price and quantity effects, but that neither respond strongly to forecasted heat shocks. Here the only responsive input is the labor of non-regular farm workers.

A second issue is that the total costs variable could be subject to measurement error, leading to attenuation of the effect of forecasted HDDs. Total costs are measured as an aggregated input bill at the farm level, and follow directly from the precise accounting exercise done to collect FADN data. Nonetheless, measurement error could be possible for total costs. We expect that the two inputs which are the most likely to be mis-measured are labor and land. In order to circumvent this issue, we show the response of total labor and land in volume (resp. in hours, and hectares) in Table A13. Both are not responsive to heat shock. We note that our measure of labor includes family labor and unpaid labor, and as such should account for the potential opportunity cost of unpaid labor that could be allocated to non-farm work. For land, we use a measure of total utilized agricultural area, which is an expansive measure and accounts for both land that is owned, and land that is rented. We also note that input bills, more narrowly defined, are also less likely to be mis-measured than total costs. The results shown above hence also act as test for the possibility of measurement errors driving our zero cost response.

Next, we focus on similar regressions for the quantity of output (simply the sum of all produced quantities), storage (sum of all stored quantities), and an output price index, also outlined in Table A10.[22][23] We first note that while aggregate output does not respond positively to forecasted GDDs, the crop-specific decompositions performed in Table A17 show both wheat, corn and colza benefit from them, even though the coefficients are noisy. As expected, realized HDDs hurt production, here a one standard deviation increase in realized HDDs (3.202) leads to a decrease in output by 1.5ppt. In comparison, and using department-level time series of yields, Gammans et al. (2017) find that an additional one-day exposure to temperature above 32°C will decrease wheat production by about 2.5ppt. The comparison between these two results is not directly straightforward. Our heating degree day cut-off is at 30°C, in order to account for the impacts such temperature can already have on corn. We also account for weather realizations and forecasts over the entire possible growing season, allowing for farmers to both switch their crop mix composition (and hence the relative density of their lands exposed to weather or winter growing seasons), and switch their growing season at the crop level (slightly moving up or down their wheat or corn specific growing seasons for example). We also note that they focus on weather accrued during the warm part of the wheat growing season–although warm weather accrued over the total growing season, and the warm months should be comparable quantities in the French context. As such, it makes sense that they find a larger effect for unavoidable hotter weather, while we find that farmers face lower consequences form their endowed weather which might not necessarily be the effective weather they end up facing. Forecasted HDDs also significantly impact output quantity. A one unit increase in forecasted HDD raises output by .9% of total output. These results provide reassurance that the revenue and cost effects we identify are indeed associated with adaptive action rather than a change in input or output prices.

We also investigate effects on storage. The positive sales response might be due to an increase in volumes sold coming either from increases in production, in production sold (a decrease in current period storage), the use of previously stored output, or the increase in output prices holding volumes sold constant. We find that storage does not respond in a significant way to forecasted HDDs.

[22] We show in Table A17 the crop-specific responses of output quantity to temperature realizations and forecasts.

[23] The output price index used in this regression corresponds to a weighted average of crop-farm level output prices as observed in the data, using relative land shares as weights. subsection B.1 gives a formal definition of the different variables.

Finally, we examine the robustness of the results to changes in the specification of the estimating equation including the addition of lags of realized weather or removing department-specific time trends. We vary the cutoffs used to define a GDD versus an HDD and add a measure of freezing degree days. We also use a disaggregated measure of realized weather, measured at the village level, combined with the same department-level measure of forecasts. In this case, realized weather is also computed using the entire range of hourly realizations, and not only four observations per day. And we investigate heterogeneity in the response to forecasts across farms. In all cases, we find results that are in line with the baseline estimates. These robustness checks are reported in Appendix D. We further show results for different sections of the data ranked by their average HDD realizations in Figure A19. When looking at profit results below, we investigate the variation of our results along the sign and magnitude of the forecast error.[24]

4.3 What Is Driving the Results? Differences in Weather Variation and Actions

So what is causing the similar response of revenues and costs to GDDs, on the one hand, and the substantially different response to HDDs on the other hand? This section describes three parts of the answer. First, we show that the variation in temperatures underlying GDDs and HDDs differs. GDDs are much more likely to be associated with small or marginal changes in temperature. HDDs, in contrast, are much more common when there are large, non-marginal changes in temperature. These different changes in temperature, in turn, drive different behaviors by farmers. In response to GDDs, farmers take more continuous actions, while they take action in response to HDDs that are lumpier or which entail opportunity costs rather than changes in costly inputs.

To demonstrate these effects, we first decompose the variation in temperature in our data into marginal and non-marginal variation for our two temperature variables (GDDs and HDDs). Our decomposition is done according to the following procedure: first, we compute a thirty-year average of hourly temperature realizations at the department level.[25] We then compute the average temperature value for each hour-day-month tuple. Using those same data, we also build the temperature standard deviation for each hour-day-month tuple. For each hourly, department-level temperature realization in 1994-2018, we then check whether that realization deviates by more than 1.96 standard deviation from the mean. If it does, we classify the realization as non-marginal, if it does not, we classify it as marginal.

In our sample, 85% of the HDDs are classified as non-marginal according to this definition, while only 13% of GDDs are classified as non-marginal.[26]

We next run our main specification but on realized and forecasted temperature variables decomposed into marginal and non-marginal (denoted in the table by "m" and "nm" respectively). The results are shown in Table A18. The main take-away from the results is that non-marginal HDD forecasts induce a positive sales response, while the rarer, marginal HDD forecasts actually lead to a reduction sales—in line with the hypothesis that higher HDDs should cause a reduction in scale for farmers. This is also in line with the adaptation strategies we outline below, which correspond to changes in the timing of the growing season, and in crop mix composition. Both are unlikely to respond to marginal changes in hot weather, and should only be undertaken as very noticeable shocks are expected. Hence, marginal heat shocks for which our main adaptation strategies are less likely, should come with negative revenue effects. We note, however, that the support for marginal HDD is small which explains that on average the effect of an additional forecasted HDD remains strongly positive and significant.

Second, we note that costs also respond to non-marginal GDD forecasts, while they do not respond in a significant way to marginal GDDs. This result is likely well identified given the significant marginal component of GDDs in our data. We see however that the revenue and cost responses to marginal GDDs overlap,

[24] The annex also contains these same results for cost and revenues as outcomes.

[25] That is, we aggregate the gridded ERA5 data for 1963-1993 at the department level using area-weights.

[26] See Table A20 for details.

and that the responses to marginal and non-marginal forecasted GDDs cannot be strongly distinguished, statistically, in the data.

The difference in variation underlying forecasted GDDs and HDDs suggests that farmers might also be taking different types of actions in response to forecasts of these different measures. We have previously highlighted in Table A10 how forecasted GDDs induce a positive input response in the form of higher input usage, while forecasted HDDs show null results across the input schedule. We now also utilize data on dates of ploughing, initial irrigation, harvest, and sowing. These data are gathered at the field level from the survey of agricultural practices (PKGC). The survey is a repeated cross-section, and as such we are not able to include farm fixed effects. We only include department fixed effects (the smallest geographic identity to which we can attach the fields, roughly corresponding to US counties). Given the coarser sampling procedure used to build this data, we cluster the standard errors at the region and year levels. Otherwise, the specification is the same as the baseline regression in Equation 1.

Table 3: Timing Response to Forecasts (Wheat - 1 month lead)

Dependent Variables:	Ploughing	First Irrigation	Sowing	Harvest
Model:	(1)	(2)	(3)	(4)
Variables				
GDD	0.0203	0.0787	0.0213	4.7×10^{-5}
	(0.0244)	(0.1107)	(0.0124)	(0.0028)
GDD (F)	0.0702	1.566***	0.1377	0.0094
	(0.0781)	(0.1326)	(0.0486)	(0.0180)
HDD	0.9084	-3.849***	0.4733	-0.1368
	(0.3785)	(0.2796)	(0.3078)	(0.0469)
HDD (F)	-1.171*	2.121*	-0.9083	0.3606
	(0.2762)	(0.6573)	(0.5233)	(0.2112)
Mean	337.5	532.1	378.6	594.3
Fixed-effects				
Department	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Fit statistics				
Observations	7,313	214	9,987	9,974
R^2	0.37425	0.54312	0.23870	0.52929

Notes: Two-way region-by-year standard-errors in parentheses. Stars indicate estimate is significantly different from zero: *$p < .10$, **$p < .05$, ***$p < .01$. Realized and forecasted rainfall in levels and squares are included as controls, as well as quadratic region-specific time trends.

We focus specifically on examining timing of production decisions. The timing of different procedures is likely to react to forecasted weather, and could have large effects on the output for a farm. Indeed, we can expect the timing of harvests to be affected by forecasts of extreme heat which could hurt the crops. On the contrary, sowing might be done earlier if one expects freeze events. Timing is also likely to have a small impact on production costs, while allowing for a higher efficiency of input usage. An example would be that forecasts can allow for a more targeted use of irrigation, fertilizers or pesticides – while keeping their related expenses constant.

We observe that the date of the first irrigation and the date of ploughing are both reactive to forecasted HDDs. Forecasted HDDs bring ploughing dates earlier, as well as sowing, although the coefficient is not significant there. Irrigation on the contrary happens later when one expects adversarial weather, perhaps in order to save water which is expected to be scarce, or to prevent too much evaporation of the irrigated water. These adaptive behaviors are likely to impact production volumes, and hence play a role in the observed, positive revenue response. We also see that realized heat shocks bring irrigation dates earlier, perhaps forcing the farmers to irrigate earlier in order to avoid crop loss.

If we use these implied responses of growing season outlines to reconstruct our weather data, we can compare the extent to which the weather faced by farms differs from the weather they are endowed with– as a result of forecast-induced adaptation. We do so in a simple way. We first compute the distribution of hours spent at each one degree temperature bin over the October-July period of each growing season, matching our favored growing season definition. We then vary these start and end dates according to the department-year specific HDD forecasts and their marginal effects on growing season timing implied by the results in Table 3. We compute a second distribution of temperature realizations using these endogenous outlines. We finally compare the evolution of the cumulative distribution functions between the endogenous season, and the endowed (fixed boundaries) season over 1994-2018. The graph below shows this variation in mass over the full range of temperatures. We see that while the implied differences in temperatures faced are small, the direction of the adaptation follows what one should expect. As farmers received information shocks about incoming hot weather, they shift the timing of their wheat season in order to increase the mass of realizations for lower temperatures, and to decrease the mass over the range of hot temperatures.

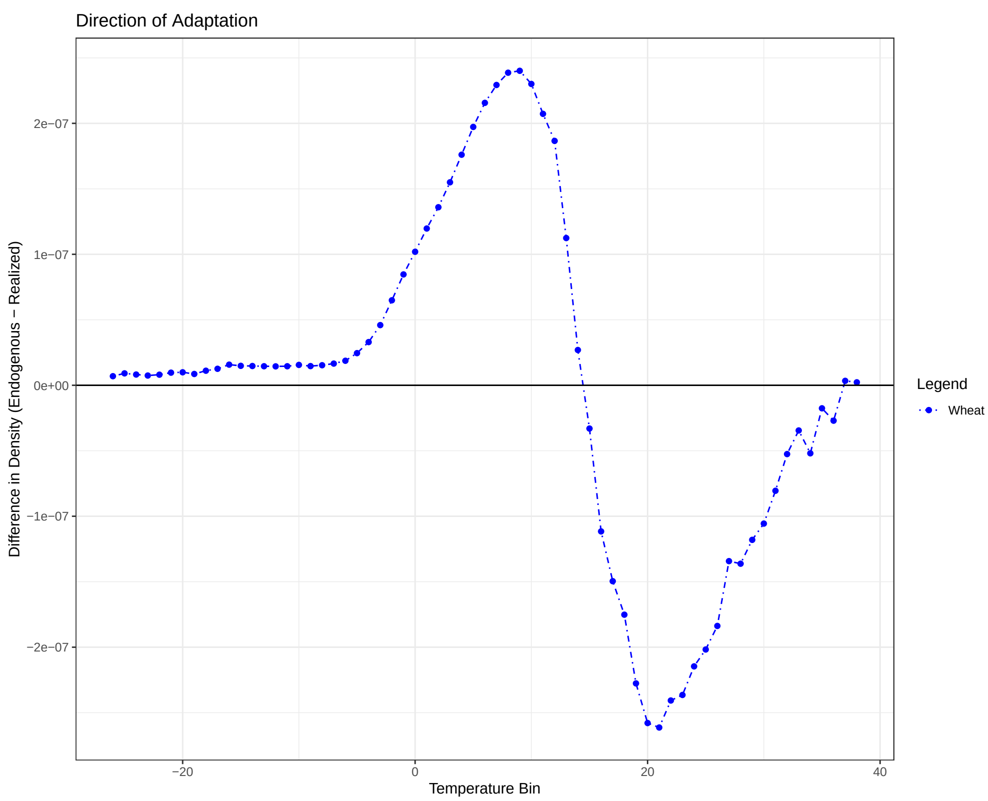

Figure 3: Difference in Temperature Distributions

Overall, these timing responses highlight potential avenues for no-cost adaptation responses to forecasted HDDs, which could explain the positive response of sales for a null response of costs.

We finally explore crop switching as a potential adaptive strategy. Crop switching might be useful for at least two reasons. The first one is that one can switch towards heat resilient crops when expecting hotter weather, hence reducing the losses faced when actually encountering bad weather. Corn and wheat, and colza and sunflower have different heat tolerance – and moving from corn to wheat and sunflower to colza might reduce one's sensitivity to heat shocks. The second is that switching crops can also allow to switch the average weather faced by farmers. We show in Figure A9 and Figure A10 that typically, wheat is sown around November, and harvested before August in France, while corn is sown around May and harvested around October. While fields used for winter wheat are unlikely to be usable for corn later in the year, using them for wheat when expecting a hot summer can help avoid facing August heat shocks, and reduce the effective amount of HDDs faced. For a given amount of HDD endowed over one's potential total growing season, forecasts can then help reduced the amount of actual HDDs faced, and as a consequence increase yields. The implied changes in the weather a farm ends up facing might also be much larger than the ones induced by crop-specific changes in growing seasons that we described above.

Table 4: Farm-Level Crop Diversity (2 months forecasts)

Dependent Variables:	Crop Count
Model:	(1)
Variables	
GDD	0.0001
	(0.0002)
GDD (F2)	8.04×10^{-5}
	(0.0006)
HDD	-0.0060
	(0.0052)
HDD (F2)	-0.0258**
	(0.0121)
Mean	3.629
Unique Farms	2,603
Fixed-effects	
Farm	Yes
Year	Yes
Fit statistics	
Observations	18,917
R^2	0.77734

Notes: Two-way department-by-year standard-errors in parentheses. Stars indicate estimate is significantly different from zero: $*p < .10$, $**p < .05$, $***p < .01$. Realized and forecasted rainfall in levels and squares are included as controls, as well as quadratic region-specific time trends.

To study the farms' crop mix response to expected weather shocks, we first simply look at the evolution of the number of crops they are growing in a given year. Here, this is simply measured as the number of crops for which they have a non-zero land use. The table shows that farms reduce their number of crops in response to expected heat. We use here two month ahead forecasts, where we find a stronger farm response, and we expect that crop mix decisions are likely based on longer forecast horizons than other production decisions.

We conclude this section by showing crop-specific land movements, which are suggestive of the potential switching patterns operated by the farms. Each coefficient corresponds to the coefficient attached to one-month ahead forecasted HDD in a regression similar to our main specification. We note that sunflower and corn see their land surface decrease in response to expected hot weather, while wheat, colza, barley ans peas/beans see their surface increase. We can expect colza and sunflower, as well as wheat and corn to be closer substitutes on the demand side, while having different heat profiles as highlighted by our Schlenker and Roberts (2009) specifications in Figure A11. Farmers might then be switching away from sunflower and towards colza, and away from corn and towards wheat in response to heat forecasts.

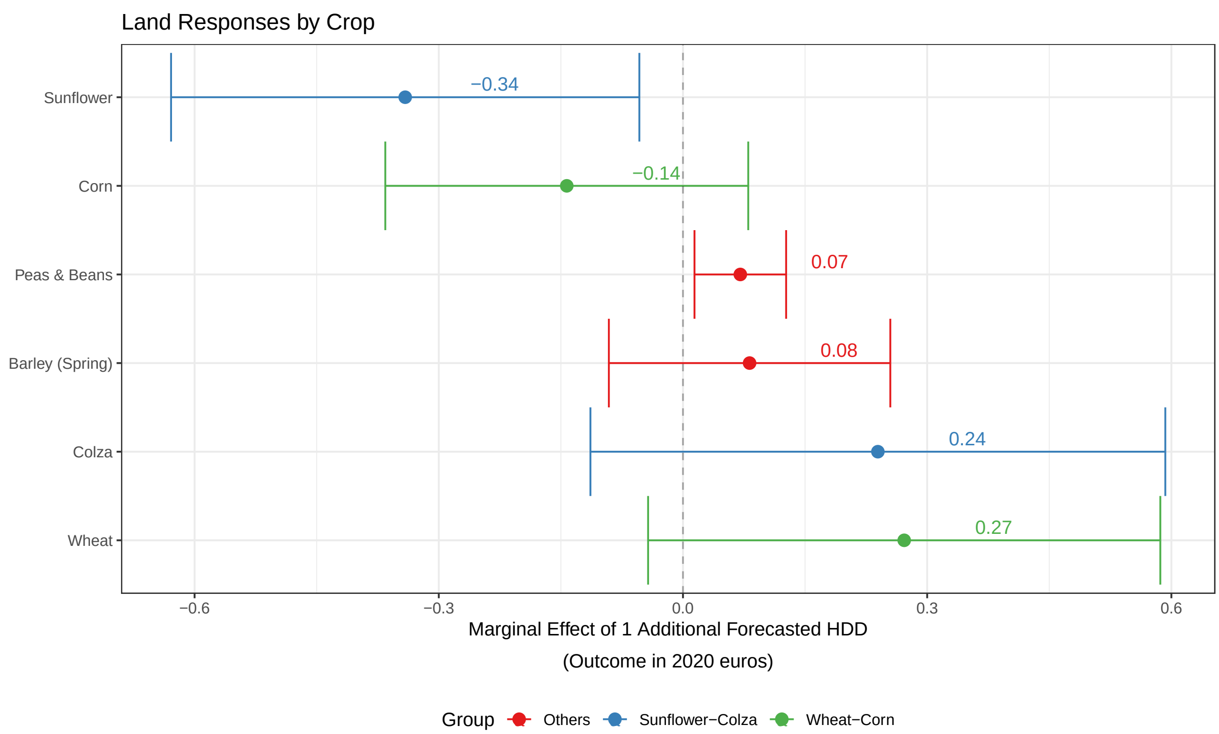

Figure 4: Decomposition of Cropland Responses

Notes: We show the results of our main specification, showing the reaction of different land area margins. Each regression contains realized and forecasts weather outcomes (rainfall and temperature), as well as region-specific quadratic time trends, farm and year fixed effects. Standard errors are clustered at the department-by-year level.

5 What are the Consequences of Adaptation?

So far, we have seen that farm revenues and costs respond to forecasts of moderate temperatures but that costs do not respond to extreme temperature forecasts while revenues increase. These effects should imply different profit effects from forecasts of moderate versus extreme temperature forecasts. Namely, moderate temperature forecasts should have little effect on profits while extreme temperature forecasts should have a large effect on profits. Even though profit effects follow from the effects on revenues and costs, they are still useful for two reasons. First, they provide a convenient, single variable summary of the net benefits of adaptation. Second, they provide a simple yet powerful test of the envelope theorem. Below, we first lay out the logic for why the regression of profits on forecasts acts as a good test of the envelope theorem before showing the results. We then use the results presented so far in the paper to offer suggestions for tests that can be done to support the use of envelope theorem-based estimators. We also look at the timing of effects on profits to examine the marginal net benefits of forecasts of different horizons, and outline the variation in profit responses along the range of HDD forecast shocks.

We then move to looking at results between periods, that is whether or not ex-ante adaptation has dynamic implications for farm profits which differ from the within-period ones. We show that farms face significant increases in production costs in the periods following the forecast, and that the overall effects of adaptation are negative. We further detail the source of these cost increases, and discuss their implications for the costs associated with crop switching.

5.1 Within-Period Effects: Static Considerations

5.1.1 Effect of Forecasts on Profits

A standard model employed in the climate economics literature considers a price-taking, profit maximizing agricultural firm, facing competitive input and output markets. The firm is further assumed to face weather shocks affecting its productivity. Denote weather-affected productivity in period t as w_t. We assume that the firm builds an expectation of w_t, which is used to make optimal production decisions. We denote this expectation ω_t, and assume that it depends on the past realization w_{t-1}, on a publicly observable forecast f_t, and on an unobserved signal v_t. Focusing on static decisions, we can then write the period t profit maximization problem, where input level x_t has to be set to solve:

$$\max_{x_t} \mathbb{E}_t\Big[\pi_t \Big| w_{t-1}, f_t, v_t\Big] \Leftrightarrow \max_{x_t} \mathbb{E}_t\Big[p_t q(x_t, w_t) - c(x_t)\Big| w_{t-1}, f_t, v_t\Big].$$

Here π_t is farm profit, p_t output price, $q(.)$ the production function which depends on weather and input choices, and finally $c(.)$ is the cost function. Assuming the problem is well behaved, the first order condition will equate expected marginal revenues and expected marginal costs.[27] That is,

$$\mathbb{E}_t\left[p_t \frac{\partial q}{\partial x}\middle| w_{t-1}, f_t, v_t\right] = \mathbb{E}_t\left[\frac{\partial c}{\partial x}\middle| w_{t-1}, f_t, v_t\right] \tag{3}$$

at optimum.

This implies a profit-maximizing choice of actions $x_t^\star = x_t^\star(\mathbb{E}_t\big[w_t\big|w_{t-1}, f_t, v_t\big])$, from which we can recover the indirect profit function $\pi(x_t^\star, w_t)$, and the static value function:

$$V(w_t) = \mathbb{E}_t\Big[\pi(x_t^\star, w_t)\Big| w_{t-1}, f_t, v_t\Big].$$

A widely invoked argument in climate econometrics states that one can find the direct effect of weather—the effect purged of the influence of adaptation behavior—on a firm by estimating the marginal effect of weather on optimal expected profits. This argument follows from the envelope theorem (see, e.g.,

[27] While we assume that output prices are not responsive to weather, or to sold output, the derivations can be extended to account for these.

Deschênes and Greenstone, 2007; Hsiang, 2016):

$$\begin{aligned}\frac{\partial V(w_t)}{\partial w_t} &= \mathbb{E}_t\left[p_t \frac{\partial q}{\partial w_t}|_{x=x^\star} + \frac{\partial x_t^\star}{\partial w_t}\left(p_t \frac{\partial q}{\partial x}|_{x=x^\star} - \frac{\partial c}{\partial x}|_{x=x^\star}\right)\middle| w_{t-1}, f_t, v_t\right] \\ &= \mathbb{E}_t\left[p_t \frac{\partial q}{\partial w_t}|_{x=x^\star}\middle| w_{t-1}, f_t, v_t\right].\end{aligned}$$

In other words, empirical estimates of $\frac{\partial V(w_t)}{\partial w_t}$, the variation in value induced by weather, is a sufficient statistic for $p_t \frac{\partial q}{\partial w_t}|_{x=x^\star}$, the marginal impact of weather on the firm absent any influence from the marginal benefits or costs of adaptation behavior. Finally, observing the variation in $\frac{\partial V(w_t)}{\partial w_t}$ across different climates, either in a cross-sectional or in a long-difference fashion, can be informative of how adaptation to local climate can impact the sensitivity to realized heat. One can then backtrack the costs of adaptation, from these estimates of benefits of adaptation, as they are to be equivalent.

We argue that observing $\frac{\partial V(w_t)}{\partial f_t}$ is a useful statistic to test whether the envelope theorem is empirically applicable. In particular, a necessary condition for the envelope condition to hold can be tested using the null hypothesis $\frac{\partial V(w_t)}{\partial f_t} = 0$.[28] To see why this is a test of the envelope condition, note that because forecasts are purely informational, we have:

$$\frac{\partial V(w_t)}{\partial f_t} = \mathbb{E}_t\left[\frac{\partial x_t^\star}{\partial f_t}\left(p_t \frac{\partial q}{\partial x}|_{x=x^\star} - \frac{\partial c}{\partial x}|_{x=x^\star}\right)\middle| w_{t-1}, f_t, v_t\right]. \tag{4}$$

If forecasts are a meaningful signal for agents, their variation will shift firm beliefs about expected weather, conditional on their private signal v_t and past realized weather w_{t-1}. If forecasts are meaningful, and if agents are both able to react to them and should react to them, we have that:

$$\mathbb{E}_t\left[\frac{\partial x_t^\star}{\partial f_t}\middle| w_{t-1}, f_t, v_t\right] \neq 0.$$

Thus, assuming actions respond to forecasts, we can conclude,

$$\frac{\partial V(w_t)}{\partial f_t} \neq 0 \Longrightarrow \mathbb{E}_t\left[\frac{\partial x_t^\star}{\partial w_t}(p_t \frac{\partial q}{\partial x}|_{x=x^\star} - \frac{\partial c}{\partial x}|_{x=x^\star})\middle| w_{t-1}, f_t, v_t\right] \neq 0.$$

We propose Equation (4) as a test for the applicability of Equation (3) in an empirical context. Under the condition that forecasts are a meaningful signal to agents (conditional on the information they already possess), then the static envelope theorem will not be applicable if the marginal adjustments to optimal profit with respect to variation in forecasts is non-zero.

In Table 5, we verify that this is indeed the case. The table shows the effects on the gross operating income of the farm. This measure is expansive, and for example encompasses subsidies received by farms. The coefficients of interest are the ones corresponding to the one month ahead GDD and HDD forecasts. All the regressions include the same controls as in the revenue and cost results (farm and year fixed effects, region-specific quadratic time trends, and the level and square of realized and forecasted rainfall). The table shows that HDD forecasts have a significant positive effect on farm profit.[29] A one unit increase in forecasted HDD increases farm profits by €$2,067$, or 2.3% of the mean. On the contrary, the effect of forecasted GDDs on profits is not distinguishable for zero. This implies that in our context, only adaptive behaviors that respond to forecasted changes in extreme weather conditions have net positive profit effects, while those

[28] While this is a necessary condition of the envelope theorem to hold, it is not a sufficient one as it only captures the size of the confounding effect of ex ante adaptation.

[29] Table A12 gives the profit results for one-month and two-months ahead leads.

responding to forecasted growing degree days have a net-zero effect. We give in Table A8 the same set of results with an additional measure of profits as value added, which confirm the signs and levels of these results.

Table 5: Profit Responses to Forecasts (1 month lead)

Dependent Variables:	Profit
Model:	(1)
Variables	
GDD	-8.322
	(13.16)
GDD (F)	3.257
	(43.14)
HDD	184.6
	(277.2)
HDD (F)	2,066.7***
	(708.1)
Mean	86,695.2
Unique Farms	2,603
Fixed-effects	
Farm	Yes
Year	Yes
Fit statistics	
Observations	18,917
R^2	0.84140

Notes: Two-way department-by-year standard-errors in parentheses. Stars indicate estimate is significantly different from zero: *$p < .10$, **$p < .05$, ***$p < .01$. Realized and forecasted rainfall in levels and squares, as well as quadratic region-specific time trends, are included as controls. The input price variable corresponds to prices observed at the store level in an agricultural input price survey run across France in order to build input price indices.

The implications for the envelope theorem are that we cannot expect the indirect—adaptive—effects of changes in extreme weather to be netted out in an optimization context. Instead, it appears that farmers are re-optimizing production decisions to respond to changes in extreme weather, resulting in non-zero profit effects. As such, this effect will be confounded with the direct effect of variations in extreme weather on production. Regressions of profit on realized extreme weather are thus likely to account for these joint channels.

5.1.2 Checking Assumptions Underlying Application of the Envelope Theorem

The envelope theorem relies on some form of differentiability of the value function, and it only applies to marginal variation in the parameters of the optimization problem. As such, two reasons why it might not apply are that: the context studied is not sufficiently well approximated by a differentiable model, or that the empirical variation in weather used for estimation is non-marginal.

Evidence relevant to the first point has already been presented in subsection 4.3. The main actions that respond to forecasts are related to the timing of ploughing, sowing and irrigation, as well as changes in crop mix composition. These actions have opportunity costs if they are not done at the right time, but they do not necessarily impose substantial additional costs on firms.[30]

In the second case, we note that hourly weather realizations are usually aggregated over time to yearly or monthly measures, the most frequent aggregates being: mean temperature, mean daily maximum temperature, growing and heating degree days, counting hours spent in specific temperature bins, or splines (Cui et al., 2023). These transformations help reduce the dimensionality of high-frequency weather data when matched to lower-frequency outcome data. In the context of agriculture, this is necessary given that yields, output, and profits are typically observed at the agricultural season level, while weather is observed daily or hourly. This aggregation requires at least two things. First that the impact of temperature is temporally separable. Here, bins, splines and GDDs/HDDs offer more flexibility than averages, but all rely on some form of homogeneity—for example that a one degree increase from $9^{\circ}C$ to $10^{\circ}C$ at 12PM is similar to an identical increase at 12AM on the same day, or that that same increase can be treated in the same way at 12AM on January 4th and March 4th of the same year.[31]

The second is that marginal and non-marginal variation in weather impact outcomes in the same way. Here we note that binning the data significantly relaxes the bite of that assumption. For GDDs (and in the same way HDDs), an hourly realization that shifts up by $4^{\circ}C$ will be treated in the same way as four hourly realizations shifting up by $1^{\circ}C$, as long as both of these shifts do not push them past the bounds used to compute the GDDs. However, we might assume agents react differently to forecasts of these respective $4^{\circ}C$ and $1^{\circ}C$ shifts, and that the first one induces a non-marginal response to optimal production decisions pushing us beyond the bounds of applicability of the envelope theorem.

Decomposing weather as in subsection 4.3 might be a useful robustness test for two-way fixed effect studies of climate change impacts. Significant differences in the effects of marginal and non-marginal temperature shocks might be indicative that the marginality assumption underlying the interpretation of the results is not warranted.

5.1.3 Timing of Effects on Profits

We next show the marginal response of profits to different HDD leads. This exercise serves two purposes. First, adding longer-horizon forecasts to the regression helps isolate news shocks. Conditional on a two-month-ahead forecast, for example, variation in the one-month-ahead forecast identifies the effect of news arriving in the intervening month. Second, further-ahead leads might allow for more adaptation if farms are faced with adjustment costs. By estimating the effect of forecasts over multiple horizons, we can trace out the marginal value of information at different forecasting horizons.

The estimates follow the distributed lag specification from Equation 2. As such, the regressions used to produce this graph are identical to the ones shown in Table 5, but for the joint inclusion of weather forecasts from lead one to five-month ahead. This means that we include realized rainfall in level and squares as controls, as well as department and year fixed effects and quadratic region-specific time trends. We cluster the standard errors at the department and year levels to account for spatial correlation in weather and agricultural patterns.

[30] There could be scarcity pricing for inputs like labor if, for example, sowing decisions become more correlated within a year. That effect would be included in our cost measure.

[31] Gammans et al. (2017) discuss the plausibility of that assumption for France. Feng et al. (2015) incidentally discuss the particular role of heat shocks during corn flowering, as potential counter-argument to this homogeneity assumption.

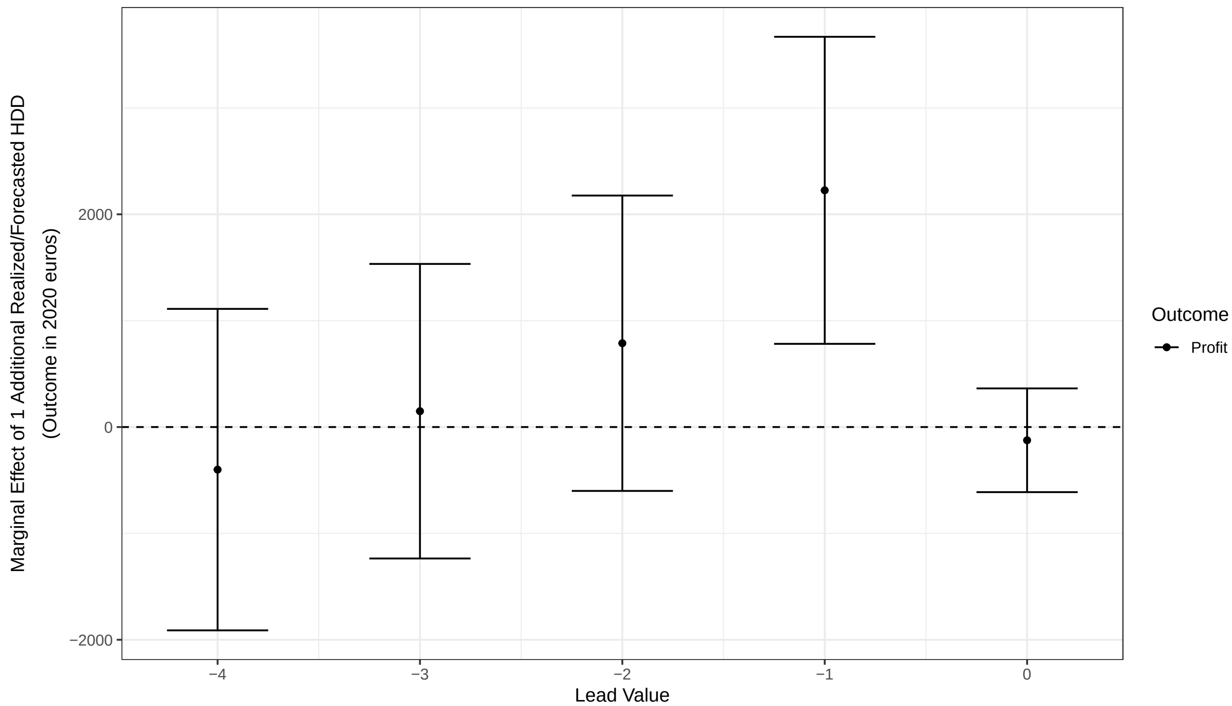

Figure 5: Profit Reactions

Figure 5 shows the results of estimation. The forecast lead is given on the x-axis with a negative sign, in order for the graph to have a chronological interpretation, a value of -1 indicates weather information released one month ahead of the event, and so on. The coefficients associated with HDD realizations are shown for a lead of 0. Five-month ahead forecasts are included in the regression but not shown in the figure because they act as controls for isolating news shocks for the four-month-ahead forecasts. Whiskers show 95% confidence intervals.

Realized hot weather has a negative effect on profits, albeit not significant at the 5% threshold. One-month ahead forecasts still have a large and significant positive impact on profits, while further leads have a decreasing value for farmers, up to four-month ahead forecasts which have an imprecise and slightly negative effect. Overall, this figure highlights that farmers mainly use one-month ahead forecasts to adapt their production schedules. The point estimates suggest that there is rising total value from forecasts up to three-months-ahead, though the marginal value is monotonically declining in forecast horizon.

5.1.4 Targeting Expected Heat

We conclude this section by decomposition the profit impacts of forecasted heating degree days, conditional on the sign and magnitude of the forecast sign error – whether forecasts under- or over-estimate the realized heat shocks of the season.

We run a specification similar in all points to the one used for Table 5, but for the decomposition of the forecasted HDD variable—forecasted HDDs located in distance bins from the realized shock. We take a bin of $[-.1, .1]$ around the realized HDD shock, and use it as our omitted category in the regression. We then identify the median negative error (realization is larger than forecast), and median positive error, and construct bins (as denominated on the figure): $bin = 1$ for the below median negative errors (below the median value of negative errors), $bin = 2$ for above median negative errors (errors closer to the realization),

$bin = 3$ for our omitted category, and $\{4, 5\}$ work similarly for positive forecast errors. These four dummies replace our $HDD(forecast)$ variable in our main regression, which remains otherwise identical. In the graph below we show the coefficient for these four variables, which then indicate the marginal profit effects of one extra forecasted HDD in the specific bin. Complete cost, revenue and profit responses are shown in Table A30.

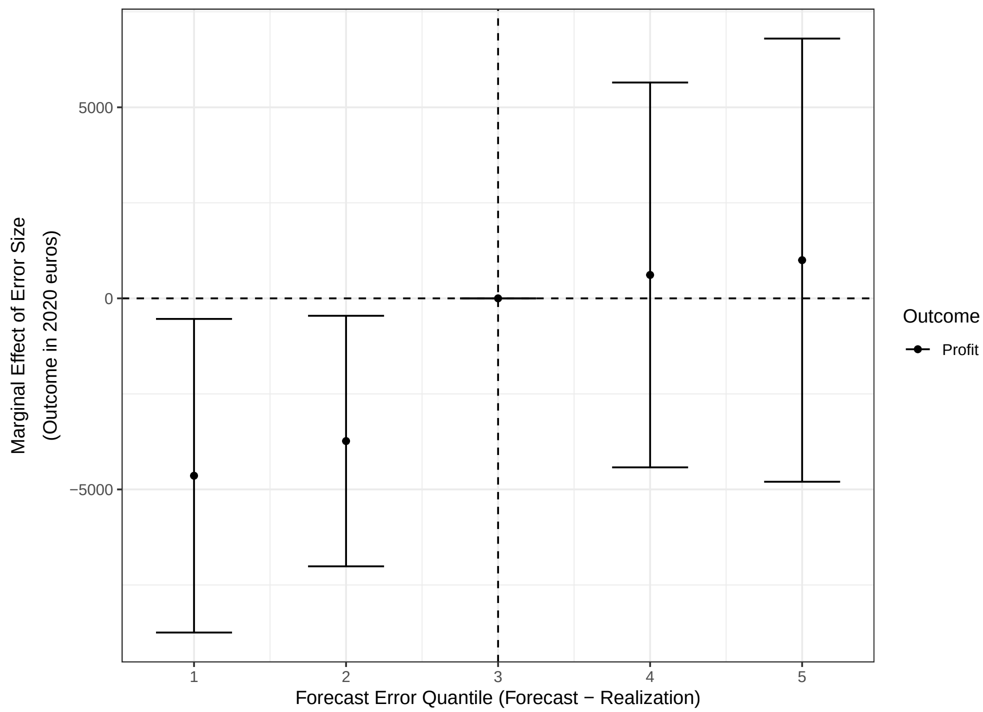

Figure 6: Profit & Sign of the Forecast Error

We see an interesting an intuitive pattern: negative forecast errors are costly to farmers relative to on-point forecasts – they likely lead to farmers under-forecasting future heat shocks and suffer from the full un-adapted impact of them. On the contrary, over-forecasts lead to no statistically significant gains in profit, although the point estimate is slightly positive.

5.2 Between-Period Effects: Dynamics of Adaptation

We next move to studying the dynamic consequences of the ex-ante adaptation generated by farmers' reaction to one-month ahead forecasts. We first describe our empirical approach, then the profit results, and finally some channels explaining these results.

5.2.1 Framework

It is useful to restate the timing of the ex-ante adaptation we are considering. We are looking at the effect of a marginal variation in the one-month ahead forecast received by farmers on farm-level outcomes, conditional on realized weather. We want to capture the costs and benefits of adaptation actions implemented which react to forecasts conditional on realized weather, and as such the costs and benefits of ex-ante adaptation. In this section, we focus on the dynamic effects of this adaptation, first the potential effects in periods after period t, and second making sure past outcomes are not responsive to the forecast received at t.

A first thing we have to note is that farmers are taking adaptation actions to maximize the present discounted value of their expected flow of profits. We only observe realized farm level outcomes in future periods, and not their counterfactual values under the other possible states of nature. We hence focus on the effects of adaptation on these realizations. This focus introduces a potential bias in our error term, in the form of the difference between the expected average outcome when implementing the adaptation action, and the realized outcome. This difference needs to be orthogonal to the forecast at t in order not to bias our results. Assuming that forecasts at t are redundant in terms of information about future weather outcomes once we control for weather realizations at t, and assuming that the forecasts are unbiased, then controlling for period t realized weather should be enough to address this potential bias when looking at future periods. These informational redundancy and unbiasedness assumptions should also be sufficient to remove any potential correlation between current forecasts and future realized and forecasted weather shocks which will also influence future farm outcomes.

When looking at past periods, $w_{d(j)t}$ will however not be sufficient to remove the correlation between past realized and forecasted weather and period t forecasted weather. We will add past realized weather as additional controls to remove this correlation which could bias the coefficients. We further add these controls for all our regressions on both past, present and future outcomes.

Our design then follows, for outcome $y_{j,t+h}$, with $t \in [-2, 3]$ for farm j, in department d and region r:

$$
\begin{aligned}
y_{j,t+h} = \sum_{k=-2}^{0} & \beta^{w}_{1,h,k} GDD_{d(j),t+k} + \beta^{w}_{2,h,k} HDD_{d(j),t+k} + g^{h,k}\left(P_{d(j),t+k}\right) + \\
& \beta^{f}_{1,h} FGDD_{d(j),t} + \beta^{f}_{2,h} FHDD_{d(j),t} + g^{h}\left(FP_{d(j),t}\right) + \\
& \zeta^{1}_{d(j)} t + \zeta^{2}_{d(j)} t^2 + \gamma_j + \eta_t + \varepsilon_{jt}
\end{aligned}
$$

We keep the same notation as before, where $GDD_{d(j),t+k}$ is period $t+k$ growing degree days for department d in which j is located, $GDD_{d(j),t+k}$ the heating degree days, the function $g^{h,k}(.)$ is a second order polynomial in realized rainfall for the period $t+k$. $FGDD$ denotes a forecasts, and we follow the same notation for the other weather variables. We include a quadratic time trend, which we allow to be region specific, and also include year and farm fixed effects.

We show in Table A21 regressions of different lead and lag values of forecasted heating degree days on the set of regressors specified here. They show that conditional on our controls, there is no significant correlation between period t forecasted degree days and both previous and future period forecasts. Such correlation would bias $\beta^{f}_{2,h}$, and prevent us from interpreting it as the consequences of ex-ante adaptation.

5.2.2 Results

We start by looking in Figure 7 at the effect of ex-ante adaptation on profits. Each coefficient corresponds to a separate regression which follows the design outlined above. The table for the results is Table A22. We recover a value of ex-ante adaptation at t of 1,908, which is close to the one outlined in the previous section. Within the period of the shock, ex-ante adaptation leads to a net increase in farm profit. However, the next two periods are followed by negative and statistically significant coefficients, with values resp. of 1,827 and 1,435. If we take these three periods as the total of the significant effects caused by adaptation at t, and a discount factor of 10%, this implies a present discounted value of ex-ante adaptation of -980€.

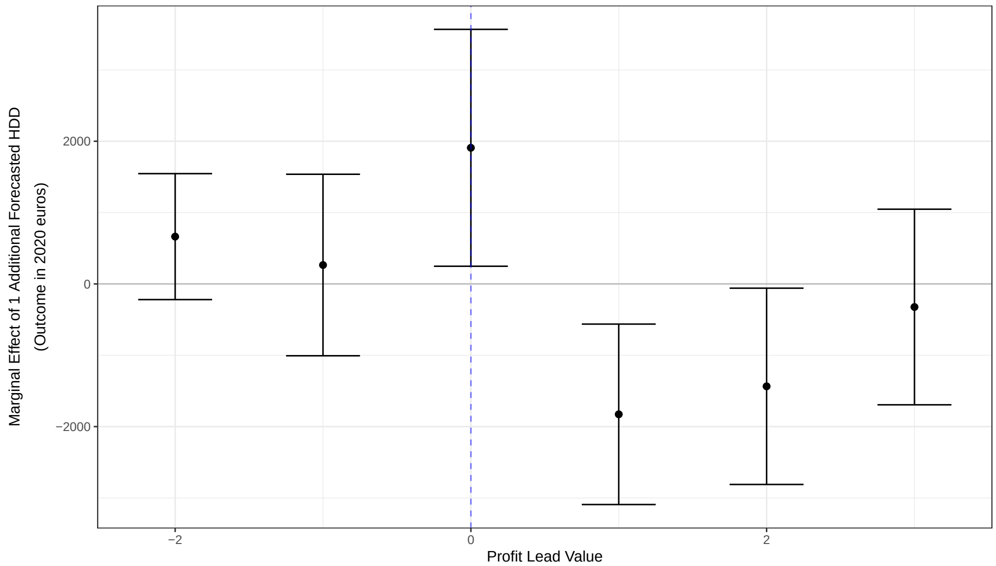

Figure 7: Dynamic Effects of Ex-Ante Adaptation on Profit

In our case, it seems that farms face decreases in profit following their ex-ante adaptation to weather shocks. If we assume that farmers are rational, and foresee the negative consequences of their adaptation actions in future periods, this should imply that the counterfactual of no-adaptation in the face of a marginal change in heating degree days is also a decrease in profit of at least the same value. One of the main adaptation channel in our context is crop switching. This dynamic pattern of immediate gains, likely driven by a lowered sensitivity of profit to heat, followed by future decreases in profits, then justified the presence of important switching costs, which have recently taken a large place in the dynamic land use literature.

We next use the same regression design to study the pattern of production costs across times, showing the graphical results in Figure 8, and the table in Table A23.

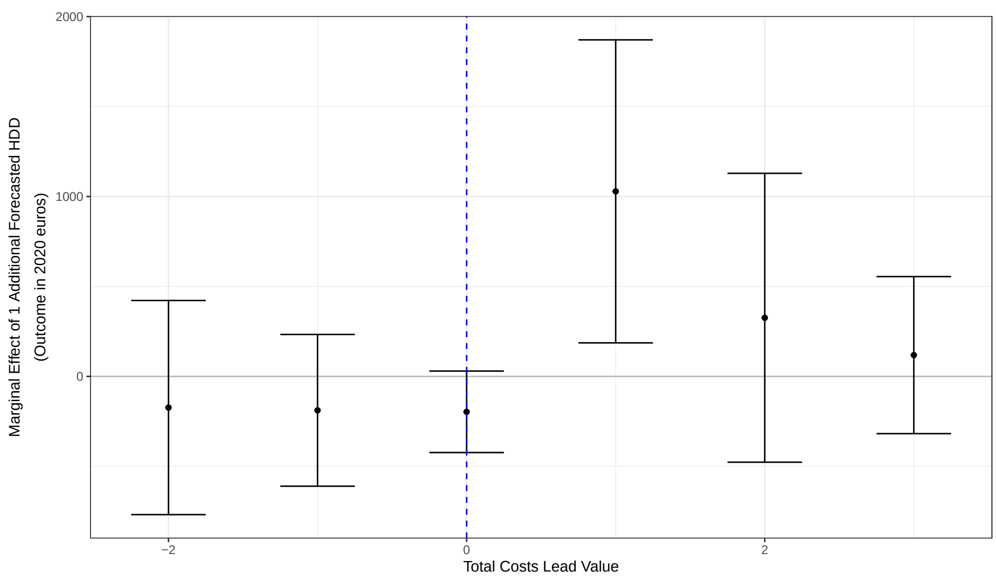

Figure 8: Dynamic Effects of Ex-Ante Adaptation on Production Costs

The costs results track well with the ones for profits. First at period t and as we had seen previously, adaptation does not come with a significant increase in production costs. However, the next period is followed by a net positive increase, and the following one still shows a positive albeit non-significant coefficient. These seem to confirm that there are costs to crop switching. Crop switching is inherently dynamic in the presence of switching costs, as crop choice has an option value. These results show crop switching is also dynamic in the sense that its benefits and costs are not realized at the same time, with gains at the moment of the shock, and losses in the following period. Both Livingston et al. (2008), Scott (2013) and Burlig et al. (2024) emphasize the effects of a field's past state for its current and future productivity. Here, we show that farms respond to disruptions in their crop choices by increasing their production costs, potentially to compensate the decreased productivity of the fields on which crop switching happened. We conclude by describing the source of these increases in costs, by looking at input specific responses in Figure 9. These are useful to illustrate how farms attempt to account for the consequences of past crop choice disruptions.

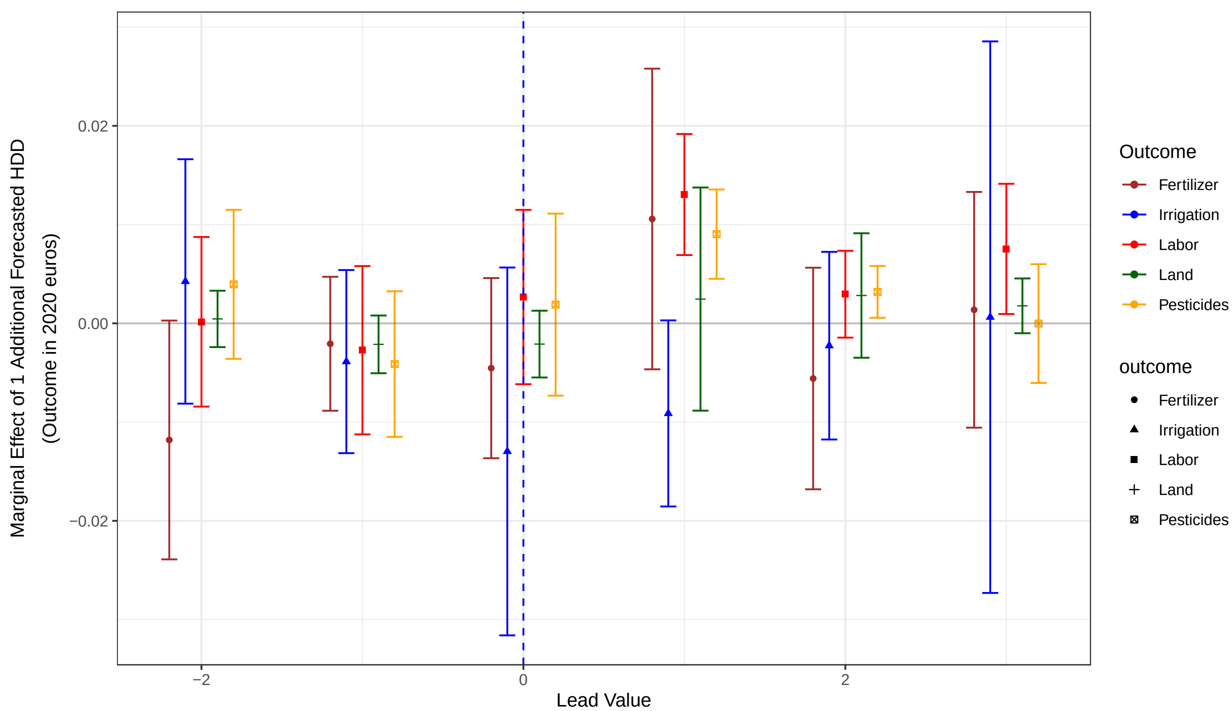

Figure 9: Dynamic Effects of Ex-Ante Adaptation on Input Use

The input responses track with the cost ones. In the period of the forecast, inputs do not respond. However, we observe positive responses in the next two periods, first pesticides, labor and fertilizers in period $t+1$, and then labor and pesticides in period $t+2$.

6 Conclusion

Climate change is expected to have significant and wide ranging impacts on human and non-human systems. Human impacts range from significant out-migrations (Cattaneo et al. (2019) for a review), to negative human health consequences (Deschênes and Greenstone (2011), Carleton (2018), Shrader et al. (2023)), to productivity and growth declines (for example Dell et al. (2012)). Recovering precise estimates of the impacts of climate change in all these areas is important to understand the scope of climate change, and to design policies that could mitigate its effects. These policies also require a good understanding of the costs of adaptation.

Here, we propose a precise analysis of the costs of adapting to heat shocks in French agriculture, a relevant context given the centrality of agriculture for climate change policies, and the relative size of France as an exporter of agricultural goods. We leverage precise accounting farm-level data over 1994-2018, which allows us to track the differential responses of farmers to growing and heating degree days, and to decompose these responses into responses to forecasts, and responses to residual surprise shocks. We highlight how—so far—French farmers have been able to use low cost adaptation strategies to heat forecasts, in order to generate large profit gains. We build confidence in these results, by decomposing the cost and revenue responses into output and input volumes and prices. We highlight the role of timing decisions, and of crop switching, in generating these gains.

This data also allows us to underline how the presence of these net profit gains from additional weather information contradict the assumptions of the envelope theorem, frequently used in the climate literature. Indeed, in our context, farmers' response to small variations in forecasted heating degree days have net positive profit impacts. We decompose the variation in weather used to build growing and heating degree days, and show how heating degree days are mostly composed of non-marginal variation ill-suited to the envelope theorem. Furthermore, the presence of no-cost adaptation strategies like timing, imply that some types of adaptation responses are likely to have non-continuous effects of profits, again going against the usual assumptions of the envelope theorem.

Finally, we show how the low cost, high revenue responses of French farmers to heat forecasts are specific to the period we study, and are unlikely to hold in a future warmer world. Focusing on the upper tail of our heating degree day distribution, we show how then, farmers lose their ability to generate profit gains from heating degree days forecasts. On the contrary, they are now forced to implement costly adaptation strategies, in the form of input increases, to protect their production. Similarly, farmers constrained in their ability to modify their crop-mix composition do not show the same ability to generate profit gains from heat forecasts.

References

Aragón, F. M., F. Oteiza, and J. P. Rud (2021): "Climate Change and Agriculture: Subsistence Farmers' Response to Extreme Heat," *American Journal of Economics: Economic Policy.*

Auffhammer, M. (2018): "Quantifying Economic Damages from Climate Change," *Journal of Economic Perspectives*, 32, 33–52.

Bareille, F. and R. Chakir (2023): "Structural Identification of Weather Impacts on Crop Yields: Disentangling Agronomic from Adaptation Effects," *American Journal of Agricultural Economics.*

Braun, T. and W. Schlenker (2023): "Cooling Externality of Large-Scale Irrigation," *NBER Working Paper.*

Burke, M. and K. Emerick (2016): "Adaptation to Climate Change: Evidence From U.S. Agriculture," *American Economic Journal: Economic Policy*, 8, 106–140.

Burlig, F., L. Preonas, and M. Woerman (2024): "Groundwater and Crop Choice in the Short and Long Run," *Working Paper.*

Canal, N. (2015): "Application a l'agriculture de la prevision saisonniere: evaluation a l'echelle de la France," *PhD Thesis, Universite Paul Sabatier Toulouse III.*

Carleton, T., A. Jina, M. Delgado, M. Greenstone, T. Houser, S. Hsiang, A. Hultgren, R. E. Kopp, K. E. McCusker, I. Nath, et al. (2022): "Valuing the Global Mortality Consequences of Climate Change Accounting for Adaptation Costs and Benefits," *Quarterly Journal of Economics*, 137, 2037–2105.

Carleton, T. A. (2018): "Crop-damaging temperatures increase suicide rates in India," *Proceeding of the Natioanal Academy of Sciences.*

Carleton, T. A. and S. M. Hsiang (2016): "Social and economic impacts of climate," *Science*, 353, aad9837.

Cattaneo, C., M. Beine, C. J. Fröhlich, D. Kniveton, I. Martinez-Zarzoso, M. Mastrorillo, K. Millock, E. Piguet, , and B. Schraven (2019): "Human Migration in the Era of Climate Change," *Review of Environmental Economics and Policy.*

Colas-Belcour, F., R. Renoult, and M. Vallance (2015): "Synthèse Eau et Agriculture," *Rapport de synthèse de missions d'expertise n°14061 CGAAER.*

Cui, X., B. Gafarov, D. Ghanem, and T. Kuffner (2023): "On Model Selection Criteria for Climate Change Impact Studies," *Journal of Econometrics.*

Dell, M., B. F. Jones, and B. A. Olken (2012): "Temperature Shocks and Economic Growth: Evidence from the Last Half Century," *American Economic Journal: Macroeconomics.*

Deryugina, T. and S. Hsiang (2017): "The Marginal Product of Climate," *NBER Working Paper.*

Deschênes, O. and M. Greenstone (2007): "The Economic Impacts of Climate Change: Evidence From Agricultural Output and Random Fluctuations in Weather," *American Economic Review*, 97, 354–385.

Deschênes, O. and M. Greenstone (2011): "Climate Change, Mortality, and Adaptation: Evidence from Annual Fluctuations in Weather in the US," *American Economic Journal: Applied Economics.*

Downey, M., N. Lind, and J. G. Shrader (2023): "Adjusting to Rain Before It Falls," *Management Science*.

Elferjani, R. and R. Soolanayakanahally (2018): "Canola Responses to Drought, Heat, and Combined Stress: Shared and Specific Effects on Carbon Assimilation, Seed Yield, and Oil Composition," *Frontiers in Plant Science*.

Feng, S., M. Oppenheimer, and W. Schlenker (2015): "Weather Anomalies, Crop Yields, and Migration in the US Corn Belt," *Working Paper*.

Gammans, M., P. Merel, and A. Ortiz-Bobea (2017): "The Impact of Climate Change on Cereal Yields: Statistical Evidence from France," *Environmental Research Letters*.

Guo, C. and C. Costello (2013): "The Value of Adaption: Climate Change and Timberland Management," *Journal of Environmental Economics and Management*, 65, 452–468.

Hsiang, S. M. (2016): "Climate Econometrics," *Annual Review of Resource Economics*, 8, 43–75.

Kala, N. (2019): "Learning, Adaptation, and Climate Uncertainty: Evidence from Indian Agriculture," *Working Paper*.

Kelly, D. L., C. D. Kolstad, and G. T. Mitchell (2005): "Adjustment Costs fro Environmental Change," *Journal of Environmental Economics and Management*.

Livingston, M., M. Roberts, and J. Rust (2008): "Optimal Corn and Soybean Rotations," *Working Paper*.

Mendelsohn, R., W. D. Nordhaus, and D. Shaw (1994): "The Impact of Global Warming on Agriculture: A Ricardian analysis," *American Economic Review*, 84, 753–771.

Merel, P., E. Paroissien, and M. Gammans (2022): "Sufficient Statistics for Climate Change Counterfactuals," *Working Paper*.

Moscona, J. and K. A. Sastry (2022): "Does Directed Innovation Mitigate Climate Damage? Evidence from U.S. Agriculture," *Quarterly Journal of Economics*.

Schauberger, B., H. Kato, T. Kato, D. Watanabe, and P. Ciais (2022): "French crop yield, area and production data for ten stable crops from 1900 to 2018 at county resolution," *Scientific Data*.

Schlenker, W., M. Hanemann, and A. C. Fisher (2005): "Will U.S. Agriculture Really Benefit from Global Warming? Accounting for Irrigation in the Hedonic Approach," *American Economic Review*, 95, 395–406.

Schlenker, W. and M. J. Roberts (2009): "Nonlinear Temperature Effects Indicate Severe Damages to U.S. Crop Yields Under Climate Change," *Proceedings of the National Academy of Sciences*, 106, 15594–8.

Scott, P. (2013): "Dynamic Discrete Choice Estimation of Agricultural Land Use," *Working Paper*.

Shrader, J. G. (2023): "Improving Climate Damage Estimates by Accounting for Adaptation," *Working Paper*.

Shrader, J. G., L. Bakkensen, and D. Lemoine (2023): "Fatal Errors: The Mortality Value of Accurate Weather Forecasts," *NBER Working Paper*.

Verhoogen, E. (2023): "Firm-Level Upgrading in Developing Countries," *Journal of Economic Literature*.

Appendix for online publication:
Costs of Climate Adaptation: Evidence from French Agriculture

A Additional Figures and Tables

A.1 Figures

A.1.1 Precision of Forecasts

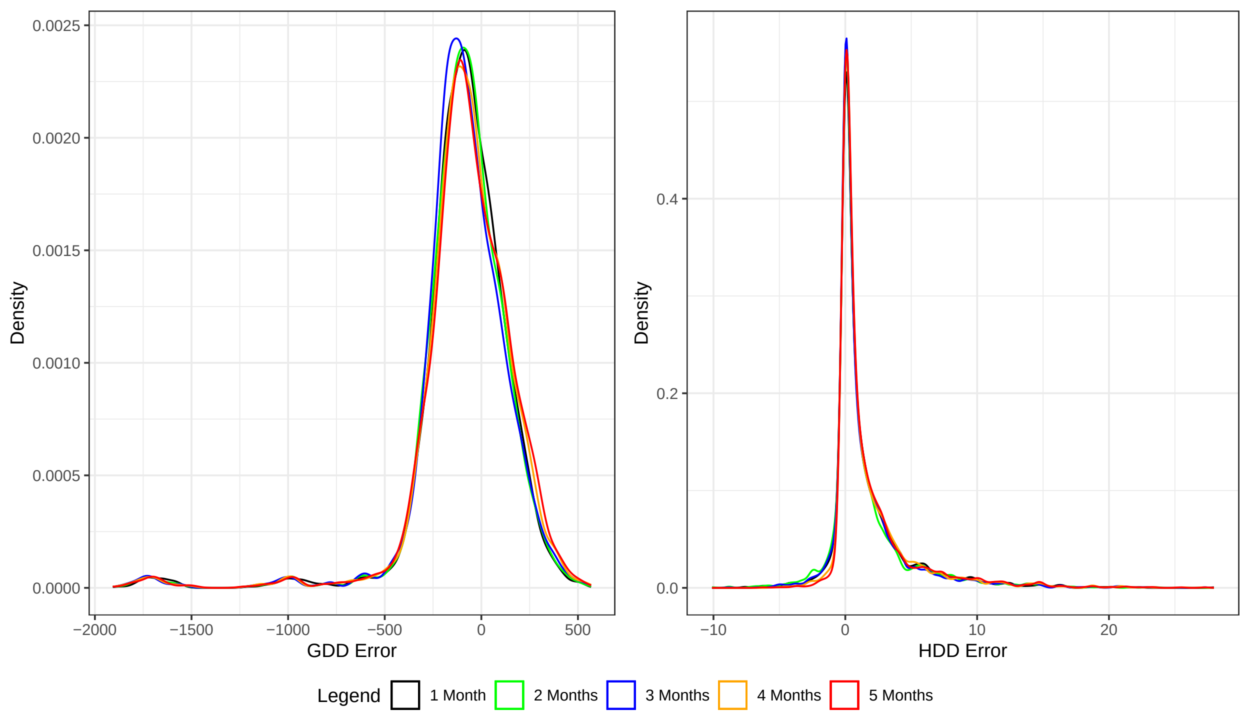

Figure A1: Distribution of Temperature Forecast Error

Notes: These estimated kernel densities show the empirical distribution of the forecast errors (realized minus forecasted), for both growing and heating degree days (aggregated over the growing season) in France over the 1994-2018 period. We later cut our sample so GDD errors are above -500, cutting the left tail of the GDD error distribution. As expected, HDDs correspond to extreme events which are harder to predict, and are under predicted in France.

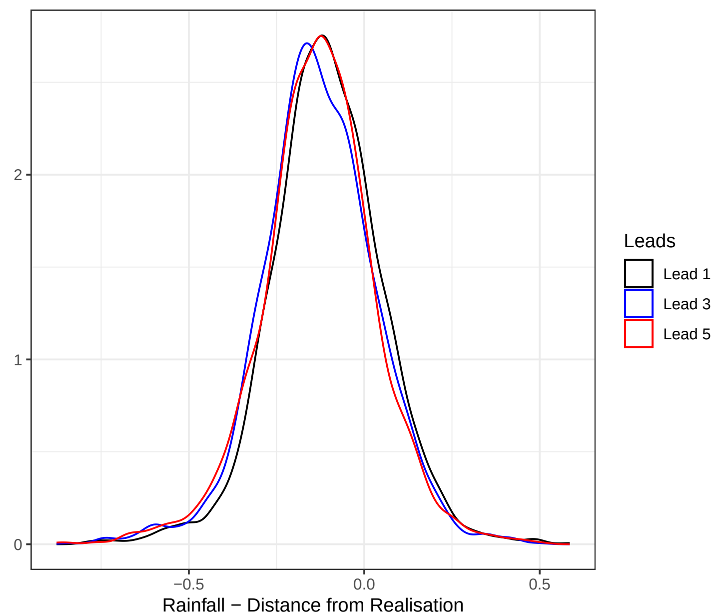

Figure A2: Distribution of Rainfall Forecast Error

Notes: These estimated kernel densities show the empirical distribution of the forecast errors for rainfall (aggregated over the growing season) in France over the 1994-2018 period.

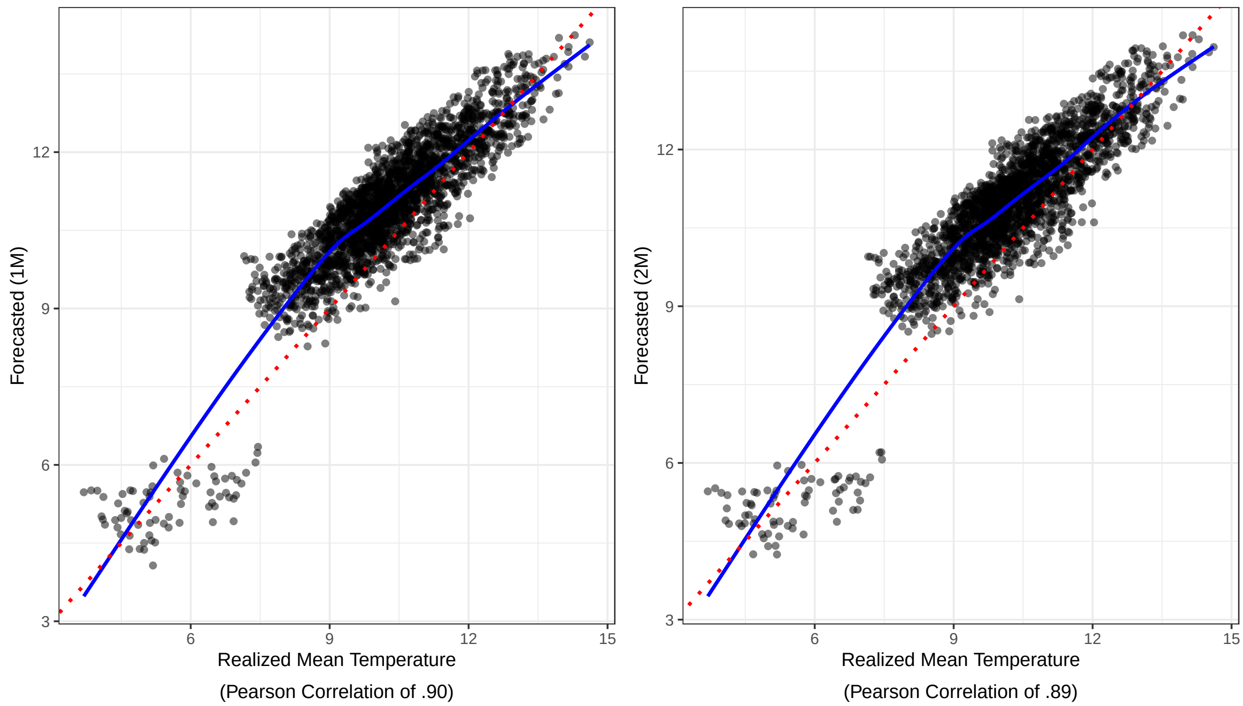

Figure A3: Correlation Realized and Forecasted Mean Temperature

Notes: We show the correlation between forecasted mean temperatures and realized ones over 1994-2018 in France. We show the correlation for the sample used for estimation, that is for the set of observations with one month ahead forecasted GDD errors above -500. The red line corresponds to the 45°line, and the blue line to a smoothed estimator matching the distribution of points.

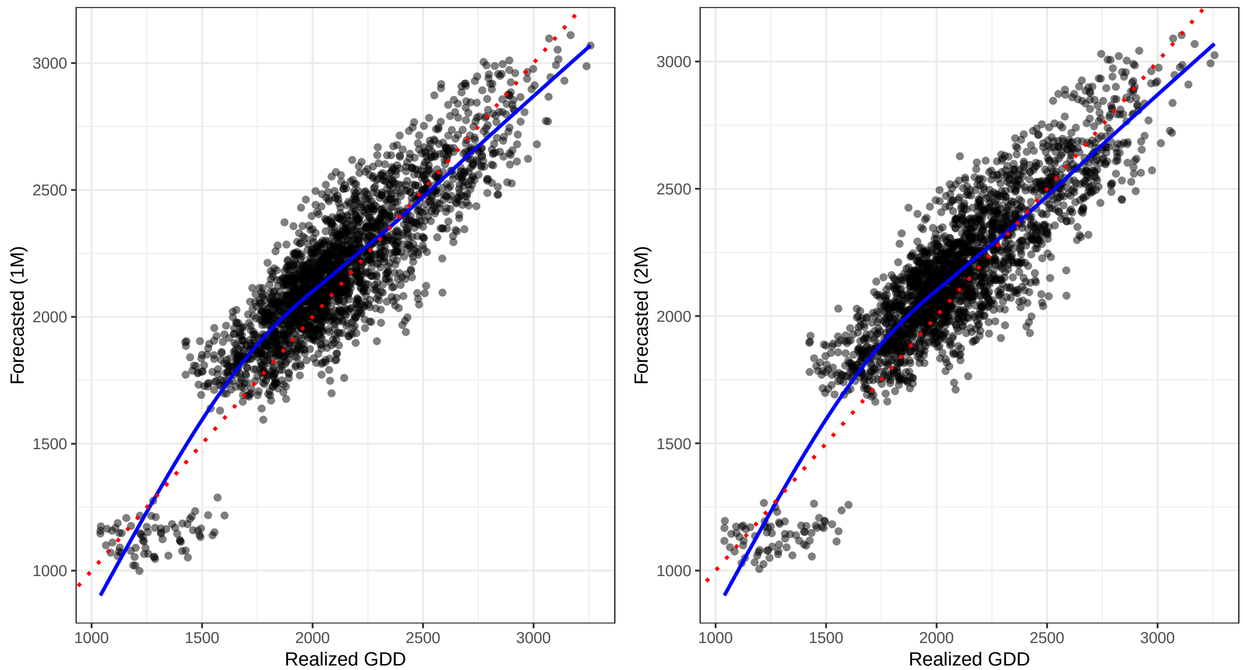

(Pearson Correlation of .88) (Pearson Correlation of .87)

Figure A4: Correlation Realized and Forecasted GDDs

Notes: We show the correlation between forecasted GDDs and realized ones over 1994-2018 in France. We show the correlation for the sample used for estimation, that is for the set of observations with one month ahead forecasted GDD errors above -500. The red line corresponds to the 45°line, and the blue line to a smoothed estimator matching the distribution of points.

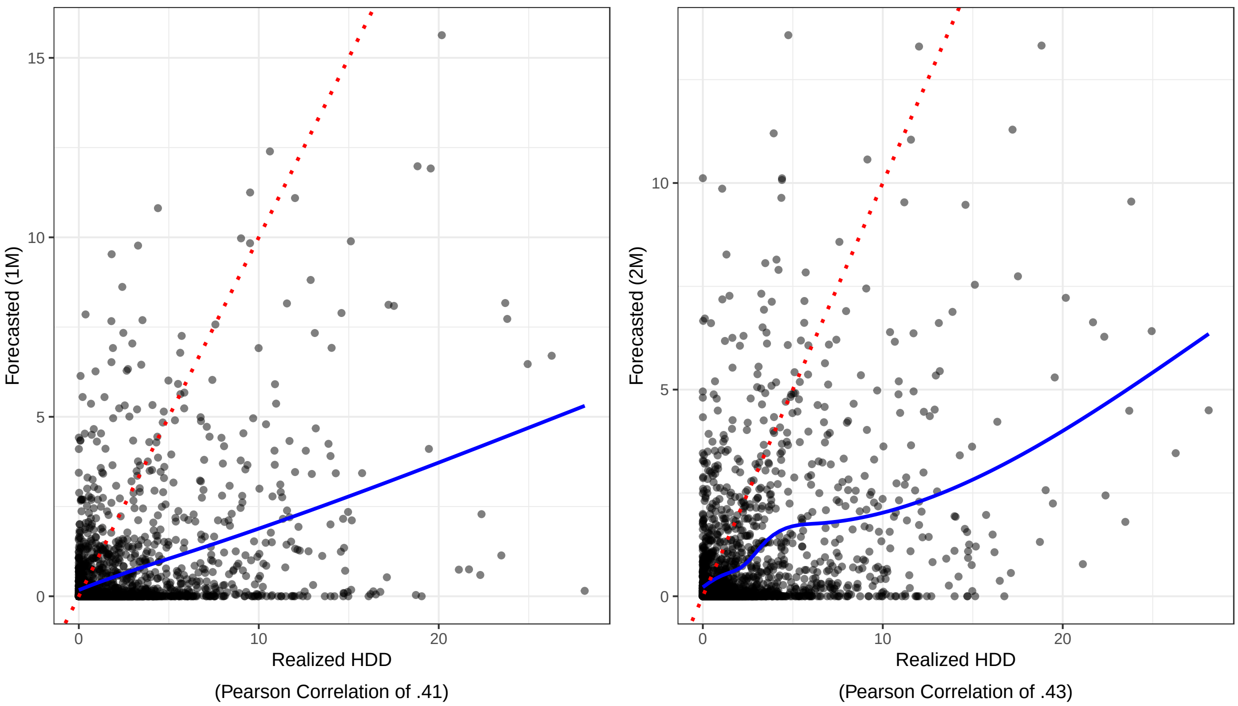

Figure A5: Correlation Realized and Forecasted HDDs

Notes: We show the correlation between forecasted HDDs and realized ones over 1994-2018 in France. We show the correlation for the sample used for estimation, that is for the set of observations with one month ahead forecasted GDD errors above -500. The red line corresponds to the 45°line, and the blue line to a smoothed estimator matching the distribution of points.

A.1.2 Climate in France

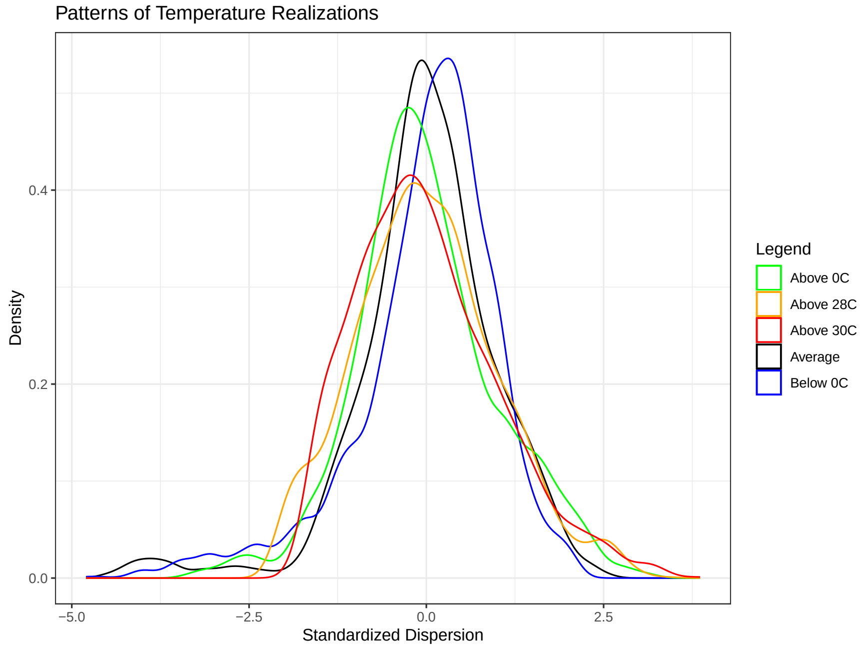

Figure A6: Distribution of the Standardized Conditional and Unconditional Mean Temperature

Notes: We show kernel density estimators giving the distribution of conditional standardized temperature realizations in France, at the department level, over 1994-2018. As expected, higher realizations have distributions with a larger spread. On average, however, realizations are quite homogeneous across the country.

These graphs are useful to highlight the spatial variation in exposure to heat in France. We first observe the divide between the South of the country more exposed to heating degree days than the center and North. Second, we see how mountainous regions in the center, around the Massif Central, the Alpes and the Pyrénées have lower growing degree days values.

The main cereal region of the country situated in the large plains below Paris up to the Massif Central have overall large growint degree day values, and low heating degree day values for the 1994-2018 period.

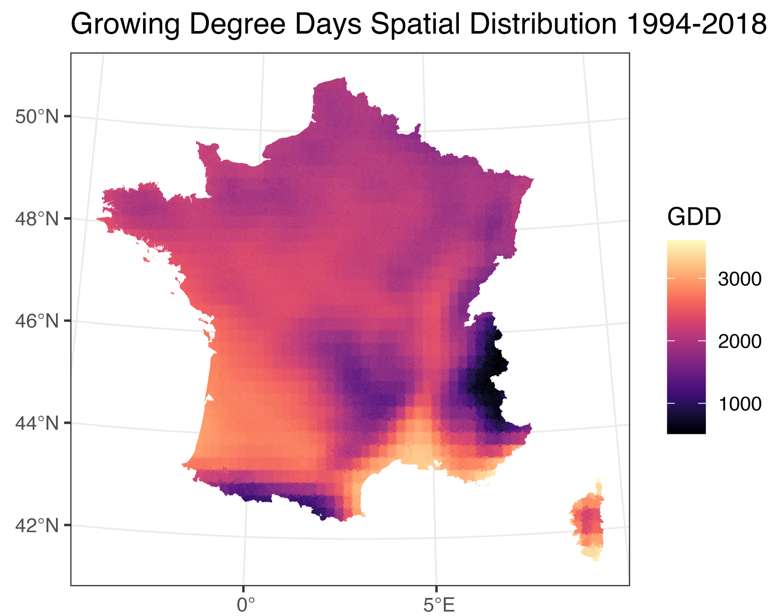

Figure A7: Distribution of Growing Degree Days

Notes: We show a map of average GDD realizations in France over 1994-2018 at the department level. The large geographical patterns are that temperature is on average higher in the South along the Mediterranean coast, and lower in mountain regions (Massif Central in the center, Pyrenees in the South at the border with Spain, and the Alpes at the border with Switzerland and Italy).

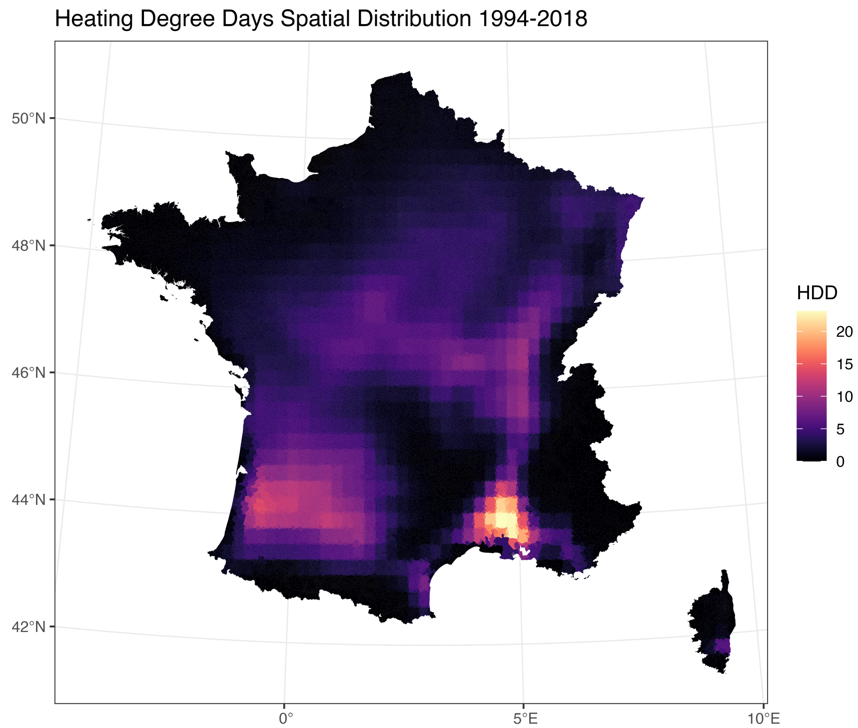

Figure A8: Distribution of Heating Degree Days

Notes: We show a map of average HDD realizations in France over 1994-2018 at the department level. HDDs are on average close to zero, with positive values in the South both around Marseilles, and in the agricultural region stretching between Toulouse and Bordeaux.

A.1.3 Growing Seasons

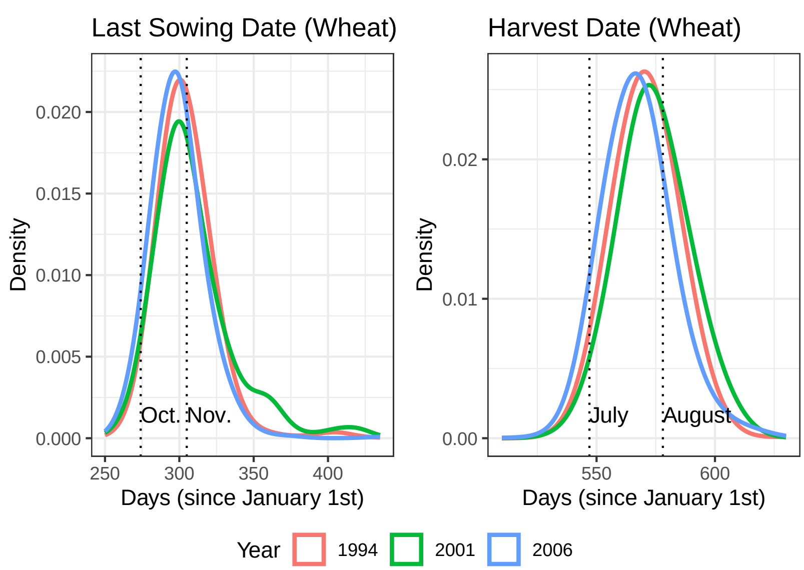

Figure A9: Wheat Growing Season in France

Notes: This figure shows the distribution of a proxy for the start and end of the wheat growing season as observed at the plot level in three cross-sectional surveys of growing practices. A unique kernel density describes spatial and cross-farm heterogeneity within a same year, while the variation across kernels shows variation across growing seasons. The x-axis shows the number of days since January 1st of the first of two calendar years overlapped within a unique agricultural season.

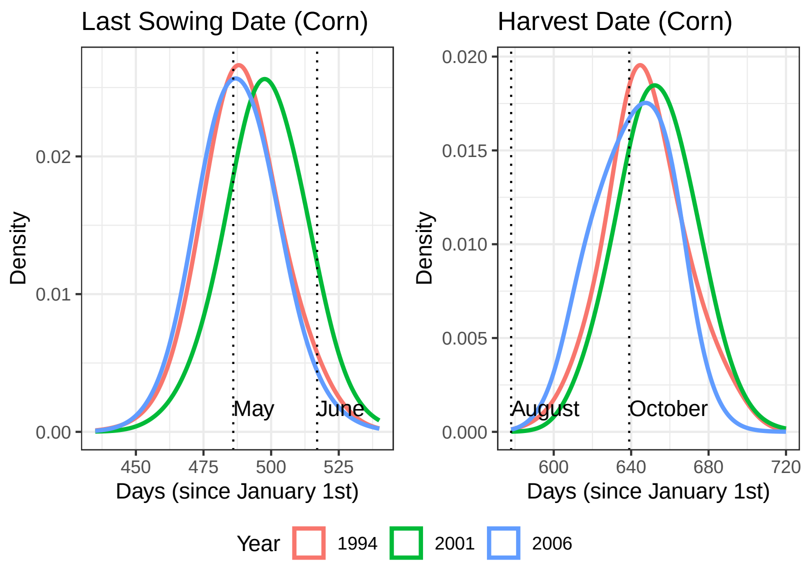

Figure A10: Corn Growing Season in France

Notes: This figure shows the distribution of a proxy for the start and end of the corn growing season as observed at the plot level in three cross-sectional surveys of growing practices. A unique kernel density describes spatial and cross-farm heterogeneity within a same year, while the variation across kernels shows variation across growing seasons. The x-axis shows the number of days since January 1st of the first of two calendar years overlapped within a unique agricultural season.

A.1.4 Yield and Output Effects of Weather Variation

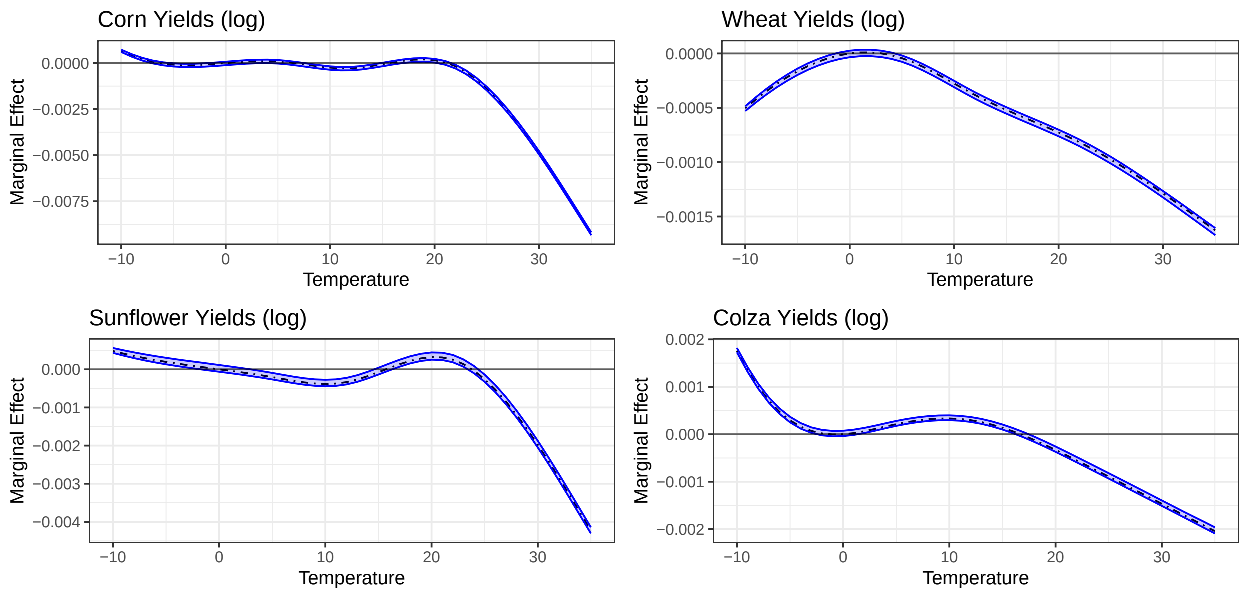

Figure A11: Temperature Effects on Crop-Specific Yields

Notes: This figure shows restricted cubic splines describing the non-linear relation between exposure to temperature during the growing season and crop yields. The regressions are farm-level versions of the specification from Schlenker and Roberts (2009). In addition to cubic splines in temperature, the regressions include farm fixed effects and region-specific quadratic time trends as controls. Standard errors are computed by bootstrapping at the farm level. Temperature effects are normalized relative to the impact at 0°C.

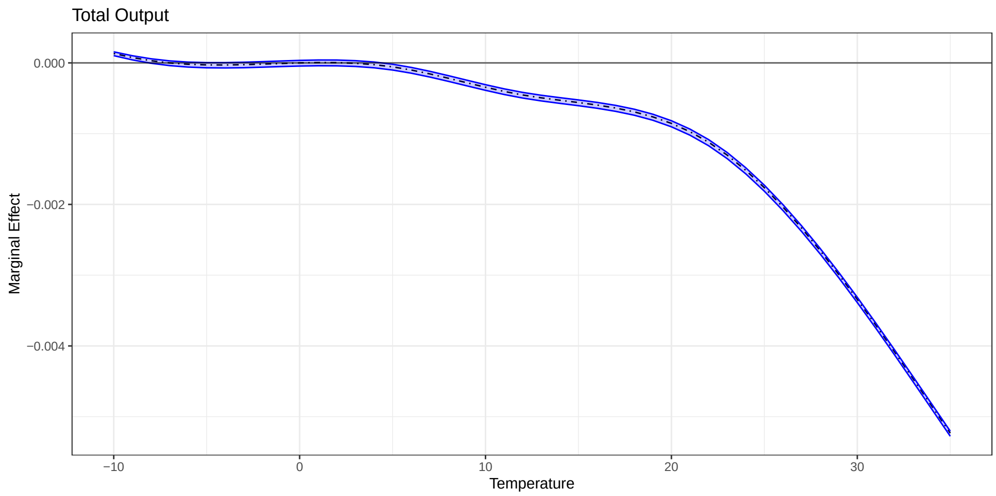

Figure A12: Temperature Effects on Farm Total Output

Notes: This figure shows restricted cubic splines describing the non-linear relation between exposure to temperature during the growing season and farm output. Output is measured as the sum of output for wheat, durum, oats, corn, sorghum, barley, rye, triticale, sunflower, colza, soy, peas and fava. Our regressions are similar to that of Schlenker and Roberts (2009).The regressions are farm-level versions of the specification from Schlenker and Roberts (2009). In addition to cubic splines in temperature, the regressions include farm fixed effects and region-specific quadratic time trends as controls. Standard errors are computed by bootstrapping at the farm level. Temperature effects are normalized relative to the impact at 0°C.

A.1.5 Multi-Product Farms

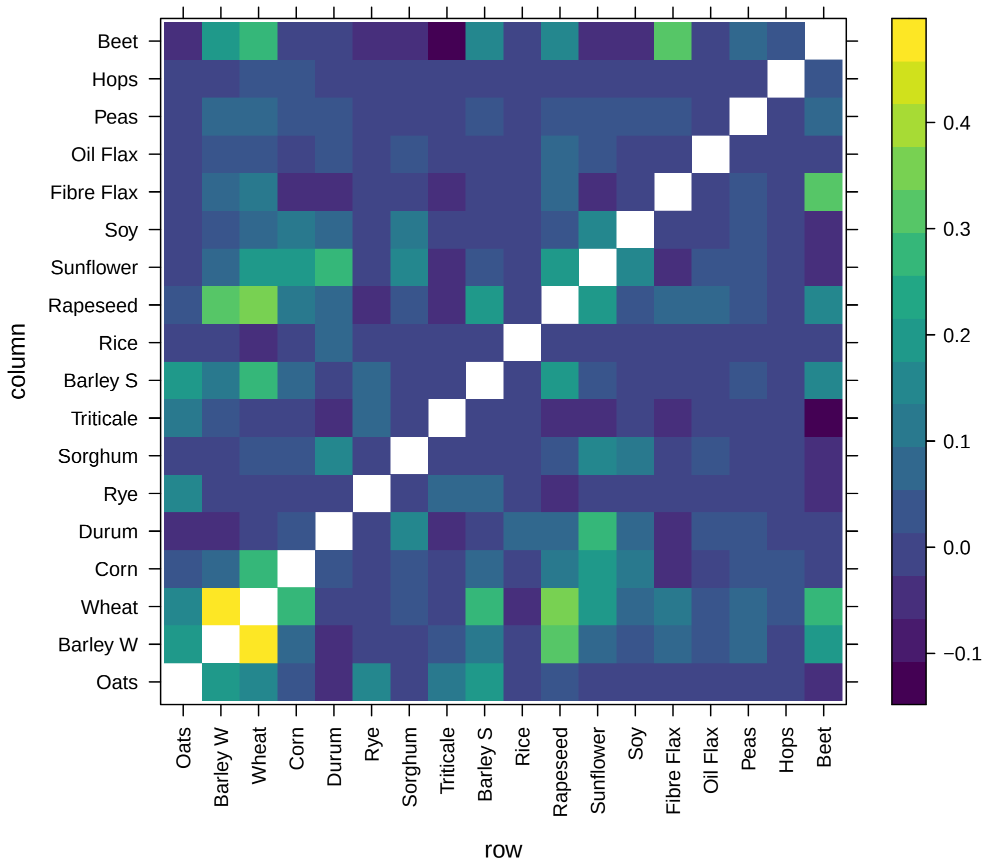

Figure A13: Conditional Probabilities to Grow Crop Pairs

Notes: This figure shows the probability that a farm in the FADN grows a given crop, conditional on growing the row-specified one.

Table A1: Descriptive Statistics - Farm-Level Crop Mix

	Cereals	Oil-Protein	Industrial
Cereals	1	0.3321	0.0999
Oil-Protein	0.9873	1	0.1321
Industrial	0.9991	0.4444	1

A.1.6 Input Price Indices

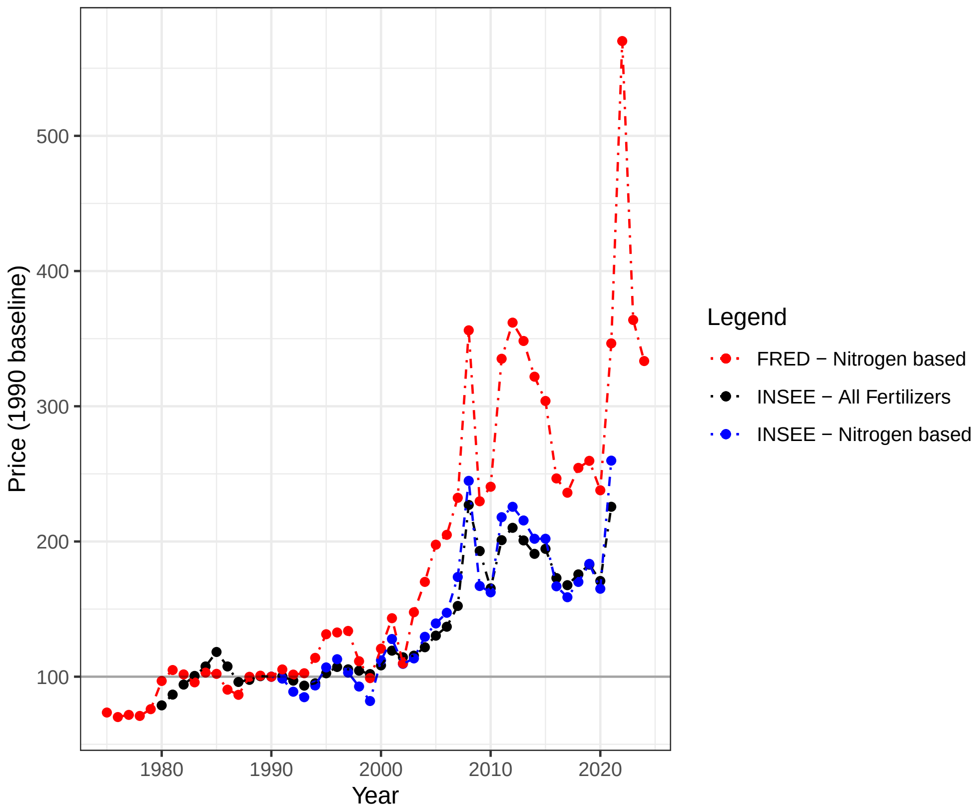

Figure A14: Comparison of Input Price Indices Series

Notes: This figure shows the input price index used for fertilizers used in the paper, and coming from the French statistical agency (INSEE), and compares this index, and its subset specific to nitrogen-based fertilizers to the Producer price index for fertilizer-based fertilizers from FRED for the USA.

A.1.7 Additional Results

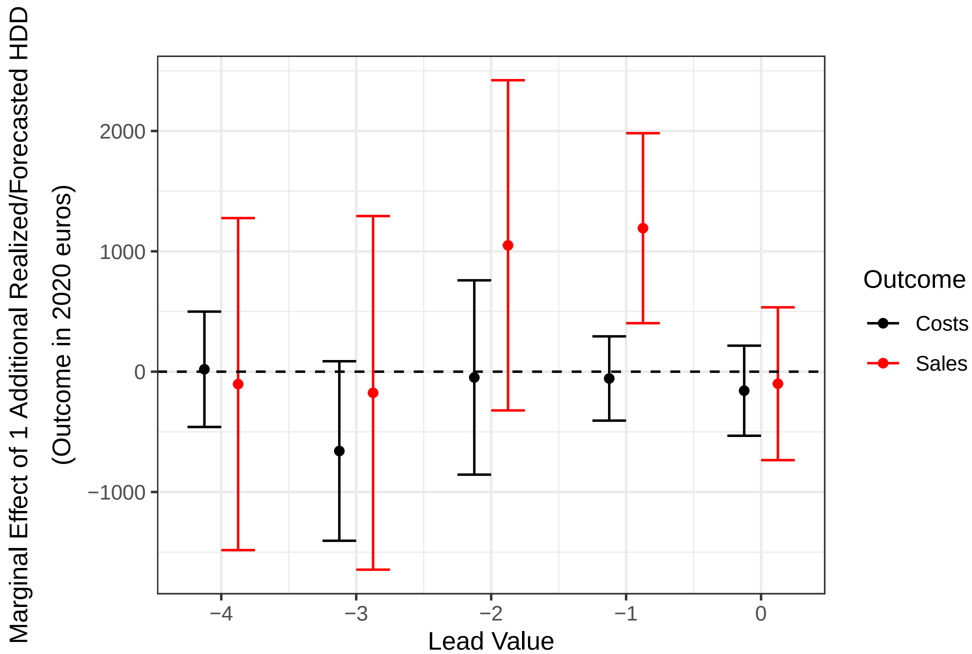

Figure A15: Effect of Forecated HDDs

Notes: We show the results of our main specification, varying the forecast lead value used in the regression. For the forecast of lead 0, we only include the realization of the weather shocks. As such, the graph compares the effect of expected HDD shocks across independent regressions, and serves as a robustness test of our results.

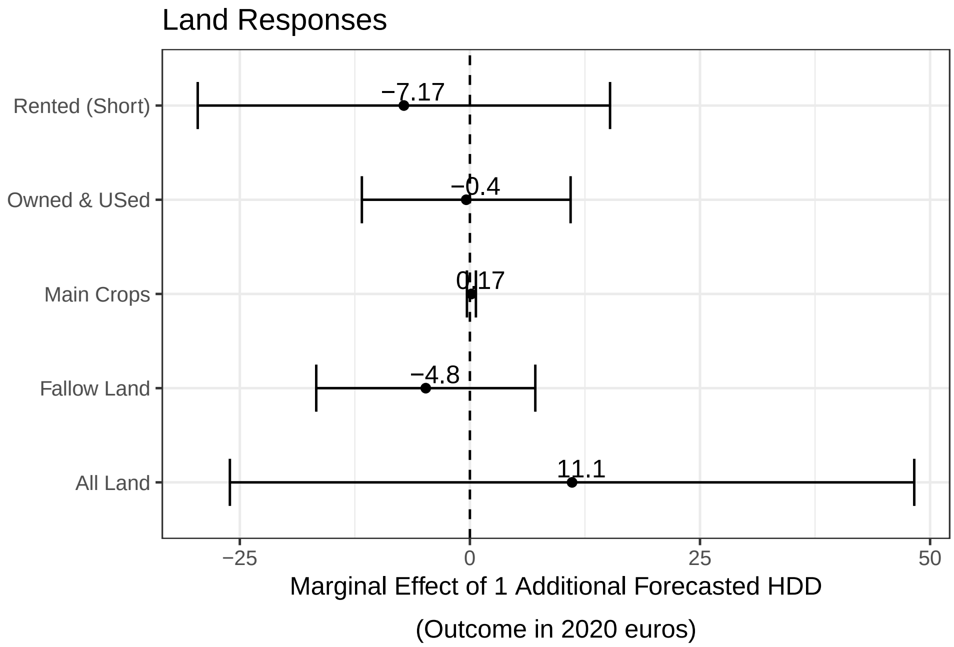

Figure A16: Decomposition of the Land responses

Notes: We show the results of our main specification, showing the reaction of different land area margins. Each regression contains realized and forecasts weather outcomes (rainfall and temperature), as well as region-specific quadratic time trends, farm and year fixed effects. Standard errors are clustered at the department-by-year level.

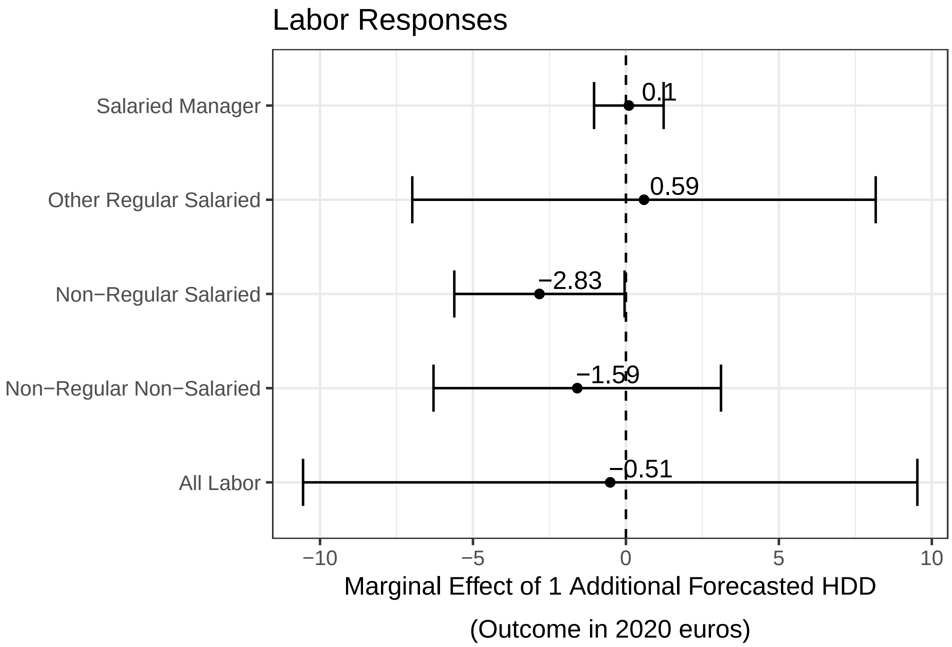

Figure A17: Decomposition of Labor Responses

Notes: We show the results of our main specification, showing the reaction of different labor margins. Each regression contains realized and forecasts weather outcomes (rainfall and temperature), as well as region-specific quadratic time trends, farm and year fixed effects. Standard errors are clustered at the department-by-year level.

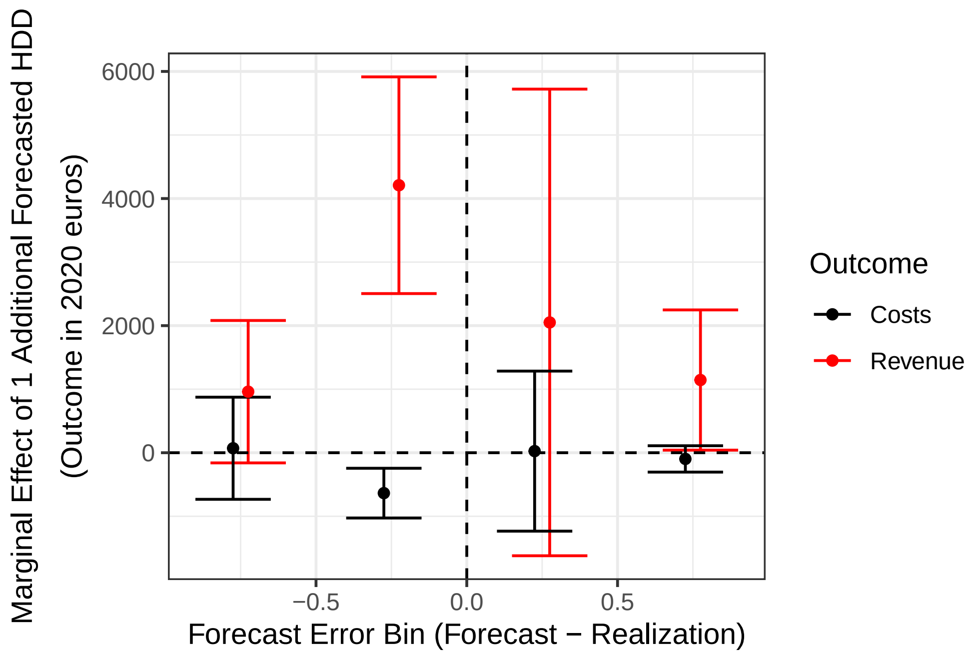

Figure A18: Cost and Revenue & Sign of the Forecast Error

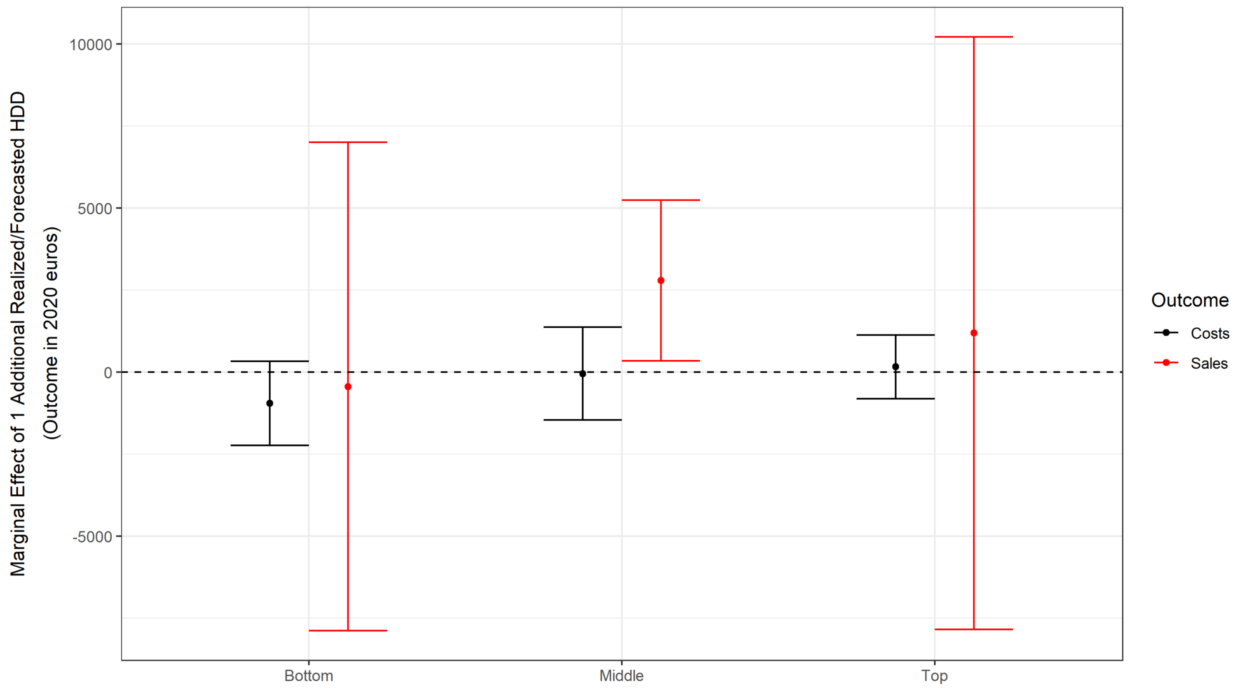

Figure A19: Limits to Adaptation

Notes: We rank French departments by their average realized HDD values within the years of our sample, and compute the marginal cost and revenue responses to a one unit forecasted HDD shock independently for each third of the dataset. Further evidence shows that the middle and top bins both feature crop-switching, hence that farmers seem to respond in the same way in these two bins but do not face the same benefits from their actions.

A.2 Tables

Table A2: Descriptive Statistics - Farm-Level Dataset

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max
Rainfall	18,917	0.755	0.172	0.372	0.634	0.848	1.634
Rainfall (F1)	18,917	0.854	0.126	0.608	0.759	0.921	1.747
Rainfall (F2)	18,917	0.874	0.129	0.649	0.780	0.944	1.812
GDD	18,917	2,140.454	263.726	1,132.679	1,957.350	2,292.319	3,239.741
GDD (F1)	18,917	2,182.291	237.644	1,053.609	2,022.225	2,305.369	3,015.025
HDD	18,917	2.362	3.202	0.000	0.243	3.131	23.808
HDD (F1)	18,917	0.438	1.120	0	0	0.4	16
Sales	18,917	155,386.200	113,739.700	1,106.436	78,738.730	199,717.400	1,333,809.000
Total Costs	18,917	123,249.100	85,427.280	6,955.080	66,903.320	156,468.300	1,035,695.000
Intermediate Inputs	18,917	103,749.600	67,148.830	5,860.800	59,074.470	132,588.800	841,067.000
Value Added	18,917	58,356.790	62,857.780	−159,167.000	18,305.820	80,436.330	875,181.400
Profit	18,917	86,695.230	70,268.490	−135,056.500	40,065.860	114,770.700	973,561.100
Price Index	18,917	263.040	296.369	1.463	140.488	230.596	11,322.960
Storage (sum)	18,917	88.646	1,877.100	−28,569	−280	500	25,085
Storage (index)	16,075	65.407	983.746	−19,787.310	−67.932	176.113	25,085.000
Output (sum)	18,917	7,793.042	5,033.159	146	4,252.4	10,170.6	58,374
Output (index)	18,917	2,400.654	2,926.340	27.588	768.565	2,846.639	58,374.000
Production (corn)	10,346	3,464.252	4,046.307	0.000	877.250	4,682.250	58,374.000
Price (corn)	9,958	150.008	39.116	36.012	122.962	172.746	1,332.181
Sales (corn)	10,346	53,507.850	69,739.220	0.000	11,162.650	68,977.980	1,155,906.000
Quantity Sold (corn)	10,346	3,437.496	4,133.216	0.000	817.000	4,617.000	68,524.000
Production (wheat)	17,336	3,840.134	3,005.148	0.000	1,665.300	5,248.115	31,452.000
Price (wheat)	17,181	163.115	37.955	16.831	136.366	184.217	1,768.937
Sales (wheat)	17,336	61,538.020	53,985.880	0.000	24,344.600	82,921.740	660,068.800
Quantity Sold (wheat)	17,336	3,798.700	3,162.210	0.000	1,560.000	5,166.000	41,140.000
Irrigation	18,917	683.145	2,664.963	0	0	0	46,053
Labor	18,917	2,684.032	1,329.925	160	1,600	3,200	13,200
Phytosanitary	18,917	23,072.220	16,499.150	0.000	11,388.300	30,872.050	189,311.100
Fertilizer	18,917	31,777.250	20,726.760	0.000	17,456.120	41,060.730	217,676.500
Land	18,917	14,504.110	8,619.045	948	8,379	18,709	79,549
Seeds	18,917	12,231.150	9,685.803	0.000	5,847.321	15,882.330	146,058.500

Table A3: Descriptive Statistics - Store-Level Dataset

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max
Pesticide Prices	2,126	615.503	718.332	23.845	257.649	696.640	9,973.849
Fertilizer Prices	3,098	3,253.702	2,958.929	3.500	2,450.439	3,783.050	125,017.500
Seed Prices	1,831	343.569	1,777.400	28.504	95.647	171.789	48,892.200

Table A4: Descriptive Statistics - Plot-Level Data

Statistic	N	Mean	St. Dev.	Min	Max
Ploughing	32,263	337.475	83.538	15	570
Sowing	39,988	378.568	99.240	15	570
Irrigation	3,775	532.128	24.309	170	630
Harvest	38,818	594.288	41.551	510	720

Table A5: Descriptive Statistics - Land Price Data

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max
Land Prices	3,405	5,753.823	3,084.605	448.000	3,717.000	6,834.000	30,097.360

Table A6: Descriptive Statistics - Weather Outcomes within Samples

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max
Panel A	Entire Sample						
Rainfall	18,926	0.755	0.172	0.372	0.634	0.848	1.634
Rainfall (F1)	18,926	0.854	0.126	0.608	0.759	0.921	1.747
Rainfall (F2)	18,926	0.874	0.129	0.649	0.780	0.944	1.812
GDD	18,926	2,140.355	263.708	1,132.679	1,957.350	2,292.319	3,239.741
GDD (F1)	18,926	2,182.468	237.730	1,053.609	2,022.225	2,305.860	3,015.025
HDD	18,926	2.361	3.202	0.000	0.242	3.131	23.808
HDD (F1)	18,926	0.438	1.120	0	0	0.4	16
Panel B	Worst Region-Years						
Rainfall	1,425	0.678	0.145	0.372	0.557	0.749	1.153
Rainfall (F1)	1,425	0.881	0.114	0.608	0.798	0.946	1.203
Rainfall (F2)	1,425	0.903	0.118	0.686	0.808	0.972	1.190
GDD	1,425	2,417.994	255.672	1,928.005	2,196.093	2,587.399	3,139.565
GDD (F1)	1,425	2,355.650	276.100	1,866.416	2,136.270	2,548.512	3,004.334
HDD	1,425	10.392	3.509	6.126	7.347	12.020	23.808
HDD (F1)	1,425	1.414	2.073	0	0.02	2.1	16

Table A7: Cost and Revenue Reactions to HDD

Dependent Variables:	Sales	Costs	Sales	Costs
Model:	(1)	(2)	(3)	(4)
Variables				
GDD	5.103	1.213	10.18	2.948
	(11.84)	(7.809)	(10.99)	(8.430)
GDD (F)	37.39	30.79**		
	(34.66)	(14.48)		
GDD (F2)			55.55	38.42**
			(38.62)	(17.50)
HDD	-158.3	-191.9	-225.0	-196.6
	(332.2)	(156.9)	(311.9)	(193.9)
HDD (F)	1,192.2***	-56.71		
	(402.7)	(178.5)		
HDD (F2)			1,049.7	-48.19
			(699.9)	(411.9)
Mean	155,386.2	123,249.1	155,386.2	123,249.1
Unique Farms	2,603	2,603	2,603	2,603
Fixed-effects				
Farm	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Fit statistics				
Observations	18,917	18,917	18,917	18,917
R^2	0.88715	0.93666	0.88715	0.93667

Notes. Two-way department-by-year standard-errors in parentheses. Stars indicate estimate is significantly different from zero: *$p < .10$, **$p < .05$, ***$p < .01$. Realized and forecasted rainfall in levels and squares are included as controls, as well as quadratic region-specific time trends. Sales correspond to total sales at the farm levels, and costs to total costs.

Table A8: Cost, Revenue and Profit Reactions to HDD

Dependent Variables:	Revenue	Costs	Profit	Value Added
Model:	(1)	(2)	(3)	(4)
Variables				
GDD	5.103	1.213	-8.322	2.597
	(11.84)	(7.809)	(13.16)	(12.68)
GDD (F)	37.39	30.79**	3.257	16.35
	(34.66)	(14.48)	(43.14)	(43.10)
HDD	-158.3	-191.9	184.6	53.36
	(332.2)	(156.9)	(277.2)	(292.0)
HDD (F)	1,192.2***	-56.71	2,066.7***	2,429.3***
	(402.7)	(178.5)	(708.1)	(809.1)
Mean	155,386.2	123,249.1	86,695.2	58,356.8
Unique Farms	2,603	2,603	2,603	2,603
Fixed-effects				
Farm	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Fit statistics				
Observations	18,917	18,917	18,917	18,917
R^2	0.88715	0.93666	0.84140	0.80546

Notes. Two-way department-by-year standard-errors in parentheses. Stars indicate estimate is significantly different from zero: $*p < .10$, $**p < .05$, $***p < .01$. Realized and forecasted rainfall in levels and squares are included as controls, as well as quadratic region-specific time trends. Sales correspond to total sales at the farm levels, and costs to total costs.

Table A9: Price and Quantity Reactions to Forecasts

Dependent Variables:	Price	Output	Storage	Price	Output	Storage
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
GDD	-0.0275	0.0000	-0.0018**	-0.0248	0.0000	-0.0017**
	(0.0291)	(0.0001)	(0.0007)	(0.0285)	(0.0001)	(0.0007)
GDD (F)	0.0771	0.0002	-0.0002			
	(0.0595)	(0.0004)	(0.0015)			
GDD (F2)				0.1225**	0.0005	0.0004
				(0.0527)	(0.0005)	(0.0016)
HDD	-0.2299	-0.0045*	-0.0021	-0.1357	-0.0049**	-0.0012
	(0.3170)	(0.0022)	(0.0153)	(0.3607)	(0.0022)	(0.0141)
HDD (F)	0.2609	0.0089*	0.0248			
	(1.561)	(0.0046)	(0.0213)			
HDD (F2)				-1.742	0.0024	-0.0027
				(1.146)	(0.0056)	(0.0270)
Mean	205.2	7,793.0	88.65	205.2	7,793.0	88.65
Unique Farms	2,603	2,603	2,183	2,603	2,603	2,183
Fixed-effects						
Farm	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
Observations	18,917	18,917	9,602	18,917	18,917	9,602
R^2	0.73239	0.94659	0.66516	0.73276	0.94670	0.66557

Notes: Two-way department-by-year standard-errors in parentheses. Stars indicate estimate is significantly different from zero: $^*p < .10$, $^{**}p < .05$, $^{***}p < .01$. Realized and forecasted rainfall in levels and squares are included as controls, as well as quadratic region-specific time trends. The price variable corresponds to a price index across all crops sold by the farm, the output is a non-weighted sum of the output quantities, and the storage is a non-weighted sum of the variation in storage across all crops also in quantity. Both output and storage are expressed in logs.

Table A10: Prices and Quantities for Outputs and Inputs (1 month lead)

Dependent Variables:	Output Price	Output (log)	Storage (log)	Input Price	Irrigation	Fertilizer	Phytosanitary
Model:	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Variables							
GDD	-0.0275	1.58×10^{-5}	-0.0018**	-0.1264	0.4377	-2.648	-1.297
	(0.0291)	(0.0001)	(0.0007)	(0.2577)	(0.2677)	(2.659)	(1.256)
GDD (F)	0.0771	0.0002	-0.0002	0.5127	-0.4981	7.640	-0.6958
	(0.0595)	(0.0004)	(0.0015)	(1.061)	(0.6717)	(8.721)	(2.423)
HDD	-0.2299	-0.0045*	-0.0021	-15.64	7.506	-86.74	-16.09
	(0.3170)	(0.0022)	(0.0153)	(10.31)	(6.729)	(63.42)	(36.68)
HDD (F)	0.2609	0.0089*	0.0248	-0.2832	-23.88	-9.140	54.64
	(1.561)	(0.0046)	(0.0213)	(11.63)	(17.36)	(77.49)	(71.49)
Mean	205.2	7,793.0	88.65	1,899.6	683.1	31,777.2	23,072.2
Unique Farms	2,603	2,603	2,183		2,603	2,603	2,603
Fixed-effects							
Farm	Yes	Yes	Yes		Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Company				Yes			
Product				Yes			
Fit statistics							
Observations	18,917	18,917	9,602	33,174	18,917	18,917	18,917
R^2	0.73239	0.94659	0.66516	0.35144	0.86328	0.89547	0.92414

Notes: Two-way department-by-year standard-errors in parentheses. Stars indicate estimate is significantly different from zero: *$p < .10$, **$p < .05$, ***$p < .01$. Realized and forecasted rainfall in levels and squares, as well as quadratic region-specific time trends, are included as controls. The input price variable corresponds to prices observed at the store level in an agricultural input price survey run across France in order to build input price indices.

Table A11: Cost and Revenue Reactions to HDD—Comparison

Dependent Variables:	Revenue		Costs	
Model:	(1)	(2)	(3)	(4)
Variables				
GDD	5.103	9.882	1.213	2.483
	(11.84)	(11.67)	(7.809)	(8.442)
GDD (F)	37.39		30.79**	
	(34.66)		(14.48)	
HDD	-158.3	-100.1	-191.9	-158.2
	(332.2)	(323.9)	(156.9)	(191.0)
HDD (F)	1,192.2***		-56.71	
	(402.7)		(178.5)	
Mean	155,386.2	155,386.2	123,249.1	123,249.1
Unique Farms	2,603	2,603	2,603	2,603
Fixed-effects				
Farm	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Fit statistics				
Observations	18,917	18,917	18,917	18,917
R^2	0.88715	0.88700	0.93666	0.93662

Notes. Two-way department-by-year standard-errors in parentheses. Stars indicate estimate is significantly different from zero: *$p < .10$, **$p < .05$, ***$p < .01$. Realized and forecasted rainfall in levels and squares are included as controls, as well as quadratic region-specific time trends. Sales correspond to total sales at the farm levels, and costs to total costs.

Table A12: Profit Reaction to Forecasts

Dependent Variables:	Value Added		Profit	
Model:	(1)	(2)	(3)	(4)
Variables				
GDD	2.597	7.932	-8.322	-3.996
	(12.68)	(14.48)	(13.16)	(14.22)
GDD (F)	16.35		3.257	
	(43.10)		(43.14)	
HDD	53.36	117.8	184.6	232.4
	(292.0)	(272.9)	(277.2)	(253.2)
HDD (F)	2,429.3***		2,066.7***	
	(809.1)		(708.1)	
Mean	58,356.8	58,356.8	86,695.2	86,695.2
Unique Farms	2,603	2,603	2,603	2,603
Fixed-effects				
Farm	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Fit statistics				
Observations	18,917	18,917	18,917	18,917
R^2	0.80546	0.80435	0.84140	0.84076

Notes. Two-way department-by-year standard-errors in parentheses. Stars indicate estimate is significantly different from zero: *$p < .10$, **$p < .05$, ***$p < .01$. Realized and forecasted rainfall in levels and squares are included as controls, as well as quadratic region-specific time trends.

Table A13: Input Reactions to Forecasts (1 month lead)

Dependent Variables:	Land	Labor	Fertilizer	Phytosanitary	Seeds	Irrigation
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
GDD	-0.1459	0.0660	-2.648	-1.297	0.3886	0.4377
	(0.2313)	(0.1441)	(2.659)	(1.256)	(1.177)	(0.2677)
GDD (F)	0.0960	0.3955**	7.640	-0.6958	-0.2525	-0.4981
	(0.7421)	(0.1786)	(8.721)	(2.423)	(2.267)	(0.6717)
HDD	5.878	-0.5341	-86.74	-16.09	-17.99	7.506
	(6.055)	(4.215)	(63.42)	(36.68)	(25.75)	(6.729)
HDD (F)	11.10	-0.5139	-9.140	54.64	-32.99	-23.88
	(18.97)	(5.123)	(77.49)	(71.49)	(52.62)	(17.36)
Mean	14,504.1	2,684.0	31,777.2	23,072.2	12,231.1	683.1
Unique Farms	2,603	2,603	2,603	2,603	2,603	2,603
Fixed-effects						
Farm	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
Observations	18,917	18,917	18,917	18,917	18,917	18,917
R^2	0.97649	0.86153	0.89547	0.92414	0.88021	0.86328

Notes: Two-way department-by-year standard-errors in parentheses. Stars indicate estimate is significantly different from zero: *$p < .10$, **$p < .05$, ***$p < .01$. Realized and forecasted rainfall in levels and squares are included as controls, as well as quadratic region-specific time trends.

Table A14: Input Reactions to Forecasts (2 months lead)

Dependent Variables:	Land	Labor	Fertilizer	Phytosanitary	Seeds	Irrigation
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
GDD	-0.0902	0.0964	-2.213	-1.473	0.3063	0.3209
	(0.2358)	(0.1368)	(2.638)	(1.268)	(1.151)	(0.2655)
GDD (F2)	-0.4784	0.6862**	7.582	2.205	-0.7147	0.1433
	(1.269)	(0.3289)	(9.079)	(3.116)	(2.809)	(0.6912)
HDD	4.587	-1.461	-91.57	-7.057	-14.96	10.05
	(4.927)	(3.618)	(73.23)	(42.08)	(25.35)	(6.736)
HDD (F2)	40.92	4.447	-65.32	-10.02	-118.5**	-33.40
	(39.92)	(10.19)	(119.0)	(84.24)	(57.32)	(23.67)
Mean	14,504.1	2,684.0	31,777.2	23,072.2	12,231.1	683.1
Unique Farms	2,603	2,603	2,603	2,603	2,603	2,603
Fixed-effects						
Farm	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
Observations	18,917	18,917	18,917	18,917	18,917	18,917
R^2	0.97649	0.86160	0.89546	0.92412	0.88027	0.86332

Notes: Two-way department-by-year standard-errors in parentheses. Stars indicate estimate is significantly different from zero: $*p < .10$, $**p < .05$, $***p < .01$. Realized and forecasted rainfall in levels and squares are included as controls, as well as quadratic region-specific time trends.

Table A15: Input Prices Reactions to Forecasts (1 month lead)

Dependent Variables:	All Store Prices	Fertilizers	Pesticides	Seeds	Land
Model:	(1)	(2)	(3)	(4)	(5)
Variables					
GDD	-0.1264	-0.1150	-0.3309	-0.0022	-0.9180**
	(0.2577)	(0.1148)	(0.3590)	(0.7071)	(0.3841)
GDD (F)	-0.5127	0.2934*	-1.019	-0.5383	0.3801
	(1.061)	(0.1668)	(1.328)	(2.557)	(0.7375)
HDD	-15.64	-0.9002	-17.44	-15.24	5.278
	(10.31)	(1.969)	(14.07)	(26.53)	(7.407)
HDD (F)	-0.2832	-2.031	3.342	3.638	-1.186
	(11.63)	(3.208)	(21.13)	(20.88)	(12.65)
Mean	1,899.6	517.0	3,151.3	310.8	4,848.3
Unique Farms	298	203	283	195	2,428
Fixed-effects					
Company	Yes	Yes	Yes	Yes	Year
Product	Yes	Yes	Yes	Yes	
Year	Yes	Yes	Yes	Yes	Yes
Farm					Yes
Fit statistics					
Observations	33,174	9,258	17,883	6,033	17,003
R^2	0.35144	0.61626	0.21604	0.37588	0.96108

Notes: Two-way department-by-year standard-errors in parentheses. Stars indicate estimate is significantly different from zero: $*p < .10$, $**p < .05$, $***p < .01$. Realized and forecasted rainfall in levels and squares are included as controls, as well as quadratic region-specific time trends.

Table A16: Input Prices Reactions to Forecasts (2 months lead)

Dependent Variables:	All Store Prices	Fertilizers	Pesticides	Seeds	Land
Model:	(1)	(2)	(3)	(4)	(5)
Variables					
GDD	-0.1425	-0.0728	-0.3654	-0.0289	-0.8877**
	(0.2250)	(0.1214)	(0.3820)	(0.6006)	(0.3727)
GDD (F2)	0.5476	0.5868**	1.177	-0.5384	0.7041
	(0.5782)	(0.2197)	(1.003)	(2.462)	(0.6087)
HDD	-13.84	-1.015	-15.04	-14.44	6.401
	(8.467)	(2.204)	(14.49)	(28.37)	(7.860)
HDD (F2)	-13.95	-3.455	-18.26	-32.05	-6.150
	(19.30)	(4.282)	(34.75)	(32.29)	(14.42)
Mean	1,899.6	517.0	3,151.3	310.8	4,848.3
Unique Farms	298	203	283	195	2,428
Fixed-effects					
Company	Yes	Yes	Yes	Yes	
Product	Yes	Yes	Yes	Yes	
Year	Yes	Yes	Yes	Yes	Yes
Farm					Yes
Fit statistics					
Observations	33,174	9,258	17,883	6,033	17,003
R^2	0.35146	0.61629	0.21611	0.37602	0.96115

Notes: Two-way department-by-year standard-errors in parentheses. Stars indicate estimate is significantly different from zero: *$p < .10$, **$p < .05$, ***$p < .01$. Realized and forecasted rainfall in levels and squares are included as controls, as well as quadratic region-specific time trends.

Table A17: Output Quantities Reactions to Forecasts (1 months lead)

Dependent Variables:	Wheat	Corn	Sunflower	Colza	Beets
Model:	(1)	(2)	(3)	(4)	(5)
Variables					
GDD	-0.5348	0.3959	0.1296	-0.4127**	1.844
	(0.6444)	(0.3648)	(0.1032)	(0.1891)	(2.867)
GDD (F)	0.4925	0.9547	-0.1866	0.5136	-6.308
	(1.095)	(1.260)	(0.2368)	(0.5582)	(4.832)
HDD	1.234	-27.73*	-2.146	9.309**	-24.04
	(8.995)	(14.56)	(2.607)	(3.566)	(58.42)
HDD (F)	61.79***	-6.591	-1.009	13.34	66.38**
	(19.64)	(9.322)	(5.208)	(12.42)	(29.84)
Mean	3,840.1	3,464.3	509.3	914.9	9,127.8
Unique Farms	2,448	1,607	1,357	1,940	346
Fixed-effects					
Farm	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes
Fit statistics					
Observations	17,336	10,346	7,203	12,394	2,267
R^2	0.90300	0.93547	0.78666	0.80582	0.92908

Notes: Two-way department-by-year standard-errors in parentheses. Stars indicate estimate is significantly different from zero: $*p < .10$, $**p < .05$, $***p < .01$. Realized and forecasted rainfall in levels and squares are included as controls, as well as quadratic region-specific time trends.

Table A18: Farm-Level Reactions: Decomposed Weather

Dependent Variables:	Revenue	Costs	Revenue	Costs
Model:	(1)	(2)	(3)	(4)
Variables				
GDD (m)	10.21	14.67	9.516	11.24
	(11.45)	(11.19)	(13.50)	(10.83)
GDD (nm)	6.281	3.595	8.043	1.765
	(11.47)	(9.024)	(10.82)	(8.555)
GDD (m) (F)	21.81	24.53		
	(34.12)	(16.99)		
GDD (m) (F2)			76.42**	27.23
			(34.60)	(17.44)
GDD (nm) (F)	1.870	31.30*		
	(49.83)	(15.75)		
GDD (nm) (F2)			44.16	27.53*
			(39.11)	(15.63)
HDD (m)	-281.3	-225.8	-308.2	-198.1
	(325.7)	(144.8)	(278.8)	(131.9)
HDD (nm)	-483.9	168.1	-218.1	239.3
	(1,735.4)	(808.5)	(1,356.9)	(743.9)
HDD (m) (F)	-1,687.3**	-78.96		
	(735.4)	(934.0)		
HDD (nm) (F)	1,713.4***	-79.26		
	(438.0)	(363.2)		
HDD (m) (F2)			753.2	-72.96
			(758.5)	(384.6)
HDD (nm) (F2)			-641.1	-637.2
			(786.4)	(423.6)
Mean	155,396.1	123,264.8	155,396.1	123,264.8
Unique Farms	2,523	2,523	2,523	2,523
Fixed-effects				
Farm	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Fit statistics				
Observations	17,581	17,581	17,581	17,581
R^2	0.88584	0.93706	0.88592	0.93706

Notes. Two-way department-by-year standard-errors in parentheses. Stars indicate estimate is significantly different from zero: $*p < .10$, $**p < .05$, $***p < .01$. Realized and forecasted rainfall in levels and squares are included as controls, as well as quadratic region-specific time trends. Here (m) indicates marginal weather, (nm) stands for non-marginal, and the F indicates one month ahead seasonal forecasts.

Table A19: Cost and Revenue Reactions to Alternative Weather

Dependent Variables:	Revenue	Costs	Revenue	Costs
Model:	(1)	(2)	(3)	(4)
Variables				
GDD	2.727	-1.260	9.877	-0.2960
	(10.58)	(8.001)	(9.127)	(8.967)
GDD (F)	21.26	35.91**		
	(37.08)	(14.78)		
GDD (F2)			54.92	37.30**
			(37.20)	(17.39)
HDD	-53.39	-106.3	-70.30	-85.43
	(178.4)	(90.81)	(150.2)	(108.3)
HDD (F)	628.3**	-13.57		
	(261.4)	(127.9)		
HDD (F2)			30.40	46.07
			(154.3)	(92.70)
FDD	22.71	-2.684	36.93	-3.235
	(30.00)	(12.43)	(29.83)	(12.01)
FDD (F)	-23.42	-96.00		
	(107.2)	(84.34)		
FDD (F2)			-7.956	-41.77
			(139.4)	(78.21)
Mean	155,637.5	123,304.2	155,637.5	123,304.2
Unique Farms	2,625	2,625	2,625	2,625
Fixed-effects				
Farm	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Fit statistics				
Observations	18,428	18,428	18,428	18,428
R^2	0.88877	0.93766	0.88866	0.93766

Notes. Two-way department-by-year standard-errors in parentheses. Stars indicate estimate is significantly different from zero: $^{*}p < .10$, $^{**}p < .05$, $^{***}p < .01$. Realized and forecasted rainfall in levels and squares are included as controls, as well as quadratic region-specific time trends.

Table A20: Weather Decomposition

	Mean
GDD	2066.18
GDD marginal	1803.06
GDD non-marginal	263.11
GDD - ratio	0.13
HDD	2.03
HDD marginal	0.25
HDD non-marginal	1.78
HDD - ratio	0.85

Notes. We specify the ratio of marginal to total GDD and HDD realization along with their respective average over all French departments for 1994-2018.

Table A21: Persistence in Heating Degree Days Forecasts

Dependent Variables:	FHDD (lag 2)	FHDD (lag 1)	FHDD	FHDD (lead 1)	FHDD (lead 2)	FHDD (lead 3)
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
GDD	-0.0024	-0.0018	1.56×10^{-17}	-1.61×10^{-5}	0.0016	-0.0005
	(0.0018)	(0.0022)	(4.29×10^{-10})	(0.0020)	(0.0025)	(0.0025)
GDD (F)	-0.0007	-0.0019	2.26×10^{-17}	0.0008	0.0058	0.0020
	(0.0045)	(0.0035)	(1.17×10^{-9})	(0.0052)	(0.0067)	(0.0035)
GDD (lag)	-0.0022	0.0009	-3.73×10^{-17}	0.0001	0.0001	-0.0002
	(0.0014)	(0.0016)	(2.1×10^{-10})	(0.0013)	(0.0012)	(0.0015)
GDD (lag 2)	0.0008	-0.0001	-4.94×10^{-17}	0.0004	-0.0006	0.0017
	(0.0015)	(0.0009)	(6.3×10^{-11})	(0.0011)	(0.0010)	(0.0021)
HDD	-0.0289	-0.0033	-1.27×10^{-15}	0.0496	0.0014	0.0530
	(0.0527)	(0.0257)	(1.5×10^{-8})	(0.0676)	(0.0457)	(0.0981)
HDD (F)	-0.1117	-0.1060	1.000***	-0.1428	-0.1850	-0.2279
	(0.1381)	(0.1344)	(1.75×10^{-8})	(0.1350)	(0.1083)	(0.3330)
HDD (lag)	-0.0014	0.0134	-9.05×10^{-16}	0.0222	0.0303	0.0085
	(0.0382)	(0.0495)	(1.71×10^{-8})	(0.0367)	(0.0841)	(0.0454)
HDD (lag 2)	-0.0060	0.0265	5.2×10^{-16}	0.0217	0.0030	-0.0326
	(0.0518)	(0.0325)	(1.87×10^{-8})	(0.0855)	(0.0419)	(0.0564)
Fixed-effects						
Farm	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
Observations	14,097	14,097	14,097	12,193	10,504	9,019
R^2	0.68571	0.73456	1.0000	0.71000	0.69887	0.70215

Clustered (Department & Year) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes. Two-way department-by-year standard-errors in parentheses. Stars indicate estimate is significantly different from zero: $^*p < .10$, $^{**}p < .05$, $^{***}p < .01$. Realized and forecasted rainfall in levels and squares are included as controls, as well as the one-period and two-period lag values for rainfall in squares and levels. We finally include quadratic region-specific time trends.

Table A22: Dynamic Effects on Profit

Dependent Variables:	Profit (lag 2)	Profit (lag)	Profit	Profit (lead)	Profit (lead 2)	Profit (lead 3)
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
GDD	-37.50	3.040	0.8610	1.672	6.464	-12.11
	(26.39)	(18.60)	(15.60)	(21.22)	(28.85)	(15.23)
GDD (F)	-28.51	19.16	22.59	35.83	-23.00	74.23
	(44.17)	(41.41)	(51.62)	(31.77)	(65.72)	(68.71)
GDD (lag)	-20.16	23.08	-8.886	-11.17	24.54	-19.38
	(13.49)	(20.02)	(10.79)	(12.80)	(15.51)	(23.39)
GDD (lag 2)	33.05	-1.883	-19.26*	17.50	7.012	-4.012
	(20.04)	(10.61)	(10.90)	(11.41)	(22.95)	(23.43)
HDD	85.82	-645.9	110.3	343.7	294.6	-271.3
	(351.0)	(399.8)	(338.7)	(271.8)	(539.8)	(536.6)
HDD (F)	663.2	265.1	1,908.2**	-1,827.7***	-1,435.3*	-323.6
	(450.4)	(649.2)	(847.0)	(645.3)	(701.9)	(699.6)
HDD (lag)	-628.0	-383.3	523.3**	747.2**	-301.0	1,425.3*
	(388.2)	(500.4)	(233.5)	(312.6)	(533.0)	(702.8)
HDD (lag 2)	-810.0	545.8**	709.9**	-239.8	1,350.0*	196.2
	(528.3)	(224.8)	(254.8)	(407.8)	(773.5)	(732.5)
Mean	89,683.8	87,948.4	86,695.2	86,835.8	86,755.7	86,847.5
Unique Farms	1,904	1,904	1,904	1,689	1,485	1,308
Fixed-effects						
Farm	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
Observations	14,097	14,097	14,097	12,193	10,504	9,019
R^2	0.83711	0.84295	0.84166	0.83855	0.83286	0.82522

Clustered (Department & Year) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes. Two-way department-by-year standard-errors in parentheses. Stars indicate estimate is significantly different from zero: $*p < .10$, $**p < .05$, $***p < .01$. Realized and forecasted rainfall in levels and squares for period t are included as controls, as well as lag realized rainfall for periods $t-1$ and $t-2$, in both levels and squares, and finally quadratic region-specific time trends.

Table A23: Dynamic Effects on Costs

Dependent Variables:	Costs (lag 2)	Costs (lag)	Costs	Costs (lead)	Costs (lead 2)	Costs (lead 3)
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
GDD	-3.575	-6.160	-1.821	9.009	-1.588	-0.6275
	(2.936)	(7.389)	(6.765)	(7.430)	(6.439)	(15.26)
GDD (F)	4.233	10.76	25.16	1.032	19.52	0.1542
	(12.40)	(15.83)	(16.93)	(15.47)	(13.90)	(13.35)
GDD (lag)	-2.732	-1.792	8.348**	-1.979	-7.641	7.856
	(5.282)	(5.148)	(3.891)	(3.211)	(5.514)	(8.929)
GDD (lag 2)	-7.519	1.596	-1.100	-6.294*	3.882	-1.251
	(5.382)	(3.260)	(3.044)	(3.288)	(7.431)	(10.67)
HDD	156.4	85.26	-283.6	-213.1	-20.03	164.1
	(103.4)	(93.48)	(175.5)	(223.0)	(178.1)	(227.9)
HDD (F)	-173.2	-188.6	-196.9	1,028.5**	325.9	118.3
	(303.7)	(215.1)	(115.6)	(429.7)	(409.7)	(222.7)
	(126.7)	(115.7)	(117.7)	(144.6)	(301.2)	(185.8)
HDD (lag)	1.882	-243.4	-88.25	124.1	218.0	-152.4
	(113.6)	(157.4)	(174.8)	(143.9)	(221.5)	(314.4)
HDD (lag 2)	-176.8	-61.66	56.04	171.7	-117.3	212.2
Unique Farms	1,904	1,904	1,904	1,689	1,485	1,308
Fixed-effects						
Farm	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
Observations	14,097	14,097	14,097	12,193	10,504	9,019
R^2	0.93847	0.94030	0.94043	0.94116	0.94228	0.94544

Clustered (Department & Year) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes. Two-way department-by-year standard-errors in parentheses. Stars indicate estimate is significantly different from zero: $^{*}p < .10$, $^{**}p < .05$, $^{***}p < .01$. Realized and forecasted rainfall in levels and squares for period t are included as controls, as well as lag realized rainfall for periods $t-1$ and $t-2$, in both levels and squares, and finally quadratic region-specific time trends.

B Data Details

B.1 Definition of Variables

Definitions for the main variables used:

- **Crop Prices.** They are measured by dividing the total value of sales of that given crop, by the total quantity sold.
- **Output Price Index.** For l_{jct} the land area allocated to crop c by farm j in period t, $\mathbb{C}_{jt}$ the crop mix of farm j in period t, and p_{jct} the output price of that same crop for that same farm, we build:

$$p_{jt} = \sum_{c \in \mathbb{C}_{jt}} \frac{l_{jct}}{\sum_{c \in \mathbb{C}_{jt}} l_{jct}} p_{jct} \tag{5}$$

  We consider the following crops for that purpose: wheat, durum wheat, oats, corn, corn (seeds), sorghum, spring barley, winter barley, rye, triticale, summer cereals, other cereals, sunflower, colza, soy, dry peas, feverole beans, protein peas.

- **Output/Storage Quantity Index.** This index is used as an aggregate measure of farm output/storage, and is intended to represent the average level of output/storage across the farm rather than a total quantity. It is built for the same set of crops as the one used for the output price index. We use the analogous formula: with q_{jct} the output/storage quantity for crop c in farm j in year t:

$$q_{jt} = \sum_{c \in \mathbb{C}_{jt}} \frac{l_{jct}}{\sum_{c \in \mathbb{C}_{jt}} l_{jct}} q_{jct} \tag{6}$$

- **Storage.** We define the variation in storage at period t - the net storage flow - as the difference between the quantity produced and the quantity sold for a specific crop.
- **Land prices.** are defined as the total value of land divided by the total quantity of land.
- **Fertilizer, pesticide and seed.** They are observed at the farm-level, correspond to deflated bills, and are defined as the difference between purchases plus beginning-of-period stocks, minus end-of-period stocks.
- **Labor.** Defined as the total number of paid hours worked over the season.
- **Intermediary Inputs.** Defined as a deflated bill. The sum of expenses for: fertilizer, seeds, pesticides, animal food, veterinary products, products for animal reproduction, packaging, fuel, maintenance products, supplies, food for workers, raw materials, purchases of services for cultivation, breeding or others, water, gas, electricity, irrigation water, lease installments, material rental, animal rental, maintenance for buildings, lands and material, studies and research, veterinary services, communication and commercials, transportation costs, travel costs, postal services, banking services, other services and costs.
- **Total Costs.** Defined as a deflated bill. Corresponds to the sum of intermediary inputs, social contributions for workers, personnel expenses, taxes, insurance.
- **Value Added.** Defined as the difference between total production value minus animal purchases, minus intermediary inputs. Total production itself corresponds to total production sold, self-consumed production, immobilised production, stored production, gains from animal boarding, land rentals, other rentals, agricultural tourism.

- **Profit** Defined as the gross operating income of the farm. Our profit variable encompasses value added, and among else also includes subsidies, expenses for insurance, and insurance indemnities.

All variables measured in euros are converted into 2020 euros using the INSEE consumer price index.

B.2 Weather Data

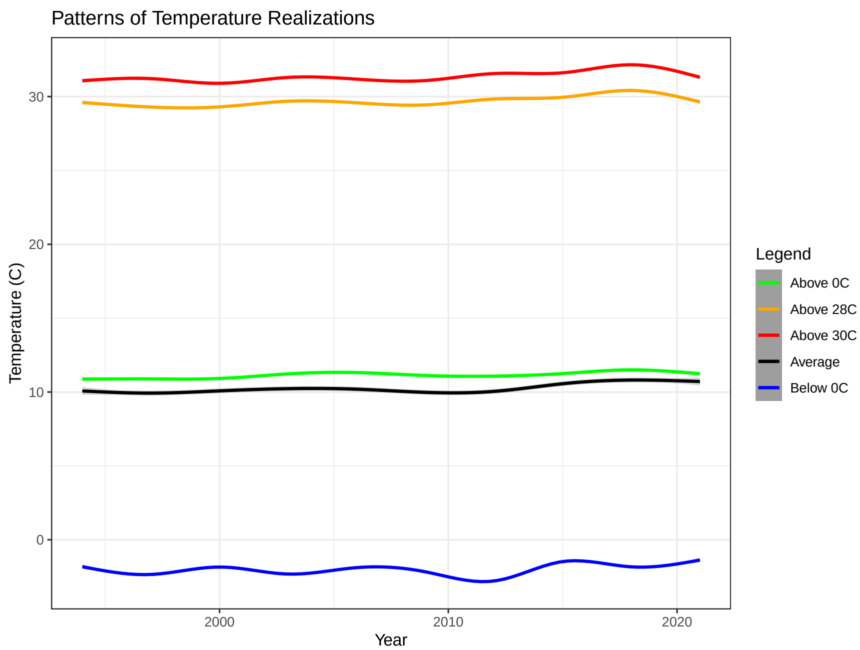

Figure A20: Unconditional and Conditional Average Temperature Realizations

C Validation of Results

Table A24: Corn-Specific Outcomes

Dependent Variables:	Output	Price	Sales	Sales Quantity
Model:	(1)	(2)	(3)	(4)
Variables				
GDD	0.4668	-0.0041	8.820	0.6642
	(0.3725)	(0.0094)	(8.334)	(0.5162)
HDD	-30.72*	0.7290**	-133.8	-13.71
	(15.22)	(0.2970)	(186.0)	(11.39)
Mean	3,464.3	150.0	53,507.8	3,437.5
Unique Farms	1,607	1,581	1,607	1,607
Fixed-effects				
Farm	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Fit statistics				
Observations	10,346	9,958	10,346	10,346
R^2	0.93538	0.72917	0.87702	0.87607

Notes. Two-way department-by-year standard-errors in parentheses. Stars indicate estimate is significantly different from zero: $*p < .10$, $**p < .05$, $***p < .01$. Realized and forecasted rainfall in levels and squares are included as controls, as well as quadratic region-specific time trends.

Table A25: Farm-Level Outcomes

Dependent Variables:	Revenue	Costs	Intermediary Inputs	Value Added	Profit	Output Price Index
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
GDD	9.882	2.483	-3.785	7.932	-3.996	-0.1557
	(11.67)	(8.442)	(5.608)	(14.48)	(14.22)	(0.1172)
HDD	-100.1	-158.2	-126.7	117.8	232.4	0.2812
	(323.9)	(191.0)	(160.6)	(272.9)	(253.2)	(1.716)
Mean	155,386.2	123,249.1	103,749.6	58,356.8	86,695.2	263.0
Unique Farms	2,603	2,603	2,603	2,603	2,603	2,603
Fixed-effects						
Farm	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
Observations	18,917	18,917	18,917	18,917	18,917	18,917
R^2	0.88700	0.93662	0.93391	0.80435	0.84076	0.70251

Notes. Two-way department-by-year standard-errors in parentheses. Stars indicate estimate is significantly different from zero: $*p < .10$, $**p < .05$, $***p < .01$. Realized and forecasted rainfall in levels and squares are included as controls, as well as quadratic region-specific time trends.

D Robustness

D.1 Heterogeneous Response to Forecasts

We show that there is little heterogeneity in farms' response to forecasted HDDs, looking at two potential important margins. First, we show that farms of different size do not seem to respond in very different ways to forecasts, then that farmers from different generations seem to have also a similar response to forecasts.

In the following graph we show the coefficient associated to the one month ahead HDD forecast when regressing farm log profit on realized and forecasted weather, and including year and farm fixed effects. We measure farm size using their gross operating income deflated to 2020 euros. While the graph shows some variation across quantiles and showcases a form of U-shape, there are no large differences in coefficient values across these quantiles. Farms are able to save the same percentage of their profit by using forecasted HDDs. This implies however that larger farmers are able to save a larger value whe using forecasts, hinting that heat shocks act as a negative multiplier to production rather than an additive shock.

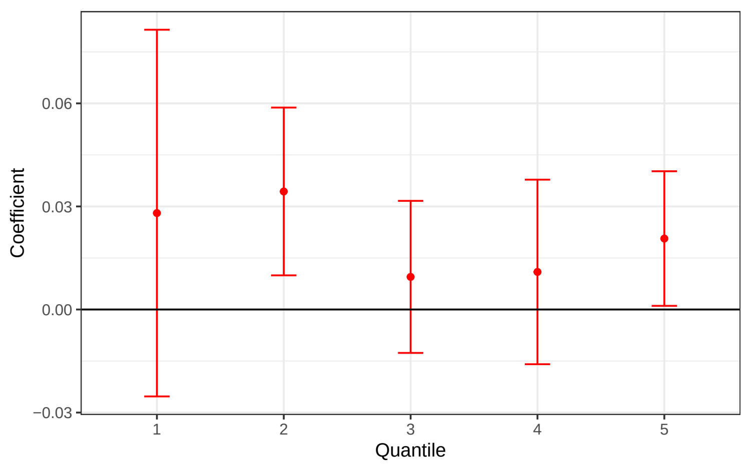

Figure A21: Varying Response to Fixed Lead of HDD

Next we rank farms by the date of birth of their manager. This analysis relies on the fact that forecasts are a relatively new technology, dating to the 1990s, and might be more easily adopted by farmers who went through their education when these were already available.

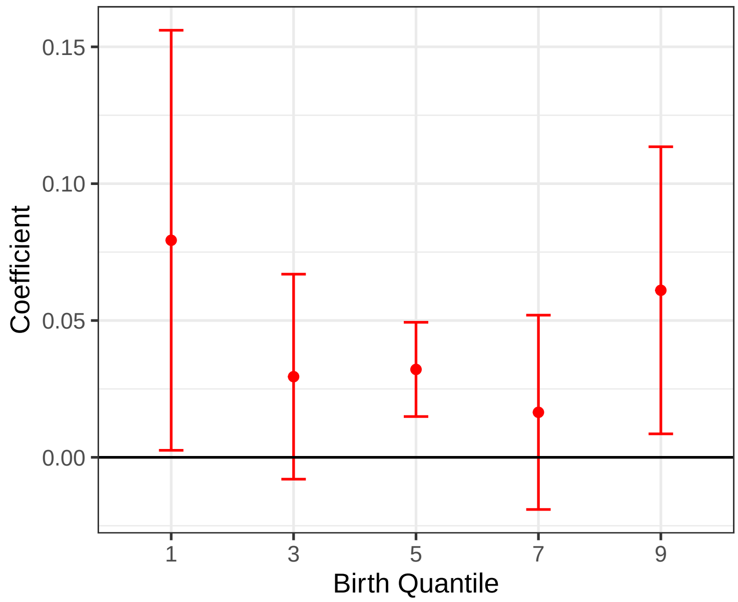

Figure A22: Profit Response to Forecasts across Manager's Age

We see a small increase in the reaction to forecasts for the youngest group of farmers, but again the difference is not stark. This is indicative that the timing of the farmers' education is not a strong predictor for their adoption of forecasts.

D.2 Including Lags

As a robustness check, we run the same regression as before but including lagged values of weather realization (rainfall in levels and squared, GDD and HDD). Again, the regressions also include farm and year fixed effects, and we cluster standard errors at the department level. Lagged realizations might matter, first because of auto-correlation in weather, for example due to the role of large weather patterns such as the North Atlantic Oscillation. Second, lags might also play a role in the setting of farmers' beliefs about the upcoming weather, as modelled and discussed by Burke and Emerick (2016). In this case, including both lags and forecasts might better account for farmers' beliefs.

As we see however, lagged realizations of growing and heating degree days play a non-significant role in driving farm profit once we control for forecasts. Forecasts of heating degree days, on the other hand, still have a large positive and significant impact on farm profit.

Table A26: Farm-Level Profit - Lead 1 with Lags

Dependent Variables:	Value Added	Profit	Value Added	Profit
Model:	(1)	(2)	(3)	(4)
Variables				
GDD	7.832	-2.904	16.74	3.510
	(14.84)	(14.43)	(16.09)	(15.45)
GDD (lag)	-3.248	-7.945	-5.777	-9.417
	(9.944)	(9.856)	(9.745)	(9.903)
GDD (F)	24.77	9.606		
	(47.09)	(45.98)		
GDD (F2)			74.09	48.48
			(48.83)	(49.66)
HDD	-22.75	133.9	-111.1	61.37
	(333.0)	(330.6)	(290.2)	(280.3)
HDD (lag)	352.6*	429.4*	390.0*	448.4*
	(200.6)	(218.6)	(220.8)	(229.6)
HDD (F)	2,202.0**	1,781.0**		
	(848.8)	(761.8)		
HDD (F2)			1,058.3	1,097.9
			(720.1)	(644.3)
Mean	58,356.8	86,695.2	58,356.8	86,695.2
Unique Farms	2,217	2,217	2,217	2,217
Fixed-effects				
Farm	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Fit statistics				
Observations	16,314	16,314	16,314	16,314
R^2	0.80934	0.84251	0.80899	0.84235

Notes. Two-way department-by-year standard-errors in parentheses. Stars indicate estimate is significantly different from zero: $*p < .10$, $**p < .05$, $***p < .01$. Realized and forecasted rainfall in levels and squares are included as controls, as well as quadratic region-specific time trends.

Compared to our main results in Table 5, we see a slight decrease in the coefficients associated to forecasted HDDs: the coefficient for value added moves from 2,430 to 2,199, and the one for profit moves from 2,005 to 1,712. As such, results are stable to the inclusion or exclusion of lags.

D.3 Alternative Weather Aggregation

We perform another robustness check, and recompute our growing and heating degree days. This time, we use 28°C as a cutoff for the classification of hourly temperature realizations as GDD or HDD. The hours spent below are now counted towards growing degree days, while those above count towards the heating degree days. We also create a measure of freezing degree days, which counts the degree-hours spent below

0°C in absolute value. This might be useful, first to identify whether we can observe interesting responses to freeze, and see whether we also observe non-marginal profit responses to their forecast. Second, low growing degree day values might correlate with abnormally low temperature, and without including freezing degree days, might capture in part the impact of freeze on agriculture. Including them will then purge GDD from its correlation with very cold events.

Table A27: Farm-Level Profit - Alternative Weather

Dependent Variables:	Value Added	Profit	Value Added	Profit
Model:	(1)	(2)	(3)	(4)
Variables				
GDD	6.967	-4.105	16.23	2.994
	(11.16)	(11.30)	(13.57)	(13.14)
GDD (F)	-12.41	-20.26		
	(44.68)	(44.57)		
GDD (F2)			54.21	33.22
			(42.57)	(42.86)
HDD	61.33	129.3	-1.853	70.28
	(177.4)	(170.9)	(154.6)	(143.3)
HDD (F)	942.8***	796.6***		
	(295.8)	(269.4)		
HDD (F2)			151.5	125.0
			(247.9)	(238.2)
FDD	68.19*	64.41	81.91**	76.22*
	(39.12)	(37.66)	(39.03)	(37.38)
FDD (F)	188.8	141.9		
	(158.4)	(155.5)		
FDD (F2)			139.2	175.6
			(124.0)	(133.9)
Mean	58,332.2	87,330.0	58,332.2	87,330.0
Unique Farms	2,625	2,625	2,625	2,625
Fixed-effects				
Farm	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Fit statistics				
Observations	18,428	18,428	18,428	18,428
R^2	0.80601	0.84326	0.80505	0.84269

Notes. Two-way department-by-year standard-errors in parentheses. Stars indicate estimate is significantly different from zero: *$p < .10$, **$p < .05$, ***$p < .01$. Realized and forecasted rainfall in levels and squares are included as controls, as well as quadratic region-specific time trends.

The results are similar to the ones from our main specification. Forecasted HDDs have a positive impact of profit, albeit a smaller one than previously. A one unit increase in forecasted HDD corresponds here to a smaller increase in temperature over the growing season, and it should be expected that it implies a smaller response. Forecasted freezing days also imply a positive profit response. As such, it seems that farmers adapt both to extremely hot and cold events, in a way that leads to non-marginal changes in their optimal profit.

In Table A19 we show the costs and revenue responses to these alternative weather calculations.

D.4 Removing Time Trends

Here we run the same regressions as in the main part of the paper, using our initial measures of weather, but removing the quadratic department-specific time trends. Results are very close to those in Table 2.

Table A28: Farm-Level Profit - No Time Trend

Dependent Variables:	Sales	Costs	Sales	Costs
Model:	(1)	(2)	(3)	(4)
Variables				
GDD	3.473	-1.273	8.786	0.8753
	(10.69)	(7.359)	(10.81)	(7.934)
GDD (F)	47.76	37.66**		
	(34.42)	(14.38)		
GDD (F2)			68.20*	45.50**
			(37.10)	(17.93)
HDD	-211.9	-184.3	-263.1	-183.4
	(303.3)	(148.6)	(309.8)	(187.2)
HDD (F)	1,342.2***	39.32		
	(454.7)	(270.2)		
HDD (F2)			1,381.6*	156.4
			(702.9)	(437.3)
Mean	155,386.2	123,249.1	155,386.2	123,249.1
Unique Farms	2,603	2,603	2,603	2,603
Fixed-effects				
Farm	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Fit statistics				
Observations	18,917	18,917	18,917	18,917
R^2	0.88667	0.93591	0.88671	0.93591

Notes. Two-way department-by-year standard-errors in parentheses. Stars indicate estimate is significantly different from zero: $*p < .10$, $**p < .05$, $***p < .01$. Realized and forecasted rainfall in levels and squares are included as controls, as well as quadratic region-specific time trends.

D.5 Disaggregated Realized Weather

Below we run our main specification using village-level realized weather and department-level forecasted weather, which is the least coarse geographic unit that can match the forecasts' grid. We note there are about 36,000 villages in France, and hence this disaggregation corresponds to a significant gain in precision. This analysis is performed on a subset of our main sample from 2002 onwards, as we cannot locate farms at a finer level than the department prior to 2002.

Table A29: Cost and Revenue Reactions to Disaggregated Realized Weather

Dependent Variables:	Revenue	Costs	Revenue	Costs
Model:	(1)	(2)	(3)	(4)
Variables				
GDD	1.505	-2.433	1.505	-2.433
	(16.72)	(5.363)	(16.72)	(5.363)
GDD (F)	42.86	52.73**	42.86	52.73**
	(45.68)	(22.50)	(45.68)	(22.50)
HDD	43.12	39.96	43.12	39.96
	(328.3)	(114.3)	(328.3)	(114.3)
HDD (F)	1,289.7***	52.97	1,289.7***	52.97
	(256.1)	(37.98)	(256.1)	(37.98)
Mean	164,532.2	131,969.1	164,532.2	131,969.1
Unique Farms	1,890	1,890	1,890	1,890
Fixed-effects				
Farm	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Fit statistics				
Observations	12,754	12,754	12,754	12,754
R^2	0.88834	0.94130	0.88834	0.94130

Notes. Two-way department-by-year standard-errors in parentheses. Stars indicate estimate is significantly different from zero: *$p < .10$, **$p < .05$, ***$p < .01$. Realized and forecasted rainfall in levels and squares are included as controls, as well as quadratic region-specific time trends. Realized weather is measure at the village level, while forecasts remain measured at the department one.

D.6 Decomposing along the Forecast Error Sign

Table A30: Cost and Revenue Responses - Heterogeneity along the sign of Forecasts Errors: Forecasts - Realization (1 month forecasts)

Dependent Variables:	Sales	Costs	Profit	Value Added
Model:	(1)	(2)	(3)	(4)
Variables				
GDD	7.588	2.326	-3.975	6.732
	(10.51)	(7.555)	(13.24)	(12.82)
GDD (F)	37.05	28.63**	3.802	18.25
	(33.90)	(12.78)	(45.14)	(45.11)
HDD	-84.49	-184.6	321.9	179.2
	(392.1)	(192.9)	(291.1)	(297.7)
Forecast Bin 1	-3,509.7**	-1,805.6	-4,641.8**	-4,290.6*
	(1,637.1)	(1,158.6)	(2,093.7)	(2,208.9)
Forecast Bin 2	-3,377.8*	-1,169.5	-3,734.6**	-3,962.1**
	(1,894.8)	(751.5)	(1,672.6)	(1,721.9)
Forecast Bin 4	895.0	-897.5	612.8	340.2
	(2,731.0)	(863.7)	(2,569.2)	(2,601.4)
Forecast Bin 5	674.7	-1,327.2	999.5	1,754.9
	(4,605.3)	(1,353.1)	(2,958.8)	(3,183.8)
Mean	155,386.2	123,249.1	86,695.2	58,356.8
Unique Farms	2,603	2,603	2,603	2,603
Fixed-effects				
Farm	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Fit statistics				
Observations	18,917	18,917	18,917	18,917
R^2	0.88717	0.93668	0.84118	0.80498

Notes. Two-way department-by-year standard-errors in parentheses. Stars indicate estimate is significantly different from zero: $*p < .10$, $**p < .05$, $***p < .01$. Realized and forecasted rainfall in levels and squares are included as controls, as well as quadratic region-specific time trends.