

Costs of Climate Adaptation: Evidence From French Agriculture*

Tristan Du Puy and Jeffrey G. Shrader

¹Columbia University

February 29, 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 to a changing climate, 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. We find that for the average farm in France, the cost of adaptation to heat shocks has been low. In contrast, the benefits of adaptation have been 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 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. Finally, we show that the future scope for low-cost adaptation is limited.

*We are grateful for comments from numerous participants in seminars at Columbia University. 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. 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 effect of climate change on the economy. One of the long-standing challenges to this work is estimating the benefits and costs of adaptation—the actions people take to prepare for or adjust to a changing climate. As argued by [Carleton et al. \(2022\)](#) and [Kelly et al. \(2005\)](#), policymakers should be interested in the sum composed of the damages of climate change and the costs of adaptation because a policy to mitigate climate change would not only prevent damage but would 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 adaptation costs 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 economist 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 inputs and input prices needed to estimate adaptation costs are rarely observed. And this lack of cost data is not limited to the agricultural setting. Recent work on temperature-related mortality estimates adaptation costs indirectly by using estimates of adaptation benefits to bound unobservable adaptation costs ([Carleton et al., 2022](#)).

We overcome this challenge by using a uniquely rich panel dataset on French agriculture from 1994-2018. The dataset contains directly observed, precise cost and revenue data at the farm-level. We observe the data for a representative set of farms repeatedly surveyed across multiple growing seasons. The data allow us to recover independent estimates of the marginal costs and benefits of adaptation to weather shocks.

More specifically, we merge the agricultural data with data on forecasts and realizations temperature. We focus on measures of growing season growing degree days (GDDs) and heating degree days (HDDs) at each farms' locations. GDDs provide a measure of exposure to moderate temperatures, while HDDs measure exposure to extremely hot temperatures. Empirically, we control for realized weather and use the forecasts to isolate the farmers' responses to weather-related information shocks. These information shocks cleanly identify the effect of changes in behavior by the farms in anticipation of upcoming weather events and therefore allow us to estimate the costs and benefits of their ex ante adaptation to weather. The farm-level panel data also allows us to include farm-level 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](#)).

The results show that farmers are able to leverage heat forecasts to generate relatively large revenue gains at small costs. 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. Again making use of the rich farm-level data, we show that adaptation to heat shocks in France has mainly taken the form of ex ante timing changes affecting different steps in the growing season. As such, French farmers have so far been able to implement low cost/high gain adaptation strategies in response to heat shocks.

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 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 this breach of the 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 GDDs, and net-positive profit responses to forecasted 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 changes. These timing changes are likely to drive non-continuous changes in farm-level profits, and are hence less suited to the assumptions of usual envelope theorems.

In a third step, we investigate whether the net positive profit responses to forecasted HDDs are likely to be good indicators of future costs of adaptation 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 first 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.

With these results, we can conclude our paper by saying that—so far—French farmers have been skillful in their use of weather forecasts. They have leveraged GDD forecasts to increase their scale of production, and HDD forecasts to generate net positive profit gains through the shifting of major production events across the growing season. These good times are unlikely to last, however. In the subset of worst HDD realizations over 1994-2018, we see that farms lose their ability to turn forecasts into profit gains. Instead, they start leveraging heat forecasts in order to downscale their production and reduce their production costs.

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

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

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

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 modelling 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 than 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

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

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 structural methods to disentangle adaptation effects from direct weather effects, but their method requires a complete enumeration and modelling 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 [section 6](#) examines limits to the adaptation found in the prior section. Finally, [section 7](#) 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 [Schauburger et al. \(2022\)](#). France also possesses the largest agricultural area within the EU, and is the third 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

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 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 further compute unconditional and conditional growing season mean temperature realizations at the department level, for our period of study. Figure A11 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.

⁴We 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).

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.⁵

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.

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.⁶ Each series is defined according to the nature of the product sold, to its sales condition and the firm that sells it.⁷ We localise the sale points at the department level, and match them with department-level weather data. We also use the Laspeyres price indices derived by the INSEE from the EPCIA in order to deflate the FADN farm-level input bills for seeds, fertilizers

⁵The 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.

⁶The relative number of goods sampled within a category is proportional to that category’s relative sales.

⁷The spread of one good’s series across firms is proportional to the firms’ market share for the sale of that specific good.

and pesticides. Specifically we use the Ipampa price index series, from 1994 to 2020. 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 ERA5 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 $[4^{\circ}, 30^{\circ}]$ 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. Growing-season precipitation is measured by summing precipitation over each day.

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 A9](#), 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.⁸ 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 SEAS5 Seasonal 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).⁹ For temperature, forecasts are produced for each 6 hour interval over the forecast horizon. They give an instantaneous point prediction of temperature, while for rainfall, they 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 exceeds 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

⁸[Gammans et al. \(2017\)](#) also define the winter wheat growing season for France as spanning October to July.

⁹ECMWF provides an ensemble of 25 forecasts, which we average.

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 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 Precision

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. 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.¹⁰ 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.

3 Empirical Strategy

To identify the cost and revenue changes with respect to expected weather changes, 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.

$$y_{jt} = \beta_1^w \text{GDD}_{d(j)t} + \beta_2^w \text{HDD}_{d(j)t} + \beta_1^f \text{FGDD}_{d(j)t} + \beta_2^f \text{FHDD}_{d(j)t} + g(P_{d(j)t}) + g_2(FP_{d(j)t}) + \gamma_j + \eta_t + \zeta_{d(j)}^1 t + \zeta_{d(j)}^2 t^2 + \varepsilon_{jt} \quad (1)$$

¹⁰As 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.

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 and FHDD) experienced by farms located in department d during the growing season. Given that weather and forecasts vary at the department level, we cluster our standard errors at that level as well.

We focus on the effects of forecasts 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 input choice 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 department specific time trends in the form of $\zeta_{d(j)}^1 t$ and $\zeta_{d(j)}^2 t^2$. These account for potential sub-national trends that would correlate with weather and our outcomes of interest.¹¹ 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 department-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

¹¹Recent 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.

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.¹² This 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, or to modify decisions such as the timing of harvest. 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.

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 a 7% change in sales and costs.

¹²We 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.

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

Dependent Variables:	Sales	Costs
Model:	(1)	(2)
<i>Variables</i>		
GDD	4.940 (9.241)	-5.687 (5.675)
GDD (F)	42.33** (16.83)	25.83*** (9.283)
HDD	-236.4 (243.8)	-135.2 (126.3)
HDD (F)	1,268.6** (566.7)	83.56 (158.3)
Mean	155,386.2	123,249.1
<i>Fixed-effects</i>		
Farm	Yes	Yes
Year	Yes	Yes
<i>Fit statistics</i>		
Observations	18,917	18,917
R ²	0.88872	0.93855

Notes: Estimates are based on [Equation 1](#), using the baseline sample. Realized and forecasted rainfall in levels and squares, as well as quadratic department-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. One-way department-level standard-errors in parentheses. Stars indicate estimate is significantly different from zero: * $p < .10$, ** $p < .05$, *** $p < .01$.

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-hour in HDD forecasted one month in advance leads to an increase in sales of €1,268. 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.¹³

¹³[Table A7](#) compares results to one-month and two-months ahead forecasts. [Figure A10](#) compares the results for

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.¹⁴ 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 A13](#). 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 A11](#) results for a larger set of inputs.¹⁵ Store prices are unresponsive to HDD forecasts, as are inputs. The positive response of fertilizers to forecasted GDDs follows from the fact that crop nutrient uptake is higher in good growing conditions ([Sheriff, 2005](#)), and acts as one sensibility test for our results. It also tracks with the previously observed increase in costs following a positive GDD forecast. 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 to forecasted heat shocks.

all available leads.

¹⁴Fertilizers 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. 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.

¹⁵[Table A12](#) shows similar results for two-month ahead forecasts.

Table 3: Prices and Quantities for Outputs and Inputs (1 month lead)

Dependent Variables: Model:	Output Price (1)	Output (2)	Storage (3)	Input Price (4)	Irrigation (5)	Fertilizer (6)	PhytoSanitary (7)
<i>Variables</i>							
GDD	-0.1755** (0.0838)	-0.2196 (0.4644)	-0.9748** (0.4070)	-0.1484 (0.1836)	0.3178 (0.2582)	-2.380 (1.511)	-1.826** (0.8668)
GDD (F)	0.3076 (0.2462)	1.561* (0.8953)	-0.7268 (0.9964)	-0.4966 (0.7285)	-0.4195 (0.6964)	12.98*** (3.216)	-0.9860 (1.947)
HDD	0.0208 (1.313)	-13.81 (14.17)	-2.921 (9.653)	-18.90* (9.794)	13.73 (8.393)	-96.00** (37.16)	-18.40 (25.56)
HDD (F)	-0.9513 (2.096)	71.81*** (19.04)	45.19** (21.50)	-0.4188 (8.421)	-6.651 (10.34)	-44.57 (72.35)	12.56 (55.83)
Mean	263.0	7,793.0	88.65	1,899.6	683.1	31,777.2	23,072.2
<i>Fixed-effects</i>							
Farm	Yes	Yes	Yes		Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Company				Yes			
Product				Yes			
<i>Fit statistics</i>							
Observations	18,917	18,917	18,917	33,174	18,917	18,917	18,917
R ²	0.71819	0.94480	0.10975	0.35655	0.87414	0.89992	0.92675

Notes: One-way department-level 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 department-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.

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 A11](#). 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, and an output price index, also outlined in [Table 3](#).¹⁶ The quantity produced exhibits a similar response to forecasted GDDs and HDDs as that of total revenue.

¹⁶We show in [Table A15](#) the crop-specific responses of output quantity to temperature realizations and forecasts.

Forecasted GDDs slightly raise output—and given the null response of output prices—explains the positive revenue response to forecasted GDDs. Forecasted HDDs significantly impact output quantity. A one unit increase in forecasted HDD raises output by 72 cwt, or about 1% 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 increases in response to forecasted HDDs (by more than 60%). This storage effect indicates that the revenue results in [Table 2](#) understate the total effect on output because some of the current period output is directed to storage rather than sold.

Fourth, 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. 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](#).

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:

$$y_{jt} = \beta_1^w \text{GDD}_{d(j)t} + \beta_2^w \text{HDD}_{d(j)t} + \sum_{\ell=1}^5 \left(\beta_{1,\ell}^f \text{FGDD}_{d(j)t}^\ell + \beta_{2,\ell}^f \text{FHDD}_{d(j)t}^\ell + g_{2,\ell}(FP_{d(j)t}^\ell) \right) + \quad (2)$$

$$g(P_{d(j)t}) + \zeta_{d(j)}^1 t + \zeta_{d(j)}^2 t^2 + \gamma_j + \eta_t + \varepsilon_{jt}$$

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

¹⁷In 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.

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

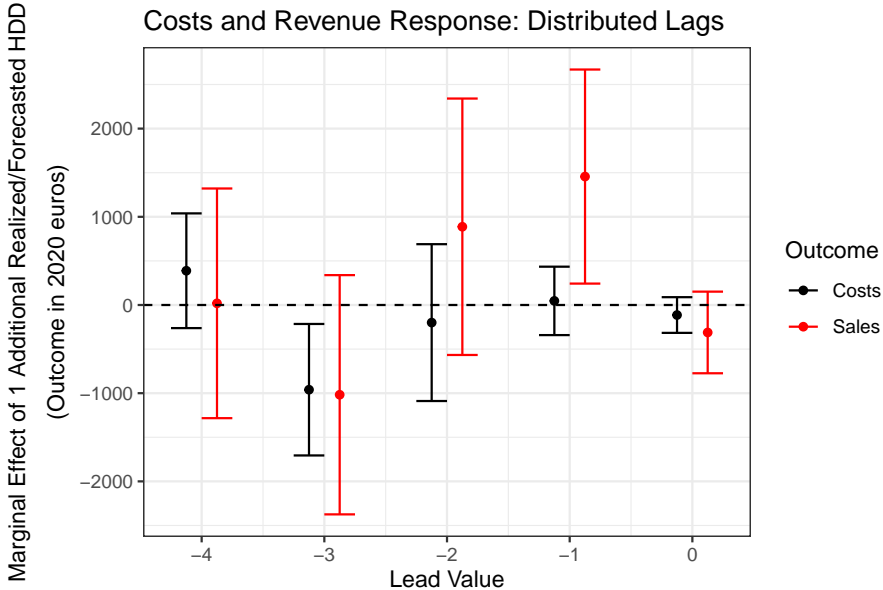


Figure 1: Costs and Revenue Reactions

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 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.

¹⁸As 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.

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 two 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.¹⁹ 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.²⁰

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 A16](#). 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. We note, however, that the effect of marginal HDDs is not significant and that the support for marginal HDD is small. Therefore, it could be the case that our data does not contain sufficient variation to capture this response.

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 non-marginal GDDs overlap, 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 3](#) how forecasted GDDs induce a positive input response in the form of higher fertilizer use, 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 (PKGIC). The survey is a repeated cross-section, and as such we are not able to include farm fixed effects. We

¹⁹That is, we aggregate the gridded ERA5 data for 1963-1993 at the department level using area-weights.

²⁰See [Table A18](#) for details.

only include department fixed effects (the smallest geographic identity to which we can attach the fields). Otherwise, the specification is the same as the baseline regression in [Equation 1](#).

Table 4: Timing Response to Forecasts (1 month lead)

Dependent Variables: Model:	Ploughing (1)	Irrigation (2)	Sowing (3)	Harvest (4)
<i>Variables</i>				
GDD	-0.2846*** (0.0005)	0.0628*** (0.0014)	-0.1230*** (0.0025)	-0.1024*** (0.0005)
GDD (F)	0.5321*** (0.0005)	-0.1022*** (0.0040)	0.2874*** (0.0071)	0.5398*** (0.0012)
HDD	-3.614*** (0.0037)	1.357*** (0.0301)	-0.8389*** (0.0537)	-1.836*** (0.0096)
HDD (F)	-1.259*** (0.0032)	-0.9459*** (0.0250)	0.2799*** (0.0447)	-1.002*** (0.0080)
Mean	20.07	21.38	25.08	38.37
<i>Fixed-effects</i>				
Department	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
<i>Fit statistics</i>				
Observations	44,784	44,784	44,784	44,784
R ²	0.46796	0.96372	0.53227	0.63566

Notes: One-way department-level 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 department-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. Indeed, if variations in timing only imply changes of a few weeks or days, we can expect input prices to remain constant over that period, and that these procedures will hence cost the same amount.

We observe that the date of the first irrigation, the date of harvest and the date of ploughing are all brought earlier (come earlier in the season) as a reaction to forecasted HDDs. These adaptive behaviors are likely to impact production volumes, and hence play a role in the observed, positive revenue response. On average, forecasted GDDs lead to the growing season being shifted forward (later)—in response to a GDD forecast, farmers plough, sow and harvest later. On the contrary, forecasted HDDs lead to the growing season being shifted backward (earlier). We expect that

farmers will want to extend the crop growth period when weather is more advantageous, and cut it short to avoid crop-damaging hot days when weather worsens. Here, the period between sowing and harvest is extended in response to forecasted GDDs, but it is reduced in response to forecasted HDDs. Responses to forecasted heat shocks are also of a much larger order of magnitude than responses to forecasted GDDs.

We also observe that all the procedures listed respond to realized HDDs: the first irrigation happens later (perhaps with the goal of saving water resources, and timing them closer to heat shocks), whereas ploughing, sowing and harvesting happen earlier. Realized GDDs induced responses of the same sign, but of a much smaller order of magnitude.

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.

5 Results: Effects on Profits

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. Finally, we look at the timing of effects on profits to examine the marginal net benefits of forecasts of different horizons.

5.1 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:

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

Assuming the problem is well behaved, the first order condition will equate expected marginal

revenues and expected marginal costs.²¹ That is,

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

at optimum.

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

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

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., Deschênes and Greenstone, 2007; Hsiang, 2016):

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

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 $p_t \frac{\partial q}{\partial w_t} \Big|_{x=x^*}$, the marginal impact of weather on the firm absent any influence from the marginal benefits or costs of adaptation behavior.

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 $\frac{\partial V(w_t)}{\partial f_t} = 0$.²² To see why this is a test of the envelope condition, note that because forecasts are purely informational, we have:

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

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:

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

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

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

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

²¹While we assume that output prices are not responsive to weather, or to sold output, the derivations can be extended to account for these.

²²While 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.

information they already possess), then the 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](#) verify that this is indeed the case. The table shows effects for two different definitions of profit: profit defined as value added measured by the difference between total production value and intermediary inputs (Column 1) and the gross operating income of the farm (Column 2). 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, department-specific quadratic time trends, and the level and square of realized and forecasted rainfall).

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

Dependent Variables: Model:	Value Added (1)	Profit (2)
<i>Variables</i>		
GDD	5.388 (8.456)	-2.871 (8.127)
GDD (F)	15.30 (17.47)	11.27 (17.59)
HDD	-13.53 (259.7)	91.07 (231.1)
HDD (F)	2,449.3*** (613.3)	2,051.2*** (566.8)
Mean	58,356.8	86,695.2
<i>Fixed-effects</i>		
Farm	Yes	Yes
Year	Yes	Yes
<i>Fit statistics</i>		
Observations	18,917	18,917
R ²	0.81036	0.84572

Notes. One-way department-level 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 department-specific time trends.

The table shows that HDD forecasts have a significant positive effect on farm profits for both definitions of profits.²³ A one unit increase in forecasted HDD increases farm profits by €2,449 when using value added as a measure of profit. When using our second measure, the effect is slightly lower at €2,051. On the contrary, the effect of forecasted GDDs on profits is not

²³[Table A10](#) gives the profit results for one-month and two-months ahead leads.

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 responding to forecasted growing degree days have a net-zero effect.

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.

For GDDs, we do not find a significant effect of forecasts. This is again consistent with the results on revenues and costs, but as noted above, given that our test is a necessary but not sufficient, we cannot conclude that the envelope condition holds for GDDs.

5.2 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: the context studied is not sufficiently 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, irrigation, and harvest. These actions have opportunity costs if they are not done at the right time, but they do not impose substantial additional costs on firms.²⁴

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.²⁵

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

²⁴There 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.

²⁵Gammans et al. (2017) discuss the plausibility of that assumption for France.

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.

5.3 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 department-specific time trends. We cluster the standard errors at the department level to account for spatial correlation in weather and agricultural patterns.

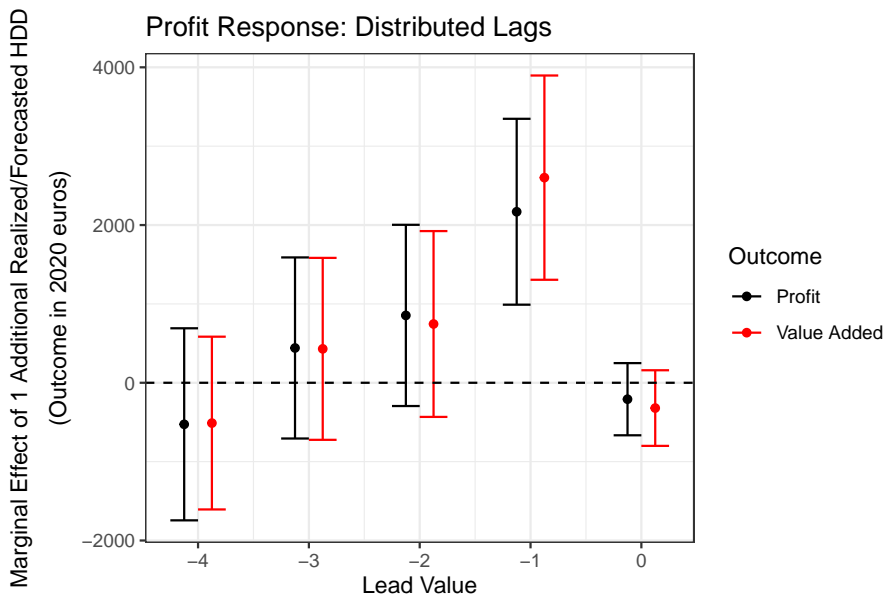


Figure 2: Profit Reactions

[Figure 2](#) 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.

6 Results: Limits to Adaptation

We have so far highlighted the relatively small costs of ex ante adaptation to heat shocks in French agriculture. For a given realized temperature, cereal, oil, and protein crop farms are able to use forecasted weather to shift their production schedule to generate relatively large revenue gains at relatively low costs. Specifically, they use forecasted weather to shift the timing of key moments of the growing season to better match the coming weather. It is important to note that this adaptation strategy is, by design, aligned with the weather variation that farmers have observed in France so far. By this, we mean two things: first, that these adaptation strategies are aligned with the range of extreme heat shocks faced in the historical record—and [Figure A11](#) shows that temperature realizations substantially above 30°C have been rare in France—and second that they are aligned with the frequency of these shocks within and across seasons. [Figure A8](#) shows that overall, heating degree days remain scarce. It is likely that both more extreme shocks, and more frequent shocks, would require different adaptation strategies.

In that sense, the results presented above can only be interpreted as indicative of low costs of adaptation to weather—and to climate change—over 1994-2018, and not as indicative of future costs of adaptation to expected climate change. The French government officially expects a 4°C heating scenario for the country over the coming century,²⁶ for which the recent past is unlikely to be informative. In order to make progress on the approximation of costs of adaptation to climate change, we examine heterogeneity of our results to different average climate conditions. Specifically, we aim at showing how the consequences of heat shocks—both in terms of their impact on profits, and on adaptation costs—change when heat shocks become more frequent and more pronounced. When the weather draws from climate become more adverse, adaptation strategies are likely to shift, potentially becoming both more costly and less effective.

To examine this heterogeneity, we compute the distribution of heating degree days in our data, and select data in the top decile. We further restrict this sample by only keeping farms observed at least twice—in order for effects to be identifiable in the presence of farm fixed effects. We then perform our analysis on this subset and compare the results to the ones from our complete sample.

²⁶See https://www.lemonde.fr/en/environment/article/2023/06/12/how-can-france-adapt-to-4-c-of-global-warming_6030922_114.html for example

Table 6: Profits and Costs Under Adversarial Weather (1 month lead)

Dependent Variable:	Profit		Costs	
Model:	(1)	(2)	(3)	(4)
<i>Variables</i>				
GDD	-2.865 (8.118)	-107.8* (55.14)	-8.211 (5.234)	-90.20** (42.12)
GDD (F1)	11.17 (17.58)	-345.5 (299.1)	23.37** (9.452)	368.3* (181.8)
HDD	89.59 (230.8)	182.1 (575.7)	-127.6 (113.4)	1,381.8*** (455.9)
HDD (F1)	2,051.0*** (566.6)	-748.8 (771.0)	-34.55 (171.7)	-403.4 (405.0)
Mean	86,687.4	77,220.1	123,270.4	128,852.2
Weather	All	Bad	All	Bad
<i>Fixed-effects</i>				
Farm	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
<i>Fit statistics</i>				
Observations	18,926	1,425	18,926	1,425
R ²	0.84571	0.89847	0.93969	0.96362

Notes. One-way department-level 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 department-specific time trends.

Results on profits and costs are shown in [Table 6](#). We want to highlight several elements in these regressions. First, the large positive and statistically significant response to forecasted HDDs disappears. On the contrary, forecasted HDDs are now associated with a negative, albeit non-significant, profit response. In very hot years, farms seem to lose their ability to use heat forecasts in order to generate additional profits. The second element we want to highlight is that we also start observing non-zero cost responses to HDDs. Realized HDDs are now associated with a large, positive cost response. A one unit increase in HDD generates a €1,382 increase in costs, or about 1% increase in costs. While one-month ahead forecasts do not generate an additional cost response, the two-month-ahead forecasts shown in [Table A20](#) generate a large negative one. There, a one unit increase in forecasted HDD leads to a decrease in costs of €2,555, or a 2% decrease in costs.

Taking stock of these results, in years with especially severe heat, farmers seem to lose their ability to use heat forecasts to generate additional profits. On the contrary, they start responding to realized HDDs with on-the-spot costly adaptation, and when receiving information about incoming shocks early enough, they seem to downscale their production by heavily decreasing

their operational costs.

We next look at output price, quantity and storage responses in [Table 7](#).

Table 7: Price and Quantity Responses Under Adversarial Weather

Dependent Variables: Model:	Price (1)	Output (2)	Storage (3)	Price (4)	Output (5)	Storage (6)
<i>Variables</i>						
GDD	-0.2379 (0.2157)	-1.972 (5.165)	1.001 (5.564)	-0.3663* (0.2105)	-10.05** (4.791)	-3.268 (4.101)
GDD (F1)	-0.1334 (0.5636)	2.214 (23.58)	-3.163 (14.80)			
GDD (F2)				0.3774 (0.4177)	-9.833 (7.146)	-3.406 (6.371)
HDD	1.978 (1.959)	10.21 (41.11)	-61.90** (27.89)	2.865 (2.317)	65.23 (41.84)	-39.97 (31.68)
HDD (F1)	-1.546 (1.659)	-54.15 (57.21)	-41.99 (48.23)			
HDD (F2)				-7.605** (3.650)	-264.6** (108.3)	-109.4 (76.32)
Mean	178.2	8,256.1	20.83	178.2	8,256.1	20.83
<i>Fixed-effects</i>						
Farm	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes
<i>Fit statistics</i>						
Observations	1,425	1,425	1,425	1,425	1,425	1,425
R ²	0.65144	0.96024	0.55139	0.64839	0.96067	0.55015

Notes. One-way department-level 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 department-specific time trends.

In the hottest years, realized HDDs do not have a statistically significant impact on profits. When we break down the profit response, things are more complex. Indeed, we have seen that farms do increase their costs of production, and [Table A21](#) details this cost response in terms of increases in labor, fertilizer and irrigation. Finally, [Table 7](#) shows that farms also response to realized HDDs by decreasing their storage. The null profit impact of heat shocks thus becomes a combination of an increase in production costs which seems efficient at protecting crop yields—indeed we observe a null response of output quantity to HDDs—and a decrease in storage, which seems to offset the increase in costs at the profit level.

Moving to two-month ahead forecasts, we observe that farms decrease their production. A one unit increase in HDD forecasts leads to a 264 cwt decrease in output, or a 3% decrease.

As we are looking at a response to forecasts, this decrease is behavioral, and orthogonal to the direct impact of heat on crop growth. This behavioral response tracks at the input level, and in [Table A21](#), we observe that farms respond jointly by increasing the amount of land they use, and largely decreasing their seeds expenses. In our main sample, ex ante adaptation was costless, did not imply any increase in input usage, and led to large profit gains. In this set of hottest years, ex ante adaptation is most salient two-month ahead, and consists of a planned decrease in production scope, which seems mostly driven by a large decrease in seed density.

The implications of heat shocks appear to be quite heterogeneous across the distribution of heat realizations. While for the time being, French farmers seem to have been able to leverage heat forecasts to generate large profit gains at a relatively low cost, the consequences of adaptation on the production schedule become more manifest as heat increases. In our sample of worse years, farmers lose their ability to generate profit gains from heat forecasts. Instead, they rely on at least two different strategies. When informed about incoming heat shocks early enough, they downscale their production—following the predictions we get from a profit maximizing firm expecting a negative productivity shock—and limit their operating costs. At the time of the shock, however, and in all likelihood in order to protect their already sown crops, farmers adopt costly adaptation strategies, increasing input usage along several dimensions, in order to limit the negative effect of the shocks on their production. When available, farmers also sell their stored production. These two countervailing effects lead to an overall neutral impact of realized HDDs on profits.

We conclude by noting that on-the-spot strategies are likely to be transitory strategies, ill-suited to a future climate where heat shocks are frequent. Indeed, increases in irrigation, and a reliance on storage, are both more adapted to one-off heat shocks, rather than a series of them. While farmers have been able to mitigate the impacts of heat in the very hot years observed over 1994-2018, this will likely not remain the case in a hotter world. There is also a fundamental limit to how much decisions about the timing of planting and harvesting can be changed.

7 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 in generating these gains.

This data also allows us to highlight 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. When informed sufficiently in advance, they also use heat forecasts to reduce their production and accordingly reduce their production costs.

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.
- 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 National 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*.
- DESCHÈNES, 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*.
- 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*.
- MENDELSON, 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.
- SHERIFF, G. (2005): “Efficient Waste? Why Farmers Over-Applly Nutrients and the Implications for Policy Design,” *Review of Agricultural Economics*.
- 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

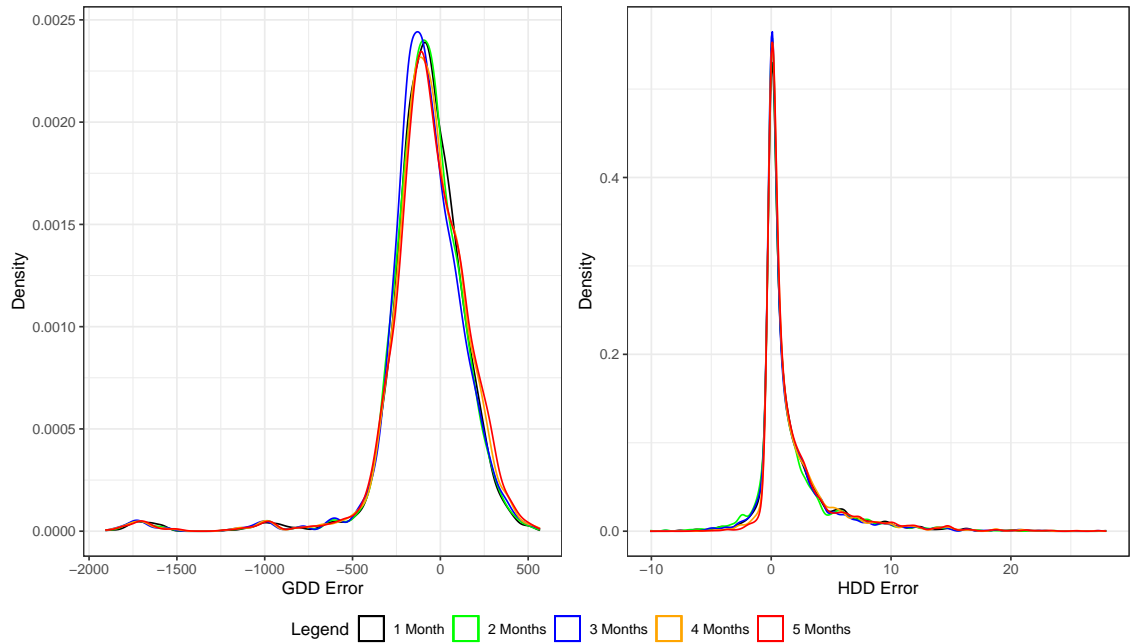


Figure A1: Distribution of Temperature Forecast Error

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.




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.

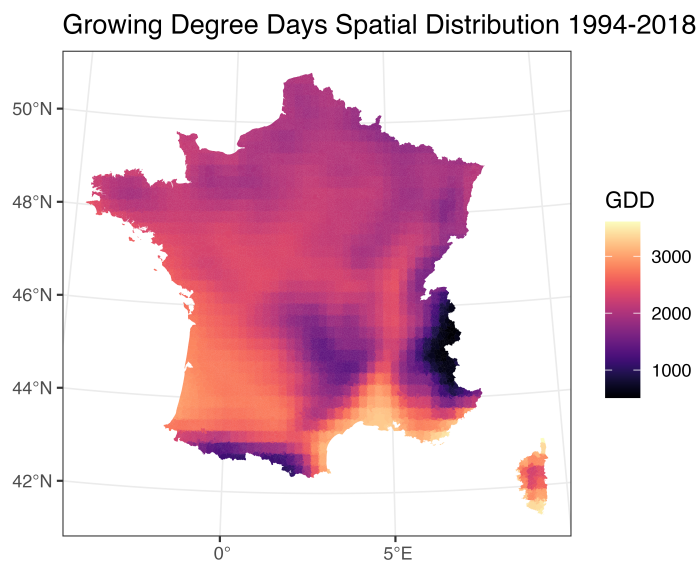


Figure A7: Distribution of Growing Degree Days

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 Alps at the border with Switzerland and Italy).

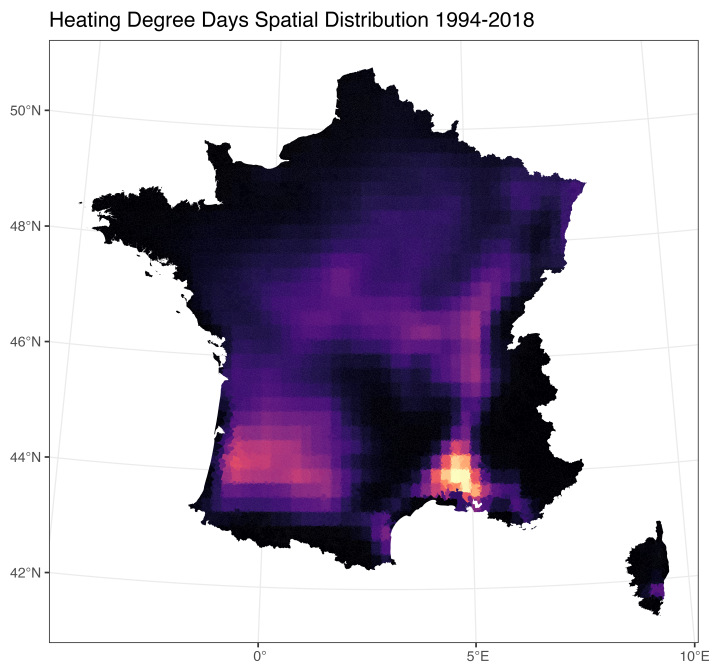


Figure A8: Distribution of Heating Degree Days

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.

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 growing degree day values, and low heating degree day values for the 1994-2018 period.

A.1.3 Multi-Product Farms




Figure A9: 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.1.4 Additional Results




Figure A10: Effect of Forecasted 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.

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	Pctl(25)	Pctl(75)	Max
Ploughing Date	44,789	20.074	16.924	0.000	0.000	29.000	99.000
Sowing Date	44,789	25.084	12.288	0.000	19.000	33.000	54.000
First Irrigation Date	44,789	21.384	38.298	0.000	0.000	35.000	99.000
Harvest Date	44,789	38.372	17.154	0.000	37.000	44.000	99.000

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
<i>Panel A</i>		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
<i>Panel B</i>		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)
<i>Variables</i>				
GDD	3.665 (9.695)	-8.131 (5.239)	9.514 (9.449)	-6.340 (5.115)
GDD (F)	39.10** (17.35)	23.00** (9.438)		
GDD (F2)			57.75** (24.18)	32.86*** (9.367)
HDD	-206.1 (265.6)	-136.1 (113.1)	-281.3 (223.5)	-143.3 (105.5)
HDD (F)	1,177.3** (555.8)	-39.15 (171.7)		
HDD (F2)			1,069.6 (732.1)	-151.8 (436.5)
Mean	155,386.2	123,249.1	155,386.2	123,249.1
<i>Fixed-effects</i>				
Farm	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
<i>Fit statistics</i>				
Observations	18,917	18,917	18,917	18,917
R ²	0.88977	0.93971	0.88977	0.93972

Notes. One-way department-level 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 department-specific time trends. Sales correspond to total sales at the farm levels, and costs to total costs.

Table A8: Price and Quantity Reactions to Forecasts

Dependent Variables: Model:	Price (1)	Output (2)	Storage (3)	Price (4)	Output (5)	Storage (6)
<i>Variables</i>						
GDD	-0.1755** (0.0838)	-0.2196 (0.4644)	-0.9748** (0.4070)	-0.1709** (0.0785)	0.0515 (0.5045)	-0.8601** (0.4257)
GDD (F)	0.3076 (0.2462)	1.561* (0.8953)	-0.7268 (0.9964)			
GDD (F2)				0.3508 (0.3119)	4.971*** (0.9960)	2.167** (1.078)
HDD	0.0208 (1.313)	-13.81 (14.17)	-2.921 (9.653)	0.6533 (1.035)	-13.67 (14.65)	-0.7248 (10.67)
HDD (F)	-0.9513 (2.096)	71.81*** (19.04)	45.19** (21.50)			
HDD (F2)				-6.921 (4.211)	-13.18 (29.20)	-19.48 (22.34)
Mean	263.0	7,793.0	88.65	263.0	7,793.0	88.65
<i>Fixed-effects</i>						
Farm	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes
<i>Fit statistics</i>						
Observations	18,917	18,917	18,917	18,917	18,917	18,917
R ²	0.71819	0.94480	0.10975	0.71815	0.94496	0.11004

Notes: One-way department-level 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 department-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.

Table A9: Cost and Revenue Reactions to HDD—Comparison

Dependent Variables:	Sales		Costs	
Model:	(1)	(2)	(3)	(4)
<i>Variables</i>				
GDD	3.665 (9.695)	7.610 (9.483)	-8.131 (5.239)	0.1841 (6.678)
GDD (F)	39.10** (17.35)		23.00** (9.438)	
HDD	-206.1 (265.6)	-120.7 (223.9)	-136.1 (113.1)	-137.3 (121.5)
HDD (F)	1,177.3** (555.8)		-39.15 (171.7)	
Mean	155,386.2	155,386.2	123,249.1	123,249.1
<i>Fixed-effects</i>				
Farm	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
<i>Fit statistics</i>				
Observations	18,917	18,917	18,917	18,917
R ²	0.88977	0.88649	0.93971	0.93584

Notes. One-way department-level 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 department-specific time trends. Sales correspond to total sales at the farm levels, and costs to total costs.

Table A10: Profit Reaction to Forecasts

Dependent Variables: Model:	Value Added (1)	Profit (2)	Value Added (3)	Profit (4)
<i>Variables</i>				
GDD	5.388 (8.456)	-2.871 (8.127)	12.70 (9.467)	3.244 (8.990)
GDD (F)	15.30 (17.47)	11.27 (17.59)		
GDD (F2)			66.99*** (23.23)	56.81*** (20.55)
HDD	-13.53 (259.7)	91.07 (231.1)	-92.13 (242.8)	12.87 (218.6)
HDD (F)	2,449.3*** (613.3)	2,051.2*** (566.8)		
HDD (F2)			901.6 (566.3)	969.3* (560.8)
Mean	58,356.8	86,695.2	58,356.8	86,695.2
<i>Fixed-effects</i>				
Farm	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
<i>Fit statistics</i>				
Observations	18,917	18,917	18,917	18,917
R ²	0.81036	0.84572	0.80975	0.84545

Notes. One-way department-level 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 department-specific time trends.

Table A11: Input Reactions to Forecasts (1 month lead)

Dependent Variables: Model:	Land (1)	Labor (2)	Fertilizer (3)	Phytopsanitary (4)	Seeds (5)	Irrigation (6)
<i>Variables</i>						
GDD	-0.1741 (0.1171)	-0.1082 (0.0900)	-2.380 (1.511)	-1.826** (0.8668)	-0.4362 (0.9441)	0.3178 (0.2582)
GDD (F)	0.4538 (0.4068)	0.2786 (0.2362)	12.98*** (3.216)	-0.9860 (1.947)	1.655 (1.342)	-0.4195 (0.6964)
HDD	-6.764 (4.358)	2.023 (2.600)	-96.00** (37.16)	-18.40 (25.56)	-17.62 (21.13)	13.73 (8.393)
HDD (F)	11.35 (14.32)	-1.099 (4.051)	-44.57 (72.34)	12.56 (55.83)	-14.04 (49.90)	-6.651 (10.34)
Mean	1,025.6	2,684.0	31,777.2	23,072.2	12,231.1	683.1
<i>Fixed-effects</i>						
Farm	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes
<i>Fit statistics</i>						
Observations	18,917	18,917	18,917	18,917	18,917	18,917
R ²	0.74053	0.87040	0.89992	0.92675	0.88451	0.87414

Notes: One-way department-level 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 department-specific time trends.

Table A12: Input Reactions to Forecasts (2 months lead)

Dependent Variables: Model:	Land (1)	Labor (2)	Fertilizer (3)	Phytopsanitary (4)	Seeds (5)	Irrigation (6)
<i>Variables</i>						
GDD	-0.1222 (0.1144)	-0.0816 (0.0935)	-1.537 (1.503)	-2.021** (0.8714)	-0.2452 (0.8774)	0.2470 (0.2458)
GDD (F2)	-0.0759 (0.3790)	0.5875* (0.3521)	14.90*** (4.251)	3.517 (2.667)	0.7237 (1.886)	-0.0683 (0.5944)
HDD	-7.198 (4.394)	1.021 (2.604)	-109.4*** (38.73)	-8.869 (23.63)	-18.88 (20.81)	14.36* (7.777)
HDD (F2)	-1.502 (15.24)	-0.6676 (6.564)	-29.22 (84.08)	-60.07 (78.33)	-92.27* (52.56)	7.961 (11.33)
Mean	1,025.6	2,684.0	31,777.2	23,072.2	12,231.1	683.1
<i>Fixed-effects</i>						
Farm	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes
<i>Fit statistics</i>						
Observations	18,917	18,917	18,917	18,917	18,917	18,917
R ²	0.74046	0.87043	0.89991	0.92674	0.88454	0.87416

Notes: One-way department-level 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 department-specific time trends.

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

Dependent Variables: Model:	All Store Prices (1)	Fertilizers (2)	Pesticides (3)	Seeds (4)	Land (5)
<i>Variables</i>					
GDD	-0.1484 (0.1836)	-0.1110 (0.0685)	-0.3402 (0.3375)	-0.1888 (0.1729)	-0.6888*** (0.2558)
GDD (F)	-0.4966 (0.7285)	0.1272 (0.1125)	-1.200 (1.418)	-0.4061 (0.4688)	-0.0825 (0.5971)
HDD	-18.90* (9.794)	0.7204 (1.757)	-23.72 (16.12)	-16.81 (16.62)	8.927* (4.874)
HDD (F)	-0.4188 (8.421)	0.4165 (3.343)	7.001 (14.88)	0.8332 (7.464)	1.485 (10.67)
Mean	1,899.6	517.0	3,151.3	310.8	4,848.3
<i>Fixed-effects</i>					
Company	Yes	Yes	Yes	Yes	
Product	Yes	Yes	Yes	Yes	
Year	Yes	Yes	Yes	Yes	Yes
Farm					Yes
<i>Fit statistics</i>					
Observations	33,174	9,258	17,883	6,033	17,003
R ²	0.35655	0.63051	0.22709	0.40074	0.97856

Notes: One-way department-level 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 department-specific time trends.

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

Dependent Variables: Model:	All Store Prices (1)	Fertilizers (2)	Pesticides (3)	Seeds (4)	Land (5)
<i>Variables</i>					
GDD	-0.1837 (0.1918)	-0.0778 (0.0704)	-0.3747 (0.3514)	-0.2784 (0.2660)	-0.7000*** (0.2510)
GDD (F2)	0.0790 (0.4944)	0.3884* (0.2038)	0.6180 (0.7069)	-1.528 (1.715)	0.1878 (0.7989)
HDD	-17.18* (9.766)	0.2372 (1.773)	-21.30 (16.03)	-15.76 (16.46)	10.04** (4.930)
HDD (F2)	-6.076 (17.08)	1.537 (3.840)	-3.071 (26.89)	-38.52 (26.77)	-0.9595 (11.41)
Mean	1,899.6	517.0	3,151.3	310.8	4,848.3
<i>Fixed-effects</i>					
Company	Yes	Yes	Yes	Yes	
Product	Yes	Yes	Yes	Yes	
Year	Yes	Yes	Yes	Yes	Yes
Farm					Yes
<i>Fit statistics</i>					
Observations	33,174	9,258	17,883	6,033	17,003
R ²	0.35656	0.63059	0.22690	0.40088	0.97859

Notes: One-way department-level 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 department-specific time trends.

Table A15: Output Quantities Reactions to Forecasts (1 months lead)

Dependent Variables: Model:	Wheat (1)	Corn (2)	Sunflower (3)	Rapeseeed (4)	Beetroot (5)
<i>Variables</i>					
GDD	-0.5188 (0.3300)	0.5044* (0.2980)	0.0946 (0.0944)	-0.5117*** (0.1247)	1.156 (2.350)
GDD (F)	0.5089 (0.6570)	1.198* (0.6946)	-0.1981 (0.2463)	0.3381 (0.3535)	-4.636 (3.730)
HDD	0.0956 (7.742)	-30.63** (14.61)	-2.412 (2.342)	12.49*** (3.867)	-3.329 (42.87)
HDD (F)	56.90*** (15.94)	4.099 (16.88)	-2.215 (3.516)	20.38*** (7.557)	78.12 (85.26)
Mean	3,840.1	3,464.3	509.3	914.9	9,127.8
<i>Fixed-effects</i>					
Farm	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes
<i>Fit statistics</i>					
Observations	17,336	10,346	7,203	12,394	2,267
R ²	0.90606	0.93989	0.79500	0.81774	0.93556

Notes: One-way department-level 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 department-specific time trends.

Table A16: Farm-Level Reactions: Decomposed Weather

Dependent Variables:	Sales	Costs
Model:	(1)	(2)
<i>Variables</i>		
GDD (m)	5.566 (10.39)	-0.3131 (6.164)
GDD (nm)	3.224 (11.86)	-3.604 (5.576)
GDD (m) (F)	23.48 (24.95)	8.479 (10.66)
GDD (nm) (F)	-0.5601 (31.40)	17.92* (9.897)
HDD (m)	-328.9 (345.9)	-245.8 (147.6)
HDD (m) (F)	-1,357.5 (1,263.8)	300.4 (528.4)
HDD (nm)	-290.8 (1,019.0)	469.1 (516.1)
HDD (nm) (F)	1,834.8* (1,083.1)	-193.2 (370.4)
Mean	155,396.1	123,264.8
<i>Fixed-effects</i>		
Farm	Yes	Yes
Year	Yes	Yes
<i>Fit statistics</i>		
Observations	17,581	17,581
R ²	0.88872	0.94025

Notes. One-way department-level 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 department-specific time trends. Here (m) indicates marginal weather, (nm) stands for non-marginal, and the F indicates one month ahead seasonal forecasts.

Table A17: Cost and Revenue Reactions to Alternative Weather

Dependent Variables: Model:	Sales (1)	Costs (2)	Sales (3)	Costs (4)
<i>Variables</i>				
GDD	0.4614 (8.438)	-10.22** (5.024)	8.798 (8.425)	-9.734* (5.139)
GDD (F1)	19.87 (20.30)	27.27** (10.78)		
GDD (F2)			57.18** (23.42)	28.98*** (9.423)
HDD	-72.50 (147.1)	-79.88 (58.19)	-89.70 (127.6)	-54.88 (55.52)
HDD (F1)	628.0*** (175.2)	8.212 (82.46)		
HDD (F2)			-10.23 (187.8)	5.031 (95.15)
FDD	18.67 (29.28)	2.115 (12.04)	34.86 (29.27)	3.225 (11.35)
FDD (F1)	119.3 (117.2)	-93.30* (55.41)		
FDD (F2)			76.95 (164.7)	-17.53 (37.19)
Mean	155,637.5	123,304.2	155,637.5	123,304.2
<i>Fixed-effects</i>				
Farm	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
<i>Fit statistics</i>				
Observations	18,428	18,428	18,428	18,428
R ²	0.89137	0.94083	0.89128	0.94085

Notes. One-way department-level 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 department-specific time trends. Sales correspond to total sales at the farm levels, and costs to total costs. Standard errors are clustered at the department level.

Table A18: 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 A19: Profits Under Adversarial Weather

Dependent Variable:	Profits			
Model:	(1)	(2)	(3)	(4)
<i>Variables</i>				
GDD	-2.865 (8.118)	-107.8* (55.14)	3.252 (8.983)	-99.23 (67.39)
GDD (F1)	11.17 (17.58)	-345.5 (299.1)		
GDD (F2)			56.72*** (20.54)	178.6** (71.88)
HDD	89.59 (230.8)	182.1 (575.7)	11.70 (218.3)	-603.3 (650.8)
HDD (F1)	2,051.0*** (566.6)	-748.8 (771.0)		
HDD (F2)			971.9* (560.6)	-985.5 (1,282.1)
Mean	86,687.4	77,220.1	86,687.4	77,220.1
Weather	All	Bad	All	Bad
<i>Fixed-effects</i>				
Farm	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
<i>Fit statistics</i>				
Observations	18,926	1,425	18,926	1,425
R ²	0.84571	0.89847	0.84543	0.89907

Notes. One-way department-level 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 department-specific time trends.

Table A20: Costs Under Adversarial Weather

Dependent Variable:	Costs			
Model:	(1)	(2)	(3)	(4)
<i>Variables</i>				
GDD	-8.211 (5.234)	-90.20** (42.12)	-6.416 (5.110)	-162.8*** (36.84)
GDD (F1)	23.37** (9.452)	368.3* (181.8)		
GDD (F2)			32.56*** (9.363)	2.308 (59.34)
HDD	-127.6 (113.4)	1,381.8*** (455.9)	-134.1 (106.0)	1,603.3*** (507.2)
HDD (F1)	-34.55 (171.7)	-403.4 (405.0)		
HDD (F2)			-156.5 (436.5)	-2,555.6*** (739.9)
Mean	123,270.4	128,852.2	123,270.4	128,852.2
<i>Fixed-effects</i>				
Farm	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
<i>Fit statistics</i>				
Observations	18,926	1,425	18,926	1,425
R ²	0.93969	0.96362	0.93970	0.96358

Notes. One-way department-level 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 department-specific time trends.

Table A21: Inputs Under Adversarial Weather (2 months forecasts)

Dependent Variables: Model:	Land (1)	Labor (2)	Fertilizer (3)	Phytopsanitary (4)	Seeds (5)	Irrigation (6)
<i>Variables</i>						
GDD	3.256 (2.270)	-1.765** (0.6998)	16.16 (13.27)	-23.64*** (8.449)	-5.393 (8.076)	-3.355 (4.824)
GDD (F2)	3.284 (2.320)	-0.7428 (0.6031)	-21.44 (16.73)	-29.33** (13.05)	21.00 (13.24)	-10.22* (5.502)
HDD	-13.99 (27.52)	22.77** (8.548)	208.0* (108.8)	108.8 (78.12)	7.787 (74.47)	90.41** (39.99)
HDD (F2)	65.35* (36.75)	18.39 (22.74)	-85.24 (431.7)	298.4 (268.1)	-642.5*** (157.8)	-43.78 (136.9)
Mean	773.1	2,488.9	30,492.5	21,560.9	12,956.9	1,256.8
<i>Fixed-effects</i>						
Farm	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes
<i>Fit statistics</i>						
Observations	1,425	1,425	1,425	1,425	1,425	1,425
R ²	0.79659	0.90295	0.92658	0.93031	0.93512	0.91531

Notes. One-way department-level 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 department-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:

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

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 :

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

- **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.

- **Profit** Defined as the gross operating income of the farm.

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

B.2 Weather Data

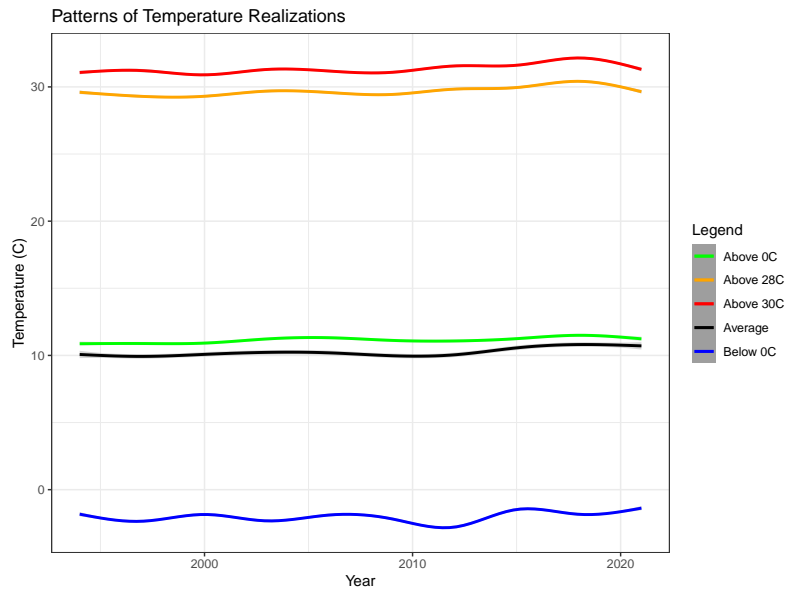


Figure A11: Unconditional and Conditional Average Temperature Realizations

C Validation of Results

Table A22: Corn-Specific Outcomes

Dependent Variables: Model:	Output (1)	Price (2)	Sales Value (3)	Sales Quantity (4)
<i>Variables</i>				
gdd	0.5611* (0.2889)	-0.0081 (0.0100)	9.598 (7.380)	0.7575* (0.4078)
hdd	-33.03** (14.64)	0.8153** (0.3098)	-153.5 (176.4)	-16.11 (11.93)
Mean	3,464.3	150.0	53,507.8	3,437.5
<i>Fixed-effects</i>				
Year	Yes	Yes	Yes	Yes
Farm	Yes	Yes	Yes	Yes
<i>Fit statistics</i>				
Observations	10,346	9,958	10,346	10,346
R ²	0.93980	0.73642	0.88250	0.88108

Notes. One-way department-level 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 department-specific time trends.

Table A23: Farm-Level Outcomes

Dependent Variables: Model:	Sales (1)	Costs (2)	Intermediate Inputs (3)	Value Added (4)	Profit (5)	Price (6)
<i>Variables</i>						
GDD	8.757 (9.647)	-7.059 (5.205)	-8.033* (4.261)	11.28 (9.475)	2.061 (9.054)	-0.1833** (0.0810)
HDD	-159.9 (245.1)	-110.7 (114.4)	-99.34 (93.29)	30.54 (253.1)	126.8 (227.6)	0.7206 (1.138)
Mean	155,386.2	123,249.1	103,749.6	58,356.8	86,695.2	263.0
<i>Fixed-effects</i>						
Year	Yes	Yes	Yes	Yes	Yes	Yes
Farm	Yes	Yes	Yes	Yes	Yes	Yes
<i>Fit statistics</i>						
Observations	18,917	18,917	18,917	18,917	18,917	18,917
R ²	0.88962	0.93968	0.93646	0.80930	0.84514	0.71770

Notes. One-way department-level 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 department-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 when using forecasts, hinting that heat shocks act as a negative multiplier to production rather than an additive shock.

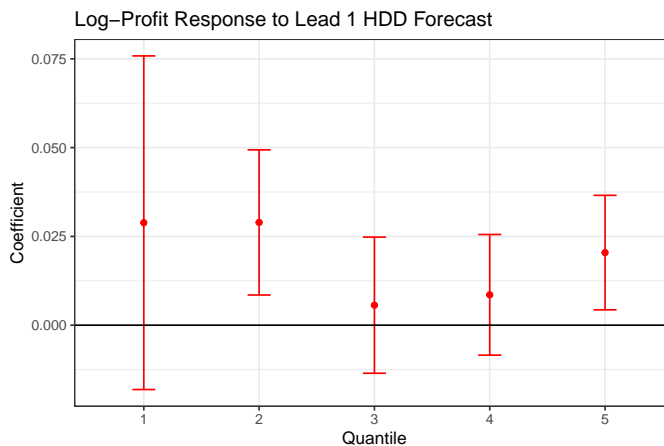


Figure A12: Varying Response to Fixed Lead of HDD

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.

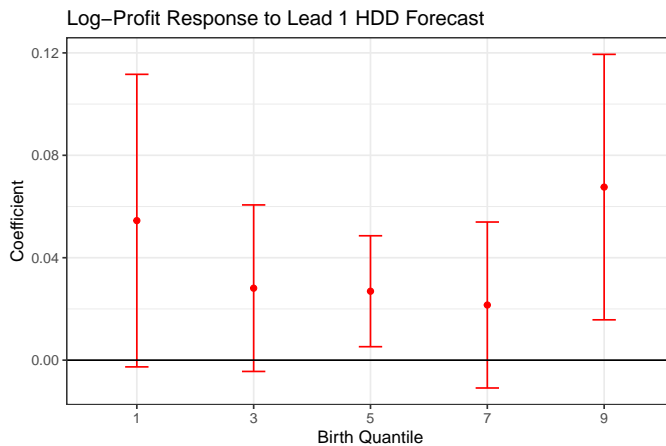


Figure A13: Profit Response to Forecasts across Manager’s Age

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 A24: Farm-Level Profit - Lead 1 with Lags

Dependent Variables: Model:	Value Added (1)	Profit (2)	Value Added (3)	Profit (4)
<i>Variables</i>				
GDD	9.573 (9.723)	0.5344 (9.239)	20.12* (11.02)	8.999 (10.41)
GDD (F)	25.72 (19.90)	18.68 (18.89)		
GDD (F2)			79.71*** (25.60)	63.54*** (23.07)
GDD (lag)	-2.743 (9.122)	-6.474 (9.197)	-4.851 (8.660)	-8.405 (8.880)
HDD	-63.17 (247.1)	75.48 (225.6)	-169.1 (247.0)	-22.15 (221.5)
HDD (F)	2,271.2*** (635.7)	1,872.2*** (581.3)		
HDD (F2)			1,052.7 (676.6)	1,182.0* (666.7)
HDD (lag)	284.4 (227.4)	368.1 (245.7)	326.7 (240.8)	389.7 (252.6)
Mean	58,356.8	86,695.2	58,356.8	86,695.2
<i>Fixed-effects</i>				
Farm	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
<i>Fit statistics</i>				
Observations	16,314	16,314	16,314	16,314
R ²	0.81435	0.84689	0.81398	0.84675

Notes. One-way department-level 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 department-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 a last 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 A25: Farm-Level Profit - Alternative Weather

Dependent Variables: Model:	Value Added (1)	Profit (2)	Value Added (3)	Profit (4)
<i>Variables</i>				
GDD	8.980 (8.606)	0.2179 (8.055)	20.30** (10.17)	9.415 (9.417)
GDD (F)	-16.08 (16.78)	-14.62 (16.79)		
GDD (F2)			59.91** (24.04)	49.54** (21.96)
HDD	32.81 (149.0)	81.40 (143.2)	-34.29 (147.9)	18.86 (144.2)
HDD (F)	931.4*** (182.9)	787.9*** (168.9)		
HDD (F2)			102.7 (164.1)	94.81 (176.0)
FDD	62.13** (30.54)	54.36* (31.32)	76.45** (33.10)	66.20* (33.55)
FDD (F)	316.8** (123.9)	269.9** (130.8)		
FDD (F2)			158.0 (124.5)	218.4* (115.8)
Mean	58,332.2	87,330.0	58,332.2	87,330.0
<i>Fixed-effects</i>				
Farm	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
<i>Fit statistics</i>				
Observations	18,428	18,428	18,428	18,428
R ²	0.81084	0.84716	0.80996	0.84672

Notes. One-way department-level 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 department-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 A17](#) 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 A26: Farm-Level Profit - No Time Trend

Dependent Variables: Model:	Sales (1)	Costs (2)	Sales (3)	Costs (4)
<i>Variables</i>				
GDD	3.473 (9.250)	-1.273 (6.461)	8.786 (9.342)	0.8753 (6.768)
GDD (F1)	47.76*** (14.61)	37.66*** (10.55)		
GDD (F2)			68.20*** (25.54)	45.50*** (16.85)
HDD	-211.9 (239.2)	-184.3 (122.5)	-263.1 (211.3)	-183.4 (123.4)
HDD (F1)	1,342.2** (642.8)	39.32 (336.5)		
HDD (F2)			1,381.6** (629.4)	156.4 (448.6)
Mean	155,386.2	123,249.1	155,386.2	123,249.1
<i>Fixed-effects</i>				
Farm	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
<i>Fit statistics</i>				
Observations	18,917	18,917	18,917	18,917
R ²	0.88667	0.93591	0.88671	0.93591

Notes. One-way department-level 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 department-specific time trends.