Testing the Application of the Envelope Theorem to Estimation of Climate Change Damages*

Tristan Du Puy and Jeffrey G. Shrader

[1]Columbia University

January 15, 2024

Abstract: Over the last two decades, a large literature has developed that uses weather data to estimate damage caused by climate change. The envelope theorem is the most important theoretical concept underlying these estimates. It allows realized weather effects to be sufficient statistics for the effects of climate change while eliminating the confounding influence of agent actions. It also allows for inference about marginal costs of climate adaptation from estimates of marginal benefits of adaptation. We test the applicability of the envelope theorem in the context of French agriculture, using precise, farm-level accounting data from a yearly repeated panel. We show that extreme heat forecasts have significant positive impacts on farm profits, contradicting the envelope theorem's implied marginality of re-optimization behavior on profits. We decompose the variation underlying extreme heat in seasonally-aggregated data and show that most of the variation is non-marginal. We also show evidence for discrete, no-cost adaptive behavior by farmers. Both might contradict the necessary conditions for the applicability of the envelope theorem.

*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. All remaining errors are ours.

1 Introduction

The envelope theorem is the most important theoretical concept underlying current estimates of climate change damages. The envelope theorem allows one to assess the impacts of weather variation by eliminating the confounding influence of agent actions (Deschênes and Greenstone, 2007; Hsiang, 2016; Deryugina and Hsiang, 2017; Merel et al., 2022). Drawing on these net-of-adaptation impacts of weather variation, and under certain conditions, one can then recover the effects of climate change. Envelope conditions also allow one to recover estimates of cost of adapting to climate change by integrating observed marginal benefits of adaptation over expected changes in climate distributions (Carleton et al., 2022).

Thousands of papers have now been written following this set of arguments.[1]. These estimates in turn underlie many prominent estimates of the social cost of carbon—an important input into climate policy. Despite the widespread use of the method, there have been almost no tests of whether the conditions implied by the envelope theorem—namely equality between marginal costs and marginal benefits of adaptive actions—actually apply in practice.[2] This paper tests whether the simple envelope conditions invoked by the climate damage literature hold in the context of agricultural damages from extreme heat—the most widely studied economic setting for climate change (Mendelsohn et al., 1994; Schlenker et al., 2005; Schlenker and Roberts, 2009).

The envelope theorem is of particular interest in the climate change context, and its use goes beyond isolating direct impacts of weather from the consequences of adaptation. Indeed, researchers usually only observe data related to profit, production, or efficiency—which provide informative about the marginal benefits of adaptation—but not cost data. In the agricultural case for example, researchers typically observe yields as in Schlenker and Roberts (2009), farm profits as in Deschênes and Greenstone (2007), or land values as in Mendelsohn et al. (1994)—all informative of benefits but not costs.[3]

As argued by Carleton et al. (2022) and Kelly et al. (2005), policymakers should really be interested in the sum composed of the damages of climate changes and the costs of adaptation to this change. Envelope conditions have been invoked by recent studies of climate damage to provide estimates of adaptation costs. 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. This comparison often takes the form of a cross-sectional analysis. Comparing sensitivity to weather variation across space, one can infer how long-term adaptation to different climates induced different sensitivity to weather, allowing for a quantification of the benefits of adaptation. These benefits can then be converted into marginal costs of adaptation, then integrated from costs of adaptation to weather into costs of adaptation to climate.

This study performs a test of the envelope theorem which relies on the combination of profit data with detailed realized and forecasted seasonal weather data. Using these two pieces, the empirical test is straightforward. The standard envelope condition used in the literature is that marginal costs and marginal benefits of adaptation are equal at optimum. In addition, the envelope theorem implies that marginal variation in the parameters of the optimization problem

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

[2] Guo and Costello (2013) is a notable exception, discussed further below. They look at the impact of climate change in a forestry context with discrete decision making.

[3] Even in non-agricultural settings, this data issue is common. The recent work by the Climate Impact Lab bounds adaptation costs by estimating the marginal benefits of adaptation then aggregating up those benefits. See, for instance, Carleton et al. (2022).

will only lead to negligible deviations from that equality. In other words, adaptive actions reacting to marginal weather changes should have a net zero impact on profit.

Controlling for realized weather, observed profit responses to forecasts will only be driven by the ex-ante adaptive actions implemented by agents, and corresponding market-level adjustments. If the envelope condition holds, the response of profits to forecasts should then be null, a directly testable hypothesis given the data available in our setting. We implement this test using a regression of profits on forecasts, controlling for realized weather. This regression cleanly identifies the effect of changes in actions in response to upcoming weather (i.e. ex ante adaptation) while holding fixed the effects of weather realizations. The test, set in the agricultural context, relies on the use of detailed seasonal forecasts, which remain a scarce data product, and their combination with detailed and well-measured farm-level profit data.

In addition to information on profits, our study also makes use of a uniquely detailed dataset on farm-level agricultural production in France that enables an unprecedentedly comprehensive accounting of both the marginal benefits of adaptation (the monetized value of changes in production practices in response to weather) as well as the marginal costs of adaptation (the monetized change in production practices times their prices). With it, we can validate the sign and significance of profit responses to forecasts, by looking independently at the effect of seasonal weather and seasonal weather forecasts on farm revenue and farm production costs. We can also provide estimates of the costs of adaptation to weather variation which do not rely on the envelope theorem. Finally, we can decompose observed responses into output volumes and output price reactions, as well as input volumes and price reactions. These estimates remain, however, specific to the French context, where climate change is expected to have milder consequences than in many other places.

Our test is implemented using the French Farm Accountancy Data Network (FADN), a detailed dataset of agricultural inputs, outputs, and prices for a panel of over 14,000 farms in France, combined with high-resolution data on weather (temperature and precipitation) and seasonal forecasts issued by the European Centre for Medium Range Weather Forecasting (ECMWF). The quality of the farm data, recovered by accountants for the monitoring the European Union's Common Agricultural Policy, is of particular importance to assess the sensitivity of profit margins to potential channels of adaptation.

Implementing the test, we clearly reject the implications of the envelope theorem. Both of the available measures of farm profit (value added and gross operating income) react positively and significantly to forecasted heat shocks, a result which is stable to the lead value used for the forecast and to different aggregations of hourly weather data into seasonal values.

The fact that our data provides independent measures of crop-specific output quantity and output revenue allows us to identify the source of the net benefits of adaptation. As such, we can confirm that the absence of cost response to forecasted heat shocks is not the result of a combined response of input levels and prices, and that neither of them significantly respond to forecasts. Along the same lines, we can identify the source of variation in revenues, which comes directly from a within-season, positive output quantity response. Specifically, wheat production strongly reacts to increases in heating degree days. In our French context, heating degree day variation is almost entirely driven by variation in hourly temperature around the 30°C threshold. Given the prevalence of wheat production in France, this explains the observed positive impact of heating degree day forecasts.

Using a repeated cross-section of plot level data on farming practices in France, we then highlight how the timing of different profit-relevant actions across the growing season reacts to forecasted heat, potentially isolating the source of revenue-enhancing and cost-neutral responses to forecasts in the farm-level dataset. We also discuss how these timing decisions are both discrete in nature, and are likely to impact farmers' objective functions in a discrete way. Indeed, poor timing—which exposes crops at a vulnerable growth stage to freeze or heat shocks—are likely to discretely shock production, and through it profits.

The sign of the profit response to forecasted weather is most likely dependent on the range of weather variation used for estimation, and on the composition of the agricultural industry studied. Here, these are bound to be specific to a French or European context. However, the mechanisms at play causing a non-zero profit response of adaptive behavior to forecasts can be generalized to other contexts: if timing responds to forecasts, and timing impacts production in a discrete way, the profit consequences of changing timing will not be well captured by the marginal approximations underlying the envelope conditions.

We further discuss the role of dynamic behavior on the part of farmers, namely storage, and dynamic weather impacts. We show that storage reacts positively to forecasted heating degree days, as farms seem to sell only part of the extra output coming from adaptation to heating degree days forecasts. We also show that the inclusion of lagged weather variables does not change the nature of our results.

This paper relates to a diverse and large economic literature which has relied on the envelope theorem. Extended to arbitrary choice sets by Milgrom and Segal (2002), and to integer and discrete choice variables by Sah and Zhao (1998), it has had applications in contract theory, game-theory, the study of non-convex production problems, and the study of robust comparative statics. The envelope theorem has also proven key to empirical analysis. Allowing for the clean derivation of estimating equations, it has guided the interpretation of reduced form results, and has helped to recover sufficient statistics linking structural models and reduced form identification methods. Its use in duality theory under the form of the Hotelling Lemma has, for example, grounded production function estimation and the recovery of optimal factor demands, as discussed in Taylor (1989). Envelope conditions also come into play for Euler equation methods used to estimate capital adjustment costs, following Hall (1978).

More recently, the envelope theorem has played a prominent role in the methods used to derive sufficient statistics in public economics, as discussed by Chetty (2009). There, envelope conditions are used to reduce the set of parameters that need to be recovered in order to identify the effect of policy schemes on welfare considerations. In other words, in cases where reduced form approaches cannot identify model primitives, the envelope theorem defines their welfare and policy-relevant aggregates, often in the form of elasticities, which can both be identified and are sufficient to draw conclusions about the consequences of the analysed policies.

As with the climate damages literature, the envelope theorem has been used to guide the interpretation of reduced form empirical results in the health and human capital literature. Bleakly (2010) uses the envelope theorem in a way similar to Hsiang (2016), in order to discuss how human capital investment coming from investment in early childhood health are netted out in experimental regressions of individual income on early childhood health. These are netted-out precisely because, at the margin and for an optimizing individual with a properly defined objective, the costs and

benefits of slightly adjusting investments in human capital cancel each other, leaving only the direct effect of improved health on income for a fixed level of human capital investment. In both the human capital and the climate damages cases, the envelope theorem is used to interpret the marginal response of an optimized object to some parameter change, and to justify that this observed marginal response is a correct measurement of the direct consequence of the parameter's variation on the objective. As such, the envelope theorem stands as a powerful tool to isolate causal channels in reduced form studies.

In a more casual literature but in a similar fashion, the envelope theorem has been used to outline the consequences of efficiency wages and monopsony power when large employers are forced to raise wages, as discussed in turn by Krugman (2015) and Dube (2015). As noted by Dube however, invoking the envelope theorem assumes that the observed changes in parameter are indeed marginal. For the case of climate damages, this implies that the identifying variation is indeed driven by marginal shocks in the data, on top of potential continuity concerns regarding the agent's policy functions. In the paper, we perform a decomposition of our weather variation into marginal and non-marginal variation. We highlight how most of the extreme heat variation in France since 1994 has actually been non-marginal. These data concerns add to the limitations related to the discontinuous consequences of certain production actions.

This paper is however first and foremost related to the literatures on climate adaptation and climate damages. This literature has long debated how to provide comprehensive estimate the consequences of climate change on human systems. This in turn has implied accounting for the costs of adaptation, or isolating measures of direct weather damages (and hence climate damages) from the effects of adaptation. Either of these approaches requires one to take a stance on the nature of the objective agents optimize, in order to then characterize adaptation. The agents' beliefs about the climate process structures what can be interpreted as adaptive behavior 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 set 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. Kala makes the point that modelling farmers' learning behavior can depend on our modelling of the objective that they maximize. Recovering the farmers' forecasting rule there is related to the choices made to define their objective. What we measure as the extent of their adaptation 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 agents' adaptive actions. Related to this literature, we show that here, not controlling for weather forecasts induces an upward bias in the measurement of the profit impact of extreme heat—extreme heat seems less damaging that it really is. We run a distributed lag model that includes all the forecasts available during the growing season (at all the possible lead values), and hence attempts to capture as much as possible of the agents' beliefs formation and adaptive behavior, and their consequences on farm profit. Running such a regression significantly increases the negative impact of extreme heat on French farm profits. This is suggestive evidence that as

we better control for the indirect consequences of adaptation on profit, we can better isolate the negative consequences of marginal variation in extreme heat.

Finally, we note that our work relates very closely to two papers which have approached climate damages without relying on the envelope theorem. The first of them is Guo and Costello (2013), which looks at the impact of climate change in a forestry context with both discrete and continuous decision making. In their context, foresters can implement continuous adaptation actions for which envelope conditions apply, and discrete actions for which they do not. They derive a model where marginal variation in weather can lead to arbitrarily large impacts on profit by inducing discrete changes in the type of tree species planted by foresters. Empirically, they find that extensive margin (discrete) adaptive actions can have large profit responses, but typically induce a net-zero profit response. Their framework relies on assumptions about replanting and harvesting cost values, and by-pass the frequent selection and endogeneity issues posed by joint continuous and discrete decision making in dynamic frameworks, as discussed in Bruneel-Zupanc (2022). 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 a simple production model and derive the implications of the envelope theorem in isolating the direct profit consequences of weather shocks. In section 3, we describe our data sources, and in section 4 we link our production model to the estimating equations used in our empirical strategy. In section 5, we go over the main results testing for the empirical applicability of the envelope theorem in the climate damages context. We decompose these results in section 6, discuss their consequences in section 7, and highlight their robustness in section 8. Finally, section 9 concludes the paper

2 Model

2.1 Derivations

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

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

Assuming the problem is well behaved, the first order condition will them imply that we

must have expected marginal revenues equal to expected marginal costs at the optimum. That is, to be at the optimum, we need that:

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

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

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

Using this value function, we can derive the impact of marginal variation in weather on optimal expected profits, where the second line follows form the envelope theorem:

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

This implies that empirically observing $\frac{\partial V(w_t)}{\partial w_t}$, the variation in optimal profit induced by weather, in a panel regression framework, is a sufficient statistic to infer about $\frac{\partial q}{\partial w_t}|_{x=x^\star}$, the net marginal impact of weather on production.

We argue that observing $\frac{\partial V(w_t)}{\partial f_t}$ is a sufficient statistic to test whether the envelope theorem is empirically applicable. Indeed, we note that because forecasts are purely informational, we have:

$$\frac{\partial V(w_t)}{\partial f_t}=\mathbb{E}_t\left[\frac{\partial x_t^\star}{\partial f_t}\big(p_t\frac{\partial q}{\partial x}|_{x=x^\star}-\frac{\partial c}{\partial x}|_{x=x^\star}\big)|w_{t-1},f_t,v_t\right] \tag{2}$$

If forecasts are a meaningful signal for agents, their variation will shift their beliefs about expected realized 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[4]:

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

Note that this implies that an absence of reaction to forecasts would not confirm the validity of the envelope conditions. However, the existence of a non-zero optimal profit response to forecasts will entail that envelope conditions do not hold. We then have that:

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

[4]We note that in a dynamic set-up, some model of capital adjustment costs imply regions of inaction. See for example Cooper and Haltiwanger (2006). The envelope theorem as used in the climate damages literature, however, is not conditioned on the existence of these inaction regions.

We propose Equation (2) as a test for the applicability of Equation (1) in an empirical context. Under the condition that forecasts are a meaningful signal to agents (conditional on the information they already possess), then the envelope theorem will not be applicable in a climate change context if the marginal adjustments to optimal profit with variation in forecasts are not null.

2.2 Why the Envelope Theorem might not apply

The envelope theorem relies on some form of differentiability of the value function, and only applies to marginal variations in the parameters of the optimization problem. As such, two avenues for its potential lack of applicability to study impacts of climate change are that: the contexts studied are not sufficiently well behaved, or that the empirical variation in weather used for estimation is non-marginal.

Milgrom and Segal (2002) discuss extensions to the classical versions of the envelope theorems for arbitrary choice sets. These extensions, however, still rely on either the absolute continuity of the objective function in the agent's action, or its equicontinuity at the initial parameter value at which the objective was optimized. It can be that such conditions do not apply to the problems at hand. In the case of agricultural production and weather, we can consider discrete timing decisions with a discrete impact on production. For example, reactions to freeze or extreme heat events require some discrete adjustments in timing which can prevent a complete loss in output.

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 help reduce the dimensionality of weather data. In the context of agriculture, this is necessary given that yields and output is typically observed at the agricultural season level, while weather is observed hourly. This aggregation requires at least two things. First that the impact of temperature is homogeneous across time. Here, bins, splines and GDDs/HDDs offer more flexibility than averages, but all rely on some for 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. The second is that to some degree marginal and non-marginal variation 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 to react differently to forecasts of these respective $4^{\circ}C$ and $1^{\circ}C$ shifts, and that the first one induces a non-marginal response to optimal production decisions pushing us beyond the bounds of applicability of the envelope theorem.

We investigate both of these avenues in section 6.2 of the paper.

3 Empirical Setting

3.1 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 choice which lowers the exposure of our data to irrigation, also reduces potential concerns of endogeneity linking irrigation and local weather, as discussed in Braun and Schlenker (2023), and also matches recommendations to avoid studying the impact of weather shocks in areas where irrigation is likely to be subsidized (Schlenker et al. (2005)).

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

Finally, in terms of exposure to weather shocks, France remains exposed to significant variation due to the frequent arrival of low-pressure areas from the Atlantic (Canal (2015)). Its Mediterranean regions also suffer from water stress over the summer as well as heavy localized rainfall. Seasonal forecasts have been developed for France since the 1990s, but quality coupled atmosphere-ocean models started to be used for seasonal forecasts only in 1999, allowing to move beyond a four months 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 with 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.[5]

[5]We note that the thresholds used for crop-specific tolerance are only indicative. Tolerance to heat varies across the growing stages of each crop, and also relate among else to drought conditions. We take these as only indicative that wheat is more heat tolerant than other crops, and is likely to be less responsive to extreme heat events observed for France in our sample. References are Tack et al. (2015), Schlenker and Roberts (2009), Shirdelmoghanloo et al. (2022), Elferjani and Soolanayakanahally (2018).

Table 1: Farm-Level Reactions - Decomposed Weather

	Share	Mean Temperature	Threshold	Reference
Wheat	47.89	10.2	34C	Tack et al, 2015
Corn	15.29	10.5	29C	Schlenker and Roberts, 2009
Barley (H)	11.38	10.0	30C	Shirdelmoghanloo et al, 2022
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 top 4 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[6].

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

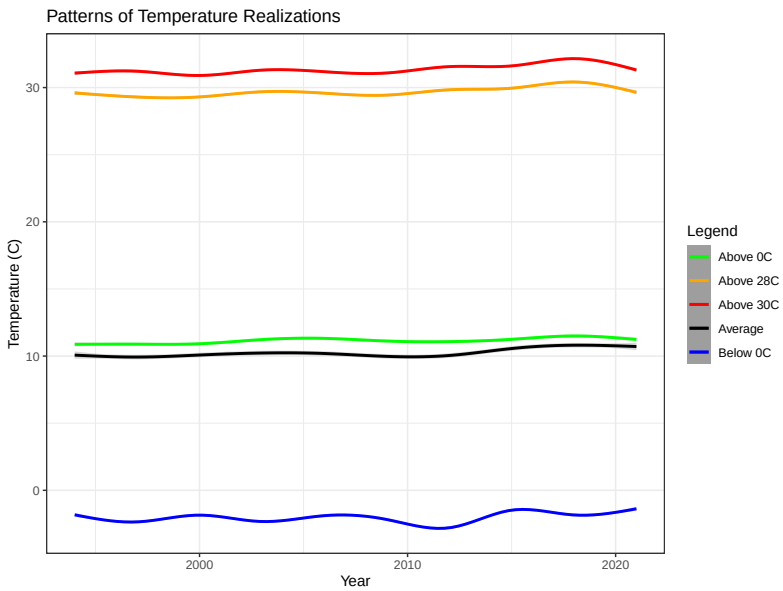


Figure 1: Unconditional and Conditional Average Temperature Realizations

Figure 1 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.

[6]These thresholds are only indicative, as crop-specific heat tolerance varies across growing stages.

Figure A3 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.

3.2 Agricultural Data

3.2.1 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. [7]

The FADN contains precise accounting data on the farm's activities. For our analysis, we focus on farms mainly producing cereal and oil crops [8]. 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 localise the farms at the department level, and match them with department-level weather data. While farms can be localized at the commune-level 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.

3.2.2 Input Price Data

Our input price data comes from three sources. From the FADN for hourly wages, 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 [9]. Each series is defined according to

[7] The FADN is not representative of all agriculture, but only of commercial farming. The definition of a commercial farm changed in 2010, but this only led to the changing of the the rules for choosing replacements for the farms leaving the sample, and not to an overhaul of the sample population itself. A commercial farm used to be defined as a farm with a unique manager, which sells more than half of their production, and whose manager's working hours corresponds 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 380k 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 (PBS or Production Brute Standard) of at least €25,000.

[8] 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.

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

the nature of the product sold, to its sales condition and the firm that sells it.[10] 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. Specifically we use the Ipampa price index series, from 1994 to 2020.

3.2.3 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.

3.2.4 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.

3.3 Climate Data

3.3.1 ERA 5

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, we aggregate the hourly data into growing-season observations [11]: growing degree days (GDD) and heat degree days (HDD) for temperature[12], and total precipitation for rainfall. 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.

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

[11] We use an extensive definition of the agricultural growing season for France, running from October of the previous year, to July of the current one.

[12] Growing degree days are computed over the $[4°, 30°]$ degree interval, and heat degree days sum the realized temperature above 30°C.

3.3.2 SEAS 5

The forecast data is taken from ECMWF's SEAS5 seasonal forecasting system. Forecast are produced on the first of each month for the following 5,160 hours [13]. For temperature, they are produced at a 6h interval, and give an instantaneous prediction of temperature, while for rainfall, they give the accumulation of rainfall every 24h. As such, the rainfall forecasts for the second day of the month will be the following: a 24h ahead forecast, a 30 to 32 days ahead forecast produced on the first day of the previous month, 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. We would ideally like to build the forecast-equivalents of our growing-season aggregates for realized weather, for different 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 (our one month lead for the rest of the paper), produced two calendar months ago, up to five months ago. As such, the forecasts that we aggregate into growing season observations are not homogeneous in terms of lead value, but are the closest equivalent of it that we can get.

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

3.3.3 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 mostly centered around zero with an upward bias in the forecast and a long tail associated with it. Rainfall also shows an upward bias in the forecasts, with distributions centered close to zero.

For the rest of the paper, we remove the observations situated in the lower tail of the temperature error distribution (error over -500 GDD), removing 120 observations in total.

We also note the distributions indicate a constant bias in forecasts (curves are shifted away from zero), but a relatively narrow distribution. This confirms that forecasts do contain meaningful information about future weather, and should shift farmers' beliefs about weather in a significant way.

4 Empirical Strategy

We propose to test whether marginal adaptation behavior in response to weather forecasts have null impacts on profit, or that:

$$\frac{\partial V(w_t)}{\partial f_t} = 0$$

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

Because we observe proxies of profit at the establishment level, and can link farms with both realized and forecasted weather, this test is quite straightforward. Using measures of cost and revenue - still at the establishment level - we can then detail this test further by looking at:

$$\mathbb{E}_t\left[p_t \frac{\partial x_t^\star}{\partial f_t} \frac{\partial q}{\partial x}|_{x=x^\star}|w_{t-1}, f_t, v_t\right] \quad ; \quad \mathbb{E}_t\left[\frac{\partial c}{\partial x}|_{x=x^\star}|w_{t-1}, f_t, v_t\right]$$

We can run the following fixed effects regression, for farm j in growing season t, located in department d:

$$Y_{jt} = \beta^w Weather_{d(j)t} + \beta^f Forecast_{d(j)t} + \eta_j + \eta_t + \varepsilon_{jt}$$

Controlling for farm and year fixed effects should residualize both forecasted and realized weather into random shocks, which allows for the identification of the β^f coefficient. In our context, $\beta^f = 0$, when Y_{jt} corresponds to farm profit, would imply a net-zero impact of forecasts on farm profit.

5 Results

5.1 Main Results

Table 2 describes the results from our main specification. We use two different definitions of profit, each shown in a different column: profit defined as the gross operating income of the farm, and value added measured as the difference between total production value and intermediary inputs. The main explanatory variables are listed in the left column, and fixed effects and control variables are described below the coefficient estimates. The coefficients of interest here are the ones corresponding to the one month ahead heating degree days (HDD) forecasts, and to one month ahead growing degree days (GDD) forecasts. All the regressions have farm and year fixed effects, and we control for both realized and forecasted rainfall (level and square), and realized GDDs and HDDs. Standard errors are clustered at the department level, which broadly correspond to counties in the US context. We compute our weather data at the department level, which justified this level of clustering. Observations are weighted using the sample weights provided in the FADN.

Table 2: Profit Reaction to Forecasts (1 month lead)

Dependent Variables:	Value Added	Profit
Model:	(1)	(2)
Variables		
GDD	4.054	-3.128
	(8.998)	(9.714)
GDD (F)	22.24	11.94
	(16.64)	(18.00)
HDD	3.911	24.03
	(246.8)	(232.4)
HDD (F)	2,430.2***	2,005.2***
	(657.5)	(559.4)
Fixed-effects		
Farm	Yes	Yes
Year	Yes	Yes
Fit statistics		
Observations	18,109	18,109
R^2	0.80318	0.83962

One-way (department) standard-errors in parentheses.
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes. Realized and forecasted rainfall in levels and squares are included as controls.

The table shows that HDD forecasts have a significant positive effect on farm profits, across definitions of profits. A one unit increase in forecasted HDD will raise farm profits by $2,420$ €, when using value added as a measure of profit. When using the narrower second measure, the effect is slightly lower, of $2,005$ €. On the contrary, the effect of forecasted GDDs on profits is non-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. Or that at the optimum, the re-optimization of production schedules to respond to changes in extreme weather have a non-zero profit effect. 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. It is interesting to note that reactions to forecasted GDDs and HDDs are very distinct, and we will later highlight how the nature of the variation for GDDs and HDDs in our data might partially explain these differences. We can also conjecture that HDDs are by definition more extreme events which might require more adaptation than GDDs.

We next show the marginal response of profits to different HDD leads - noting that further-ahead leads might allow for more adaptation, while being also less precise. We also note that forecast leads of more than four months only became available in France in 1999, as discussed in section 3. As such, we should also expect less reactions to the hindcasts calculated with a 5 month ahead lead.

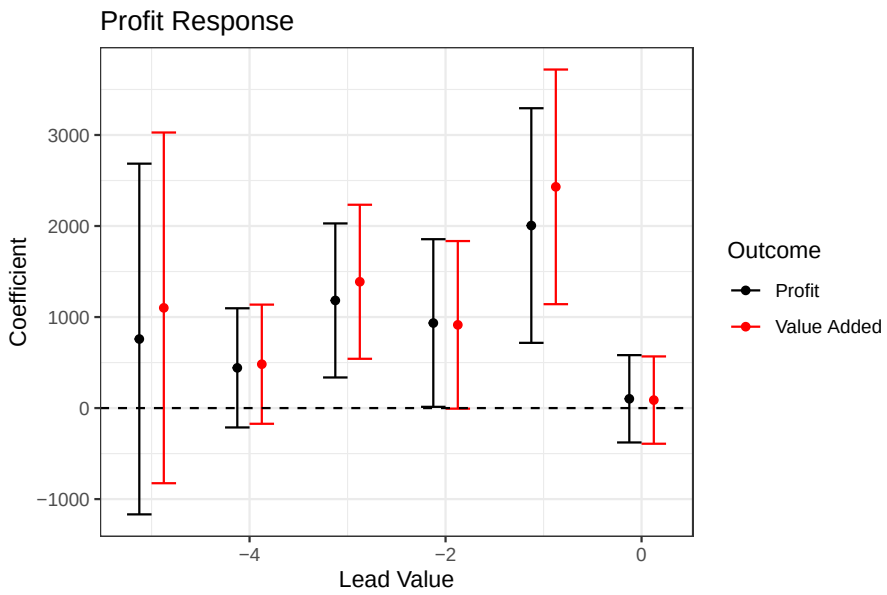


Figure 2: Profit Reactions

The regressions used to produce this graph are identical to the ones shown in Table 2, but for the variation in the lead time used for the forecasted rainfall (level and square), GDDs and HDDs. This means that we include realized weather as controls, as well as department and year fixed effects, and cluster the standard errors at the department level to account for spatial correlation in weather and agricultural patterns. We display the coefficient values associated to the forecasted HDDs. As such, the value displayed for the lead value of 1 corresponds to the one displayed in Table 2. Finally, we indicate lead values 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 result shown for the 0 lead value corresponds to a regression including only realized weather, without the corresponding forecasted values. All the coefficient values for leads ranging from 1 to 3 months ahead are statistically different from zero and positive. The coefficients for lead values 4 and 5 are still positive but lose in precision. This can account both for the lower precision of forecasts with such a lead time, or the timing of adaptive actions. We take the robustness of the positive significant coefficient attached to forecasted HDD as supporting the fact that adaptive behavior to forecasted extreme heat events has a non-zero profit effect.

As a validation check for the nature of our data, and the behavior of the farmers in our sample, we include in the annex two tables describing farm responses to realized weather only - mirroring the typical panel regression studied in the climate damages literature. Table A7 describes

the reaction of corn producers to realized weather in France for the period 1994-2018. We show the results for corn, as corn is a crop known for its sensibility to weather. There, we see that indeed corn production (in volume) is negatively affected by HDDs, and in a statistically significant way. A one unit increase in HDD decreases corn production by 21.36 cwt. In Table A8 we show more general farm-level outcomes. There, we see that on average realized HDDs do no impact farm profit in a significant way in France.

5.2 Cost vs Revenue Responses

Next, we move to decomposing the profit response to forecasted GDDs and HDDs into cost and revenue effects. We observed that forecasted GDDs have a net-zero effect on profits, while forecasted HDDs have a net positive impact on them. The envelope theorem is a statement on that net effect, but allows for the adaptive response to have effects on costs and revenues that cancel each other. As such, the net-zero impact of forecasted GDD might still come with adaptive behaviors - something interesting in itself. It is also interesting to identify the source of the positive profit effect of forecasted HDDs, which might stem from reductions in costs, or from increases in revenues.

Table 3 displays regressions using a similar design as those in Table 2, but with different outcomes. These outcomes are respectively one measure of revenues (total sales), and two measures of costs of production. The first one, total costs, is meant to be expansive. It includes the sum of intermediary inputs, social contributions to workers, personnel expenses, taxes and insurance. The second one only includes expenses on intermediary inputs, and might be better measured in that respect.

These results are useful to gain insights about the type of adaptation behavior which generates the positive profit responses to forecasted HDD and GDD seen before. Indeed, we could expect that 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 the harvest. In a simple model where local adverse weather corresponds to a negative TFP shock, but output prices are not responsive to local conditions, we would expect forecasts to reduce the scale of production, leading to a negative response in revenues and to reductions in costs. On the contrary, changes in the timing of the harvest could allow for a positive revenue response, and a neutral cost one. Finally, if weather shocks are input-biased in their effects, or if they are not, but if some of the inputs are lumpy in nature, then a weather shock could impact the relative input shares used in production.

Regarding forecasted GDDs, we could expect that the net-zero effect on profits either comes from a lack of adaptive reaction to them, or to marginal responses for which cost and revenue variations net each other out.

Table 3: Cost and Revenue Reactions to Forecasts (1 month lead),

Dependent Variables:	Sales	Total Costs	Intermediate Inputs
Model:	(1)	(2)	(3)
Variables			
GDD	1.714	-4.223	-7.969**
	(9.203)	(5.858)	(3.841)
GDD (F)	49.94***	33.91***	23.22***
	(15.14)	(10.03)	(7.929)
HDD	-168.0	-159.6	-113.4
	(233.4)	(117.5)	(87.31)
HDD (F)	1,329.7**	35.16	11.07
	(652.8)	(338.6)	(247.0)
Fixed-effects			
Farm	Yes	Yes	Yes
Year	Yes	Yes	Yes
Fit statistics			
Observations	18,109	18,109	18,109
R^2	0.88616	0.93581	0.93301

One-way (department) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes. Realized and forecasted rainfall in levels and squares are included as controls.

Here, we see that revenues respond positively to forecasted weather. An increase by one degree-hour in HDD forecasted one month in advance leads to an increase in sales by 1,329 €. On the contrary, a similar forecast has no statistically significant effect on production costs, measured with either of our two variables. This table hence highlights that positive profit responses to forecasted HDDs are driven by sales, rather than costs. This positive sales respond might in turn correspond to different realities: increases 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 investigate this further in the next section.

We also note that forecasted GDDs impact both sales and production costs, however in such a way that the two effects cancel each other out as predicted by the envelope theorem. Here, forecasted GDDs raise sales, and also raise production costs in a slightly lower fashion, albeit with the same order of magnitude.

We next perform the same exercise as with profits, and show the variation in cost and revenue responses to forecasted HDDs as we vary the lead value used for these forecasts. The regressions are similar to the ones displayed in Table 3.

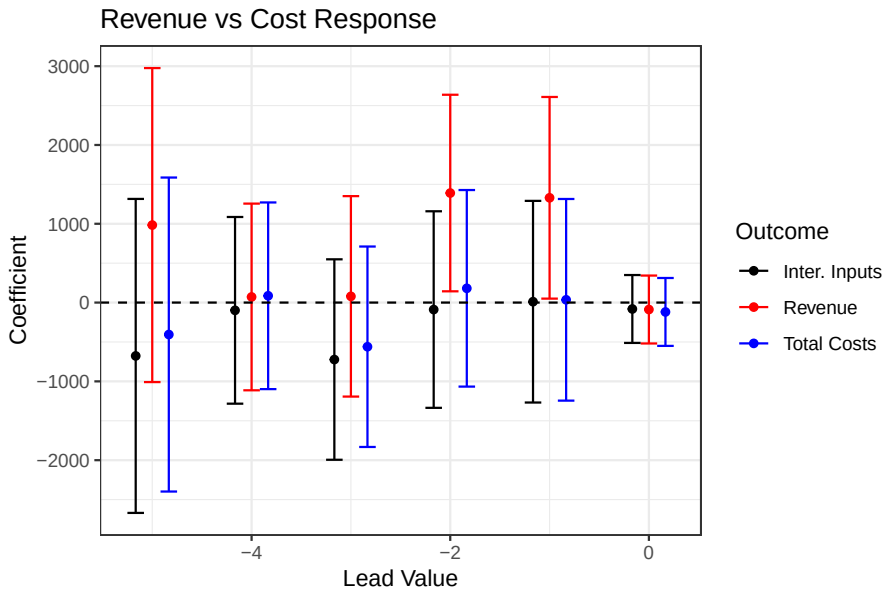


Figure 3: Cost and Revenue Reactions

We see that that the positive revenue responses to forecasted HDDs holds for leads one and two months ahead, and loses statistical significance as we move to leads further ahead in time. Costs responses to forecasted HDDs remain non-statistically significant.

5.3 Timing of Adaptation

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

$$Y_{jt} = \beta^w Weather_{d(j)t} + \sum_{l \in [1,5]} \beta_l^f Forecast_{d(j)t}^l + \eta_j + \eta_t + \varepsilon_{jt}$$

This regression aims at getting at 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 ommitted variable bias which associates to our unique set of forecast coefficients a composite of the effect of the different forecasts. As long as we control for realized weather, this omitted variable will not be an issue for our test, and for the interpretation of the forecast coefficients as representative of the agent's reaction to the arrival of new weather information via the receiving of forecasts. 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 larger profit response.

Here, each β_l^f indicates the profit response to the arrival of forecasted weather information at lead l, holding both the information received at other periods and the weather realizations constant. We note that 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 A2 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.

We plot below the coefficients associated with heating degree days: β_l^f for $l \in [1, 5]$, and show β^w for the lead value of 0.

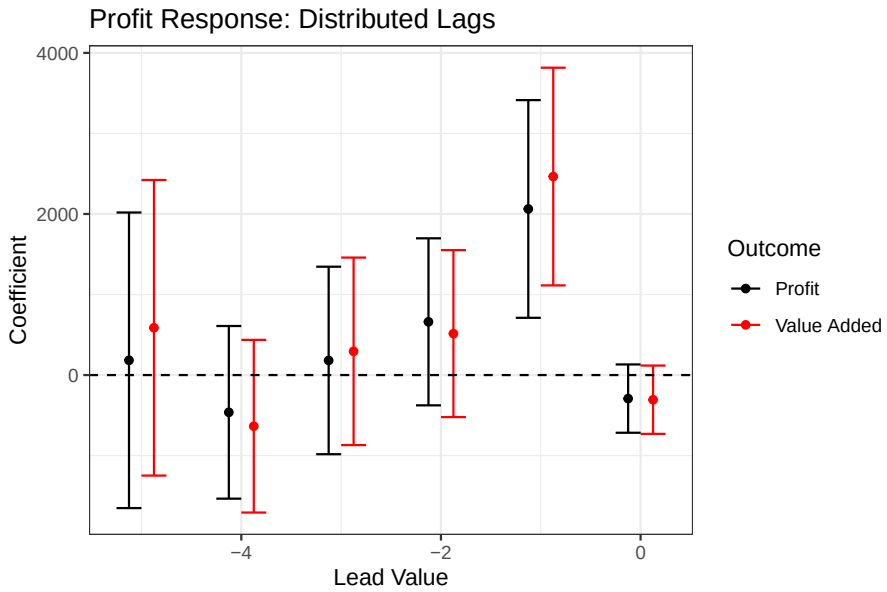


Figure 4: Cost and Revenue Reactions

Contrary to Figure 2, we see that only the coefficient associated to forecasts with a one month lead is statistically significant. Other coefficients decrease in value, and become more centred around 0 as we more further back in time. This is interesting, and should indicate that most of the profit-relevant adaptation observed in our sample is done in the month leading up to the weather events.

We also observe that the coefficients associated to realizations become more negative and almost statistically significant at the 5% level. This might mean that including more leads allows to better isolate the direct effect of extreme heat on agricultural profits.

6 Where the envelope theorem breaks

In the coming section, we explore the potential causes that could explain the net positive profit response to forecasted HDDs. We first decompose the positive revenue response highlighted in Table 3, and show that this positive revenue response seems to come mostly from a positive response of production volumes to forecasted HDDs, rather than from changes in storage behavior, or changes in output prices. We also show that the null response of production costs to forecasted HDDs is not driven by a joint reaction in input prices and input volumes - both are unresponsive to forecasted extreme weather.

We then analyse a potential source of zero-cost adjustments to production schedules which

might explain how farms can use information about incoming heat without significantly changing their costs of production. Specifically, using plot-level data in a repeated cross-section, we show how the timing of the major steps of the growing season (first irrigation, harvest and sowing) all respond to forecasted HDDs. We can expect these changes in timing to have significant effects on production, while being relatively costless - under the condition that input prices do not vary in a significant way across the windows of time considered (e.g. irrigating for the first time at date t, or two weeks later costs the same). It is also likely than forecasted extreme events warrant larger responses in timing that forecasted variations in growing degree days.

Finally, we analyse another potential explanation for the differential response to heating and growing degree days. We decompose the source of variation in growing and heating degree days into marginal and non-marginal variation, and show that most of the variation in our GDD measure is marginal, while most of the variation in our HDD measure is non-marginal. Being that the envelope theorem applies to marginal variations in the optimisation problem's parameters, this could also explain why the adaptive responses to HDDs have non-marginal impacts on profit.

6.1 Decomposing the Cost and Revenue Responses to HDDs

6.1.1 Revenue Responses: Quantity Produced, Quantity Stored and Output Prices

In Table 4 we run regressions with a similar design as previously, but focus on output quantities. Given that farms are multi-product firms, we use different variables to highlight the response of production volume to forecasted HDDs. We build two aggregate measures which aim at getting at the farm's aggregate production across product lines. The first measure we use is an output index, which attempts at constructing a weighted sum of production volumes. For farm j in period t, denoting l_{jct} the land allocated to crop c, and $\mathbb{C}_{jt}$ the farm's crop mix, the output index is:

$$q^i_{jt} = \sum_{c \in \mathbb{C}_{jt}} \frac{l_{jct}}{\sum_{c' \in \mathbb{C}_{jt}} l_{jct'}} q_{jct} \tag{3}$$

The index will respond both to increases/decreases in production at the crop level q_{jct}, and to variations in the land allocated to crops. As such, the index gives an idea of directed changes in production - whether crops allocated a larger share of land see their output increase.

The alternative to this index is to simply sum output quantities across crops. Our second farm-level quantity outcome is thus a simple sum. We build these two variables for 19 cereal, oil and protein crops[14]

Finally, we run the regression for crop-specific output volumes for the three most common crops within the cereal, oil and protein categories in France: wheat, corn and rapeseed.

[14]The crops accounted for are: wheat, durum wheat, oats, corn, corn grown for seeds, sorghum, spring and winter barley, rye, triticale, summer cereals, other cereals, sunflower, rapeseed, soy, dry peas, feverole, beans and protein peas.

Table 4: Quantity Reactions to Forecasts (1 month lead)

Dependent Variables:	Output Index	Output (sum)	Wheat	Corn	Rapeseed
Model:	(1)	(2)	(3)	(4)	(5)
Variables					
GDD	0.0054	-0.2456	-0.6113	0.4370	-0.4720***
	(0.3634)	(0.5144)	(0.3777)	(0.3188)	(0.1388)
GDD (F)	0.1559	1.801*	0.6885	0.4545	0.7710*
	(0.7008)	(0.9522)	(0.6336)	(0.7532)	(0.3953)
HDD	-5.384	-7.776	-0.3035	-18.52	10.24***
	(8.173)	(12.32)	(7.666)	(12.29)	(3.232)
HDD (F)	-9.052	74.57***	65.93***	-13.91	6.541
	(11.45)	(18.91)	(16.28)	(17.00)	(10.98)
Fixed-effects					
Farm	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes
Fit statistics					
Observations	18,109	18,109	16,613	9,763	11,913
R^2	0.83169	0.94131	0.90151	0.93466	0.79922

One-way (department) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes. Realized and forecasted rainfall in levels and squares are included as controls.

Our focus is on the coefficients associated to forecasted GDDs and HDDs. We first note that the output index has an inprecise zero, while the output sum is positive and statistically significant. This indicates that the positive revenue response to forecasted HDDs observed in Table 4 is at least in part driven by a positive response of production volumes. The crop-specific output quantities shows that the main driver of that positive response seems to be wheat, a crop which is less affected by heat shocks than corn and colza. We also note that there is no such strong positive response of output, either at the farm-level, or at the farm-crop level, to forecasted GDDs. There is a slightly positive response but of a far smaller order of magnitude, which here seems to be driven in part by colza. Overall, this table indicates that the positive profit response to forecasted HDDs is actually driven by a direct change in production schedules that increases output.

Next, we confirm this role of output response by looking at output prices. We build an output price index at the farm-level, in a similar fashion as for the output quantity index, and for the same set of crops. We also include regressions for the output prices of our three main crops: wheat, corn and colza. We use log prices rather than levels, as we are interested in proportional changes to prices.

Table 5: Output Prices (log) Reactions to Forecasts (1 month lead)

Dependent Variables:	Output Price Index	Wheat	Corn	Rapeseed
Model:	(1)	(2)	(3)	(4)
Variables				
GDD	-0.0005**	0	0	0.0001**
	(0.0002)	(0)	(0)	(0)
GDD (F)	0.0009***	-0.0003***	0	0
	(0.0003)	(0)	(0.0001)	(0.0001)
HDD	0.0099**	-0.0004	0.0047***	-0.0008
	(0.0047)	(0.0010)	(0.0017)	(0.0013)
HDD (F)	0.0005	0.0014	-0.0002	-0.0005
	(0.0046)	(0.0024)	(0.0021)	(0.0032)
Fixed-effects				
Farm	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Fit statistics				
Observations	18,917	17,181	9,958	12,037
R^2	0.59231	0.78829	0.73969	0.85520

One-way (department) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes. Realized and forecasted rainfall in levels and squares are included as controls.

We see that there is no statistically significant response of output prices to forecasted HDDs - which might indicate that agricultural markets are too large, and output prices determined at a level aggregate enough that local weather shock do not impact output prices too much. We note however that forecasted GDDs have a negative impact on wheat prices, although we do not see any positive response of wheat output to them. We also note that realized HDDs for corn have a positive impact on prices, while corn volumes do respond negatively to them but not in a significant way. This may indicate that we lack sufficient variation to identify these joint effects in a significant way. However, these output price responses seem to be indicative that the net positive response to forecasted HDDs is not driven by an positive price response.

Finally, we look at stored output quantity, to see whether variation in storage behavior might complement the production reactions discussed earlier. Again, we run regressions with the same design, and as for output quantity, build two measures of farm-level aggregate storage: an index built in the same way as the output quantity index, and the output price index, and a sum of total storage at the farm-level.

Table 6: Output Storage Reactions to Forecasts (1 month lead),

Dependent Variables:	Storage Index	Storage (sum)	Wheat	Corn	Rapeseed
Model:	(1)	(2)	(3)	(4)	(5)
Variables					
GDD	-0.4745*	-0.9174**	-0.8195***	-0.1380	-0.0077
	(0.2568)	(0.3614)	(0.3008)	(0.3709)	(0.1251)
GDD (F)	-0.5780	-0.9219	-1.421*	-0.3441	0.1454
	(0.7822)	(0.9375)	(0.7936)	(0.9053)	(0.4126)
HDD	-14.41***	0.7910	6.656	-10.29	6.506**
	(4.836)	(8.841)	(9.310)	(7.544)	(2.597)
HDD (F)	10.24	38.55*	23.68	10.99	11.75*
	(11.26)	(20.63)	(15.41)	(15.45)	(6.333)
Fixed-effects					
Farm	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes
Fit statistics					
Observations	15,421	18,109	16,613	9,763	11,913
R^2	0.09463	0.10748	0.10103	0.06793	0.05912

One-way (department) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes. Realized and forecasted rainfall in levels and squares are included as controls.

Here we note that forecasted HDDs actually have a positive response on the sum of stored output at the farm level, which is indicative that the extra output produced from adaptive responses to forecasted HDDs is partly stored, and partly sold. Here, the storage behavior seems to be driven in part by rapeseed. This positive storage response to forecasted HDDs means that storage actually dampens the immediate positive profit response to forecasted HDDs observed in Table 2.

To conclude this section, we note that the positive revenue response to forecasted HDDs seems to be driven by a direct response of production volumes, neither amplified or mitigated by variations in output prices, and partly dampened by concomitant increases in storage.

6.1.2 Cost Responses: Input Levels and Input Prices

In the following section, we perform a similar exercise for the cost side of production. We decompose the null response of production costs to forecasted HDDs into an input level and input price response.

We first look at input volumes. The regressions follow the same two-way fixed effects framework as in the previous sections, and we run them on the same sample of farms. We include six inputs. The first two are directly observed in volume: total agricultural area used (measured

in ares), total labor hours (including paid and non-paid hours, from family members, the farm manager and from employees). The last four correspond to deflated input bills. We use the country-wide input-specific Laspeyres indices computed by the French statistical agency (INSEE) in the Ipampa series to deflate these bills into quantities. We note that these price indices do not possess any geographic variation, because the original agricultural input price study run in 1995 to prepare the establishment of an agricultural input price survey concluded that prices were sufficiently homogeneous across France for country-wide indices to be sufficient. We also note that all these inputs are directly measured at the farm-level. While we use input price data later, we do not use these input prices to deflate farm bills because this survey is not stratified in order to be representative of geographic variation in prices. The actual geographic variation observed in the data corresponds both to product-constant variation in prices, and to variation in the products sold matching geographic variation in cropping patterns (the North of France for example grows beetroots which do not use the same pesticides as cereals, grown mostly in the center of the country). As such, we do not believe that this survey data would allow us to extract significant additional variation compared to the Laspeyres indices built by INSEE on that same price survey.

Regarding forecasted HDDs, we only observe two statistically significant responses: a positive response of land, and a negative response of irrigation. The irrigation response is expected, as irrigating before a heat shock is not recommended. The land response is more unexpected, but seems to indicate that when farmers expect the season to have relatively worse weather, they compensate this expected decrease in yields by increasing their scale of production. Here, a one unit increase in forecasted HDDs raises total used agricultural land by 26 ares.

We also note the positive response of fertilizers to forecasted GDDs. This is interesting as it indicates that there are complementarities between expected good weather and the use of yield-enhancing inputs. A similar positive response is observed for labor hours, and while they are not statistically significant, land and seeds also react positively.

Table 7: Input Quantities Reactions to Forecasts (1 month lead)

Dependent Variables:	Surface	Labor	Fertilizer	Phytosanitary	Seeds	Irrigation
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
GDD	-0.2363*	-0.0795	-3.556*	-2.403***	-0.1699	0.4835*
	(0.1389)	(0.1830)	(1.936)	(0.8928)	(1.157)	(0.2856)
GDD (F)	0.2015	0.5390*	7.613*	0.9767	0.6364	-0.6536
	(0.3969)	(0.2801)	(4.074)	(2.248)	(1.163)	(0.7406)
HDD	-1.697	0.0089	-77.52**	-7.306	-5.502	10.90
	(5.069)	(5.195)	(36.86)	(24.71)	(23.33)	(7.009)
HDD (F)	26.15*	0.9919	18.53	67.10	-23.73	-32.49***
	(15.69)	(7.058)	(74.90)	(83.30)	(45.34)	(11.89)
Fixed-effects						
Farm	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
Observations	18,109	18,109	18,109	18,109	18,109	18,109
R^2	0.71177	0.86123	0.89514	0.92311	0.87838	0.86350

One-way (department) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes. Realized and forecasted rainfall in levels and squares are included as controls.

Overall, we observe two countervailing input quantity reactions: an increase in land use and a decrease in irrigation. The net zero response of production costs observed earlier could then indicate that these two effects cancel each other. It is also possible that increases in land use are not so costly, if farms use more of the land they already own (the cost would then only correspond to extra fuel expenses), while the variation in irrigation itself could also only have a small impact on costs if electricity is cheap.

We conclude this decomposition section by looking at input prices. Here, we recover land prices from a land price survey, and fertilizer, pesticides and seed prices from the agricultural input price survey (EPCIA) run every year in France. For these, our data is at the product series, store and year level. A product series here corresponds to a specific brand, product type and product packaging (quantity). The regressions run for store prices are run with department, product and year fixed effects, and standard errors are clustered at the department level - which is the level at which we build our weather data.

Table 8: Input Price (log) Reactions to Forecasts (1 month lead)

Dependent Variables:	All Store Prices	Fertilizer	Phytosanitary	Seeds	Land
Model:	(1)	(2)	(3)	(4)	(5)
Variables					
GDD	0	0	-0.0001	0	0
	(0)	(0.0001)	(0.0001)	(0)	(0)
GDD (F)	0	0.0003	-0.0001	0.0001	0
	(0.0001)	(0.0002)	(0.0002)	(0.0002)	(0.0001)
HDD	-0.0030	-0.0032	-0.0020	-0.0006	-0.0020
	(0.0023)	(0.0027)	(0.0030)	(0.0027)	(0.0017)
HDD (F)	-0.0071	-0.0015	-0.0094	0.0079*	0
	(0.0072)	(0.0052)	(0.0128)	(0.0045)	(0.0018)
Fixed-effects					
Department-Firm	Yes	Yes	Yes	Yes	
Product	Yes	Yes	Yes	Yes	
Year	Yes	Yes	Yes	Yes	Yes
Farm					Yes
Fit statistics					
Standard-Errors	Department	Department	Department	Department	Department
Observations	33,174	9,258	17,883	6,033	5,716
R^2	0.87603	0.64152	0.67052	0.71827	0.97591

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes. Realized and forecasted rainfall in levels and squares are included as controls.

These results indicate an absence of input price responses to forecasted HDDs, or to almost any weather realizations. The only positive response observed is that of land prices to forecasted HDDs, but is negligible in terms of magnitude.

To conclude this section, it seems that overall the net zero response of production costs is driven by a lack of response in input quantities - for which prices remain constant. It then remains to explain why we observe a positive production response to forecasted HDDs without a significant change in input levels. We show next that other adaptation margins might also play a role in driving this response.

6.2 Sources of Adjustment

For this next section, we use data gathered at the field level from the survey of agricultural practices (PKGC). This survey is a repeated cross-section, and as such we are not able to include farm fixed effects. We only include department fixed effects (the smallest geographic identity to which we can attach the fields), as well as year fixed effects. We cluster standard errors at the department level,

and include realized and forecasted rainfall as a control (in levels and squares).

We focus specifically on timing questions, as the timing of different production procedures is likely to react to forecasted weather. 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 price.

We observe that the date of the first irrigation, the date of harvest and the date of sowing are all pushed back (come later in time) as a reaction to forecasted HDDs. These adaptive behaviors are likely to impact production volumes, and hence play a role in the positive profit respond, while not impacting costs. Timing can typically be thought of as the optimization of a discrete variable, with a discrete impact on production volumes - and hence as a part of the production schedule which is not well suited to the use of the envelope theorem. We note that this change in the timing of the first irrigation might also be linked to the observed decrease in irrigation as a response to forecasted HDDs that we discussed in the previous section. We also observe that all the procedures listed respond to realized HDDs: the harvest happens earlier as a response to extreme heat shocks, the first irrigation happens later (perhaps with the goal of saving water resources, and timing them closer to heat shocks), sowing is done earlier, and ploughing is in accordance also done earlier.

Table 9: Timing Response to Forecasts (1 month lead)

Dependent Variables:	Ploughing	First Irrigation	Harvest	Sowing
Model:	(1)	(2)	(3)	(4)
Variables				
GDD	-0.0378*	0.0100	-0.0449	-0.0380
	(0.0218)	(0.0153)	(0.0357)	(0.0233)
GDD (F)	0.0108	-0.0804***	-0.0173	0.0198
	(0.0205)	(0.0157)	(0.0297)	(0.0202)
HDD	-0.6973***	0.3424**	-0.5287***	-0.5976***
	(0.2089)	(0.1477)	(0.1614)	(0.1283)
HDD (F)	0.1958	0.7348**	0.7869***	0.4816**
	(0.3120)	(0.2820)	(0.2641)	(0.2069)
Fixed-effects				
Department	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Fit statistics				
Observations	44,784	44,784	44,784	44,784
R^2	0.44831	0.96153	0.60822	0.50481

One-way (Department) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes. Realized and forecasted rainfall in levels and squares are included as controls.

Finally, in the next section we discuss a more practical reason which might lead to the limited empirical relevance of the envelope theorem to study the impacts of climate change: the fact that the aggregation of temperature realizations into GDDs and HDDs bundles together marginal and non-marginal weather variation, while the envelope theorem is only meant to apply to marginal shocks.

6.3 Sources of Weather Variation

For this section, we decompose the variation in temperature in our data into marginal and non-marginal variation, for our two temperature variables (resp. GDDs and HDDs). Our decomposition is done according to the following procedure: first, we compute a thirty years average of hourly temperature realizations at the department level. That is, we aggregate the gridded ERA5 data for 1963-1993 at the department level using area-weights. We then compute the average temperature value in degree Celsius for each hour-day-month tuple. Using that same data, we also build the temperature standard deviation for each hour-day-month tuple.

For each hourly department-level temperature realization in 1994-2020, 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.

We then compute GDD and HDD values for each growing season, using respectively the marginal and non-marginal realizations.

We first show how most of the variation in the GDD variable is marginal, while most of the variation in the HDD variable is non-marginal.

Table 10: 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.

Next we show the correlations between the different measures of GDD and HDD at the department year level. We clearly see a low correlation between marginal and non-marginal shocks, both for GDDs (.16 between marginal and non-marginal GDDs) and HDDs (.25 between marginal and non-marginal HDDs). Here, observations are department and growing season level realizations - and the correlation hence accounts for correlation within space and season in the magnitude of the respective variations. We see that indeed, the correlation within marginal shocks, and within non-marginal shocks is stronger than across them within GDD or HDD. Non-marginal GDD shocks have a .60 correlation with non-marginal HDD shocks, while marginal GDD realizations have a .34 correlation with marginal HDDs. This is to be expected as heat shocks are correlated across time and should lift temperatures across the day, moving GDD realizations non-marginally at low temperature hours, and moving HDD realizations non-marginally in the hottest hours around noon and early afternoon.

Table 11: Correlation in Decomposed Weather Realizations

	GDD non-marginal	GDD marginal	HDD non-marginal	HDD marginal
GDD non-marginal	1.00	0.16	0.60	0.20
GDD marginal	0.16	1.00	0.13	0.34
HDD non-marginal	0.60	0.13	1.00	0.25
HDD marginal	0.20	0.34	0.25	1.00

Notes. The table indicates the correlation between the different measures of HDD and GDD for department-year realizations over 1994-2018.

While GDD realizations show some correlation between marginal and non-marginal shocks, there is none for the HDD shocks. By definition, the HDD measure is meant to account for extreme events, and this non-marginal nature of HDDs is thus to be expected. However, because the impact of extreme heat events is a central focus to the climate change literature, the non-marginal nature of these extreme shocks is an issue if they are to be studied using the envelope theorem.

Finally, we run our main specification but on decomposed realized and forecasted temperature variables, rather than their aggregate counterparts. We show the results in the following table. The first point is that only non-marginal HDD forecasts induce a positive sales response. Noting that the support for marginal HDD is small, this is either indicative that farmers do not respond to marginal variation in HDD forecasts, or that our data does not contain sufficient variation to capture this response. Second, we note that farmers also respond to non-marginal GDD forecast, 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, indicating a profit response close to zero.

Table 12: Farm-Level Reactions - Decomposed Weather

Dependent Variables:	Sales	Total Costs	Intermediate Inputs
Model:	(1)	(2)	(3)
Variables			
GDD (m.)	12.90	12.14	0.1082
	(9.841)	(9.022)	(4.943)
GDD (m - F)	33.97	26.93	8.434
	(23.65)	(16.72)	(10.78)
GDD (nm)	5.876	-3.610	-7.518
	(11.54)	(8.405)	(5.565)
GDD (nm - F)	17.82	38.18***	22.50**
	(25.19)	(13.94)	(10.91)
HDD (m)	-298.9	-90.00	-167.3
	(331.2)	(149.9)	(113.6)
HDD (m - F)	-1,787.1	-131.5	-333.9
	(1,449.5)	(916.9)	(547.5)
HDD (nm)	-810.9	-46.07	11.36
	(909.7)	(734.4)	(580.0)
HDD (nm - F)	2,060.8**	167.2	-94.08
	(982.0)	(575.2)	(423.7)
Fixed-effects			
Farm	Yes	Yes	Yes
Year	Yes	Yes	Yes
Fit statistics			
Observations	17,584	17,584	17,584
R^2	0.88539	0.93628	0.93345
Within R^2	0.00208	0.00225	0.00221

One-way (department) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes. Realized and forecasted rainfall in levels and squares are included as controls. Here (m) indicates marginal weather, (nm) stands for non-marginal, and the F indicates one month ahead seasonal forecasts.

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

7 Consequences

We can gain some insight regarding the consequences of the lack of empirical applicability of the envelope theorem in this climate context, by comparing profit responses to HDDs and GDDs when we include or omit weather forecasts as partial controls for the farmers' adaptation behavior. Note that this test is only informative, as it is well likely that the inclusion of the forecast is an incomplete proxy for the information signal used by farmers to adapt to weather. Even if they were to be perfect controls for such signals, they would only control for ex-ante adaptation, and as such be incomplete controls for farmers' entire adaptative practices. To the extent that forecasts capture the variation in farmers' beliefs about upcoming weather, they will purge the effect of weather realizations on profits from ex-ante adaptation. The coefficient associated with realizations in this context will capture on-the-stop adaptation and direct effects of weather.

We investigate this by running the same panel fixed effect regressions on the same sample as used previously, including and excluding the weather forecasts. Again, realized and forecasted rainfall in level and squares are included as controls, along with farm and year fixed effects. Standard errors are clustered at the department level.

Table 13: Profit Reaction to HDD

Dependent Variables:	Value Added		Profit	
Model:	(1)	(2)	(3)	(4)
Variables				
GDD	4.054	8.339	-3.128	-0.1146
	(8.998)	(9.830)	(9.714)	(10.44)
GDD (F)	22.24		11.94	
	(16.64)		(18.00)	
HDD	3.911	87.78	24.03	101.8
	(246.8)	(244.7)	(232.4)	(231.9)
HDD (F)	2,430.2***		2,005.2***	
	(657.5)		(559.4)	
Fixed-effects				
Farm	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Fit statistics				
Observations	18,109	18,109	18,109	18,109
R^2	0.80318	0.80203	0.83962	0.83903

One-way (department) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes. Realized and forecasted rainfall in levels and squares are included as controls.

We see here that coefficients increase in value when forecasts are omitted, both for GDDs and HDDs.

While in the context of France, and for our period of study, realized HDDs do not have a negative impact on profit, the value of associated coefficients is close to zero when controlling for forecasts, while they are quite large and positive - while still non-significant - when we do not control for forecasts. The positive value of HDD forecasts is indeed confounding the impact of realized HDDs.

We can expect that as long as HDD forecasts have a positive profit effect, their ommission will bias upwards the effect of realized HDDs, and the same should stand for realized GDDs. This implies that not controlling for farmers' beliefs and only regressing outcomes on realized weather is likely to understate the negative consequences of realized HDDs.

In the Annex in Table A5, we include the similar results for the cost and revenue regression. There, we see that both cost and revenue responses to HDDs become more negative when we include forecasts. For both GDDs and HDDs, not including forecasts biases upwards the coefficients associated to the weather realizations.

8 Robustness

8.1 Heterogeneous Response to Forecasts

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

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

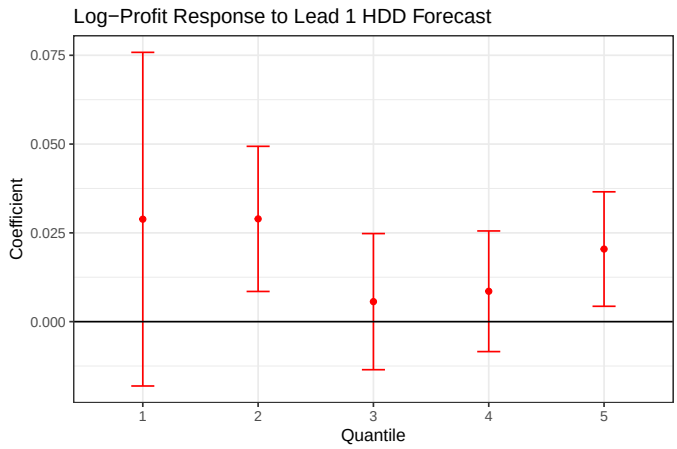


Figure 5: Varying Response to Fixed Lead of HDD

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.

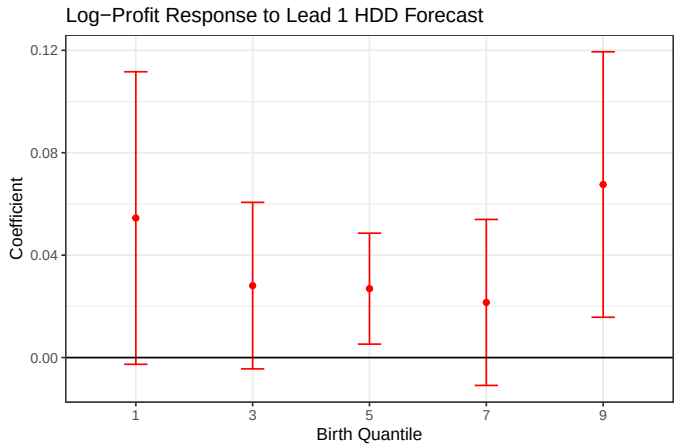


Figure 6: Profit Response to Forecasts across Manager's Age

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.

8.2 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 14: Farm-Level Profit - Lead 1 with Lags

Dependent Variables:	Value Added	Profit
Model:	(1)	(2)
Variables		
GDD	10.34	4.542
	(9.942)	(10.48)
GDD (F)	26.58	14.93
	(18.04)	(19.27)
HDD	-98.77	-58.46
	(235.6)	(233.0)
HDD (F)	2,199.7***	1,712.7***
	(690.6)	(609.4)
GDD (lag)	5.401	6.065
	(11.15)	(12.01)
HDD (lag)	332.6	344.6
	(253.7)	(279.7)
Fixed-effects		
Farm	Yes	Yes
Year	Yes	Yes
Fit statistics		
Observations	15,506	15,506
R^2	0.80695	0.84063

One-way (department) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes. Realized and forecasted rainfall in levels and squares are included as controls.

Compared to our main results in Table 2, 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.

8.3 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 15: Farm-Level Profit - Alternative Weather

Dependent Variables:	Value Added	Profit
Model:	(1)	(2)
Variables		
GDD	8.864	0.9403
	(8.746)	(9.073)
GDD (F)	-2.583	-7.674
	(17.36)	(17.80)
HDD	6.027	17.02
	(124.5)	(123.4)
HDD (F)	905.7***	713.5***
	(174.2)	(154.6)
FDD	52.45*	55.16*
	(29.69)	(32.80)
FDD (F)	180.0*	108.1
	(91.99)	(76.60)
Fixed-effects		
Farm	Yes	Yes
Year	Yes	Yes
Fit statistics		
Observations	18,428	18,428
R^2	0.80514	0.84213
Within R^2	0.01144	0.00713

One-way (department) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes. Realized and forecasted rainfall in levels and squares are included as controls.

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 A6 we show the costs and revenue responses to these alternative weather calculations.

9 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 key to understanding the scope of climate change, and to design policies that could mitigate its effects. Doing so – however – remains difficult. For a given location, actual climate change is only observed over long periods of time, out of which it is hard to extract exogenous identifying variation. Researchers have thus focused on estimating the impacts of marginal weather variations, and on how to extrapolate them into climate change damages. Producing credible extrapolations is a difficult task, as discussed by Auffhammer (2018) and Cui et al. (2023). Isolating the direct effect of marginal weather variations is difficult as well. Indeed, when we look at any outcome that is the result of a an optimized agent problem, we can expect that the agent will respond to the variations in weather, and as such that the observed variation in the outcome that follows the changes in weather is a composite of a direct effect of weather, and an indirect effect of agent adaptation. Hsiang (2016) developed a coherent framework with a clear set of conditions necessary to perform both these tasks: recover direct effects of weather variation, and extrapolate them. Following the sufficient statistics literature, he bases the first of these two steps on the use of the envelope theorem. For marginal variations in weather, indirect consequences of agent re-optimization to weather will drop out, and two-way fixed effects regressions of an optimized outcome on random weather shocks will yield correct estimates of the direct consequences of weather.

In this paper, we propose a simple test for the conditions underlying the use of the envelope theorem – namely that agent re-optimization has negligible impacts on the analysed outcome. We test this in the most studied context of climate change damages: the impact of weather shocks on agricultural profits. We do so in the French context, a relevant one given the size of France in world cereal, oil and protein crop production. Using detailed forecasted and realized temperature data, we argue that we can isolate the profit effect of part of the agent's ex-ante re-optimizing behavior to forecasted heat shocks. We show that agents re-optimize their production schedules in a way that positively and significantly impact their profit, contradicting the neglibility assumption underlying the use of the envelope theorem. We further decompose this effect, and show that it stems from a positive production response to forecasted heat shocks. Finally, we propose different avenues for the non-applicability of the theorem in our context. First of all, we show that agents use no-cost adaptation strategies which are likely to impact their profit in a discrete way - specifically we show how the timing of different production actions throughout the growing season are impacted by forecasted heat shocks. Second, we show that aggregating temperature realizations into growing and heating degree days bundles together marginal and non-marginal temperature shocks. In our

context at least, heating degree days variation is mostly driven by non-marginal variation for which the envelope conditions are not well suited.

We discuss the consequences of the breach of the envelope theorem conditions. If agent re-optimization counteracts some of the negative impact of heat shocks, imperfectly controlling for agent adaptation is likely going to underestimate the negative consequences of heat shocks, later extrapolated into downward biased impacts of climate change. We propose three methodological avenues to gauge and control the empirical non-applicability of the envelope conditions. First, performing a decomposition of the variation in weather driving the results in the study at play, and testing whether outcomes react in the same way to marginal and non-marginal variations. Second, discuss the prevalence of discrete agent actions which would be ill-suited to the use of the envelope theorem. Finally, include proxies that capture the agent's beliefs about upcoming weather, in order to control for at least some of their adaptation.

References

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*.

Bleakly, H. (2010): "Health, Human Capital, and Development," *Annual Review of Economics*.

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

Bruneel-Zupanc, C. (2022): "Discrete-Continuous Dynamic Choice Models: Identification and Conditional Choice Probability Estimation," *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 Natioanal Academy of Sciences*.

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

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

Chetty, R. (2009): "Sufficient Statistics for Welfare Analysis: a Bridge between Structural and Reduced Form Methods," *Annual Review of Economics*.

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

Cooper, R. W. and J. C. Haltiwanger (2006): "On the Nature of Capital Adjustment Costs," *Review of Economic Studies*.

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

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

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

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

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

Dube, A. (2015): "The envelope (theorem) please: Profits, efficiency wages, and monopsony," *Personal Blog Post*.

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*.

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

Hall, R. E. (1978): "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence," *Journal of Political Economy*.

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*.

Krugman, P. (2015): "Notes on Walmart and Wages (Wonkish)," *New York Times*.

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

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

Milgrom, P. and I. Segal (2002): "Envelope Theorems for Arbitrary Choice Sets," *Econometrica*.

Sah, R. and J. Zhao (1998): "Some Envelope Theorems for Integer and Discrete Choice Variables," *International Economic Review*.

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.

Shirdelmoghanloo, H., K. Chen, B. H. Paynter, T. T. Angessa, S. Wescott, H. A. khan, C. B. Hill, and C. Li (2022): "Grain-Filling rate Improves Physical Grain Quality in Barley Under Heat Stress Conditions During the Grain-Filling Period," *Frontiers in Plante Science*.

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*.

Tack, J., A. Barkley, and L. L. Nalley (2015): "Effect of Warming Temperatures on US Wheat Yields," *Proceedings of the National Academy of Science*.

Taylor, C. R. (1989): "Duality, Optimization and Microeconomic Theory," *Western Journal of Agricultural Economics*.

Appendix for online publication
TITLE

A Appendix

A.1 Figures

A.1.1 Precision of Forecasts

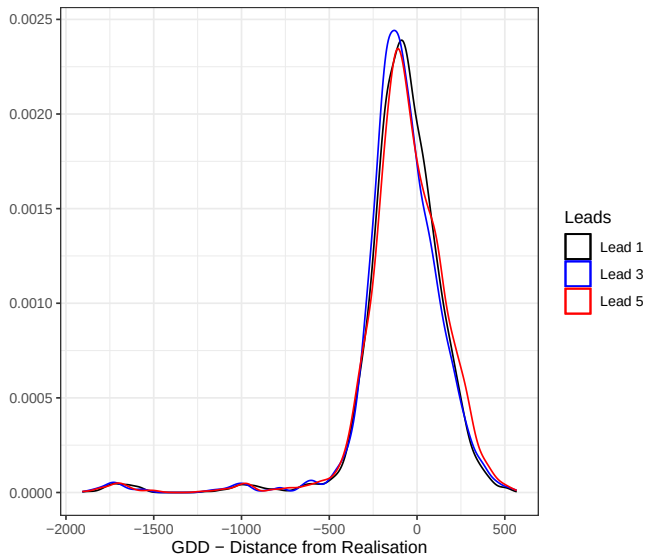


Figure A1: Distribution of Temperature Forecast Error

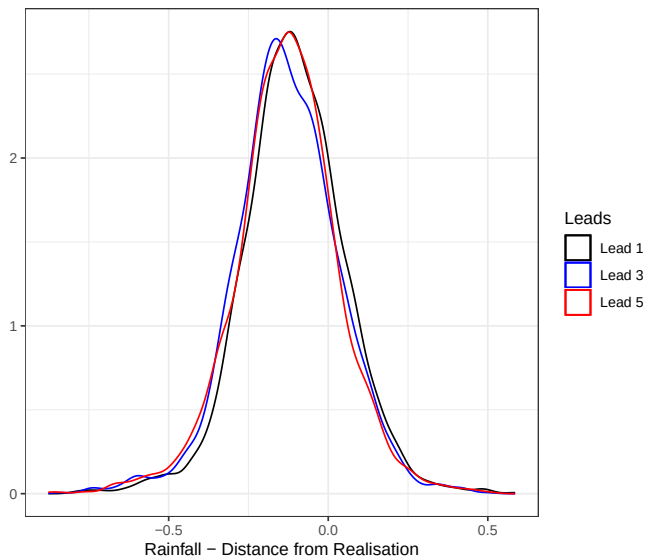


Figure A2: Distribution of Rainfall Forecast Error

A.1.2 Climate in France

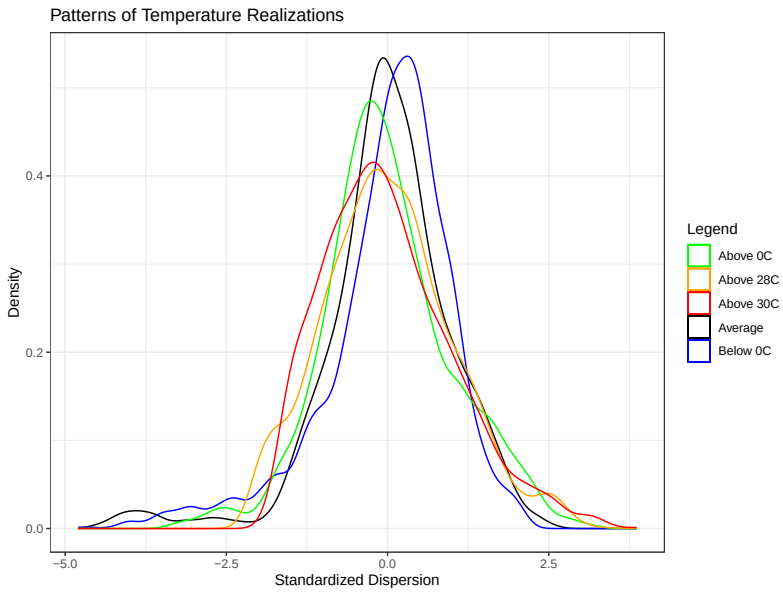


Figure A3: Distribution of the Standardized Conditional and Unconditional Mean Temperature

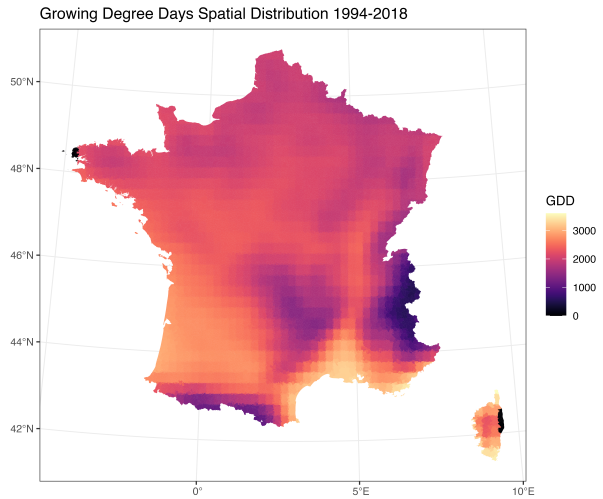


Figure A4: Distribution of Growing Degree Days

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

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

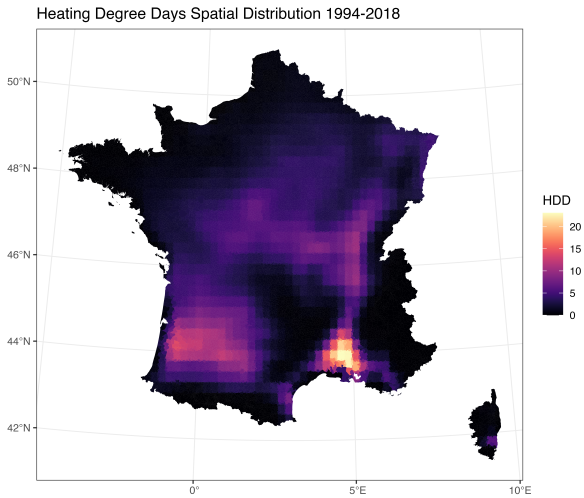


Figure A5: Distribution of Rainfall Forecast Error

A.2 Tables

Table A1: Descriptive Statistics - Farm-Level Dataset

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max
Rainfall	18,109	0.751	0.170	0.372	0.632	0.846	1.634
Rainfall (F1)	18,109	0.851	0.125	0.608	0.758	0.919	1.747
Rainfall (F2)	18,109	0.871	0.128	0.649	0.778	0.941	1.812
GDD	18,109	2,144.799	265.776	1,132.679	1,958.584	2,300.500	3,239.741
GDD (F)	18,109	2,190.808	239.034	1,053.609	2,038.088	2,318.042	3,015.025
HDD	18,109	2.355	3.165	0.000	0.247	3.131	23.808
HDD (F)	18,109	0.449	1.140	0	0	0.4	16
Sales	18,109	154,556.200	113,847.700	1,106.436	77,747.120	198,427.300	1,333,809.000
Total Costs	18,109	122,486.900	85,822.060	6,955.080	66,072.370	155,301.200	1,035,695.000
Intermediate Inputs	18,109	103,126.700	67,460.700	5,860.800	58,424.110	131,622.800	841,067.000
Value Added	18,109	58,108.540	62,612.640	−159,167.000	18,256.480	80,092.380	875,181.400
Profit	18,109	86,456.970	70,170.230	−135,056.500	39,799.620	114,723.600	973,561.100
Price Index	18,109	266.102	301.199	1.463	140.668	231.812	11,322.960
Storage (sum)	18,109	89.441	1,862.606	−28,569	−280	500	25,085
Storage Index	15,421	66.127	952.323	−19,787.310	−65.144	173.518	25,085.000
Output (sum)	18,109	7,738.650	5,027.456	146	4,201	10,104	58,374
Ouptput Index	18,109	2,332.846	2,826.531	27.588	761.109	2,751.504	58,374.000
Production (corn)	9,763	3,348.302	3,930.548	0.000	860.000	4,507.500	58,374.000
Price (corn)	9,388	150.367	39.122	36.012	123.483	173.047	1,332.181
Sales (corn)	9,763	51,811.830	67,650.980	0.000	10,847.290	66,075.980	1,155,906.000
Quantity sold (corn)	9,763	3,323.763	4,010.641	0.000	795.170	4,451.200	68,524.000
Production (wheat)	16,613	3,841.230	3,011.147	0.000	1,663.000	5,248.780	31,452.000
Price (wheat)	16,467	163.333	38.021	16.831	136.707	184.243	1,768.937
Sales (wheat)	16,613	61,586.760	54,131.060	0.000	24,348.130	82,958.100	660,068.800
Quantity sold (wheat)	16,613	3,798.452	3,170.614	0.000	1,559.000	5,160.290	41,140.000

Table A2: 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 A3: 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 A4: Descriptive Statistics - Land Price Date

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 A5: Cost and Revenue Reactions to HDD

Dependent Variables:	Sales		Total Costs	
Model:	(1)	(2)	(3)	(4)
Variables				
GDD	1.714	6.080	-4.223	-2.902
	(9.203)	(9.469)	(5.858)	(5.902)
GDD (F)	49.94***		33.91***	
	(15.14)		(10.03)	
HDD	-168.0	-88.00	-159.6	-118.9
	(233.4)	(219.8)	(117.5)	(116.7)
HDD (F)	1,329.7**		35.16	
	(652.8)		(338.6)	
Fixed-effects				
Farm	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Fit statistics				
Observations	18,109	18,109	18,109	18,109
R^2	0.88616	0.88598	0.93581	0.93575

One-way (department) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes. Realized and forecasted rainfall in levels and squares are included as controls.

Table A6: Cost and Revenue Reactions to Alternative Weather

Dependent Variables:	Sales	Total Costs	Intermediate Inputs
Model:	(1)	(2)	(3)
Variables			
GDD	1.833	-3.230	-8.214**
	(8.150)	(6.342)	(4.045)
GDD (F)	32.08*	41.88***	26.93***
	(17.25)	(10.88)	(8.070)
HDD	-117.3	-130.5**	-90.98*
	(121.7)	(64.26)	(50.85)
HDD (F)	676.8***	35.81	15.81
	(221.2)	(154.7)	(111.9)
FDD	20.41	-5.865	-3.528
	(26.58)	(14.04)	(13.61)
FDD (F)	47.60	-21.39	-15.58
	(117.6)	(49.08)	(47.97)
Fixed-effects			
Farm	Yes	Yes	Yes
Year	Yes	Yes	Yes
Fit statistics			
Observations	18,428	18,428	18,428
R^2	0.88828	0.93685	0.93376
Within R^2	0.00362	0.00199	0.00211

One-way (department) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes. Realized and forecasted rainfall in levels and squares are included as controls.

B Annex

B.1 Definition of Variables

Definitions for the main variables used:

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

$$p_{jt} = \sum_{c \in \mathbb{C}_{jt}} \frac{l_{jct}}{\sum_{c \in \mathbb{C}_{jt}} l_{jct}} p_{jct} \tag{4}$$

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

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

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

- **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 Validation of Results

- Crop-specific results to show we get the regular intuition, and matches Schlenker Roberts.
- Farm-level outcomes with only realisations.
- Lead 1 Farm-level outcomes with lagged (1 year) weather realization.
- Storage - per crop type.

Table A7: Corn-Specific Outcomes

Dependent Variables:	Production Volume	Price	Sales Value	Production Sold
Model:	(1)	(2)	(3)	(4)
Variables				
GDD	0.4767	-0.0012	10.02	0.6082
	(0.3181)	(0.0083)	(7.347)	(0.4254)
HDD	-21.36*	0.7295**	-107.9	-10.02
	(11.26)	(0.3019)	(159.7)	(10.28)
Fixed-effects				
Farm	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Fit statistics				
Observations	9,763	9,388	9,763	9,763
R^2	0.93460	0.73025	0.87690	0.87517

One-way (CDEPT) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes. Realized rainfall in levels and squares are included as controls.

Table A8: Farm-Level Outcomes

Dependent Variables:	Sales	Total Costs	Intermediate Inputs	Value Added	Profit	Output Price Index
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
GDD	6.080	-2.902	-6.924*	8.339	-0.1146	-0.2324**
	(9.469)	(5.902)	(3.890)	(9.830)	(10.44)	(0.1097)
HDD	-88.00	-118.9	-81.20	87.78	101.8	2.757*
	(219.8)	(116.7)	(84.90)	(244.7)	(231.9)	(1.540)
Fixed-effects						
Farm	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
Observations	18,109	18,109	18,109	18,109	18,109	18,109
R^2	0.88598	0.93575	0.93296	0.80203	0.83903	0.69207

One-way (CDEPT) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes. Realized rainfall in levels and squares are included as controls.