Costs of Climate Adaptation: Evidence From French Agriculture*

Tristan Du Puy and Jeffrey G. Shrader

[1]Columbia University

August 12, 2025

Abstract

The cost of adapting to a changing climate is one important input into climate mitigation and adaptation policy. But identification challenges and a lack of cost-related data mean that such estimates are rare. In this paper, we use uniquely rich data on agriculture in France to provide novel, direct estimates of the marginal cost of preemptively adapting to temperature shocks. The dataset is a farm-level panel with measures of outputs, inputs, and prices from 1994–2018. To recover the effects of adaptation, we merge this farm data with measures of realized and forecasted weather. Controlling for realized weather, we use forecasts as information shocks to estimate costs and benefits of ex ante adaptation. In the year that heat shocks arrive, we find that the cost of adaptation is low for the average farm in France. In contrast, the benefits of adaptation are large. Dynamically, however, the costs of adapting to that original shock rise over subsequent years, eventually leading the marginal benefits of adaptation to equal marginal costs, a result consistent with a dynamic envelope theorem but inconsistent with a static sufficient statistic approach widely used when estimating climate damages. This pattern of adaptation costs is driven by the behavior that farms engage in when responding to forecasts. In the year of the shock, they use the forecasts to change the timing of key growing decisions as well as their crop mix, rather than to change costly inputs. These decisions create dynamic switching costs, which we can relate to soil nutrients and pest control. Finally, we show that forms of adaptation currently implemented by farmers are less likely to remain relevant in a warmer world.

JEL Codes: D22, D24, Q12, Q54
Keywords: Climate Change, Adaptation, Agriculture, Forecasts

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

1 Introduction

The potential consequences of climate change are a central concern for our century. Adaptation costs—the cost of actions people take to prepare for or adjust to the climate—are important to consider for both mitigation of greenhouse gas emissions and adaptation itself. The total damage from climate change comes not only from the direct effects of climate shocks but also the cost of adaptation.[1] And as global warming continues to occur, adaptation costs will also determine how much adaptation occurs and thus how much climate damage people will be willing or able to avoid. Despite the importance of adaptation costs, direct estimates of them are currently lacking for important sectors of the economy. In this paper, we provide novel evidence related to the costs of adaptation in agriculture, one of the sectors of the economy most directly affected by climate.

Estimating adaptation costs faces multiple challenges. First, there is simply a lack of data on costs. Crop yields, profits, and farm land values are relatively commonly observed and a large literature has used them to estimate the damages from changing weather (see, e.g., Mendelsohn et al., 1994; Schlenker and Roberts, 2009). Output data have also been used to estimate the marginal benefits of adaptation which have then been used to bound adaptation costs using a revealed preference approach (Hultgren et al., 2022). But the input levels and prices needed to directly recover adaptation costs are rarely observed. This lack of cost data is not limited to the agricultural setting. For example, recent work on the effect of climate change on mortality also bounds unobserved adaptation costs using adaptation benefits (Carleton et al., 2022).

Second, there are identification challenges. Adaptation actions are endogenous and high-dimensional. A common definition of climate adaptation used, for instance in **?**, states that adaptation is "adjustment in natural or human systems in response to actual or expected climatic stimuli or their effects, which moderates harm or exploits beneficial opportunities."[2] One identification challenge is to separate the effect of endogenous adaptation actions from the direct effect of climate shocks—the effect that climate shocks have on the farmer holding their adaptation decisions fixed. Another identification challenge comes from the large number of potential adaptation actions. To be informative about the overall cost of adaptation, one would like to understand the costs of many potential adaptation actions.[3]

To overcome the data challenge, we use a uniquely rich panel dataset on the French agriculture sector for 1994–2018. The dataset contains accounting measures for both revenue and—crucially—costs directly at the farm level. To gain identification, we leverage variation in both realized and forecasted weather. The forecasts, in particular, allow us to isolate how farm-level revenue and costs respond to information about upcoming weather events, and we can do so while holding actual weather realizations constant. We interpret the revenue and cost responses to changes in information as measuring the marginal benefits and costs of ex ante adaptation actions. To interpret the estimates as the costs and benefits of *climate change*, we consider how the frequency of these responses will change as the climate (the distribution of weather) changes.

The intuition for identification is as follows: weather has multiple effects on a farm. Most directly, weather affects farm productivity because weather is an important input into crop production (for instance, sunny days help corn grow, but too much heat will kill corn). The farmer can also adapt to weather in many ways. They can take actions after the weather occurs to try to help their crops fare better (for example weeding more intensively if weather conditions have been conducive to weed growth). And they can take

[1] In other words, adaptation costs are part of what determines optimal climate policy stringency because mitigation of global warming would allow people to save on paying those costs. More generally, adaptation costs are a form of defensive expenditure, and reducing defensive expenditures is a benefit from reducing a negative externality (Freeman et al., 2014; Deschênes et al., 2017).

[2] Section 2 presents the formalization of this definition used in this paper. In words, that formalization says that adaptation actions are changes in farm inputs (broadly defined) in response to changes in growing-season weather conditions. We discuss the difference between responses to weather variation and responses to climate change further below.

[3] Again connecting this to defensive expenditures, estimates of the willingness to pay to avoid a negative externality that come from defensive expenditures will better reflect total willingness to pay the more complete the set of defensive expenditures observed.

action in anticipation of the arrival of weather. Farmers have access to sophisticated forecasts of weather, helping them act in advance of a weather event to try to make the effect of the weather more favorable (for example, apple farmers who anticipate an early season freeze can hire helicopters to fly over their orchards to blow warmer air down onto the apple trees). In our empirical strategy, we estimate the effect of forecasts and weather realizations on farm outcomes. The strategy identifies the last of these effects: the effect of adaptation actions taken in advance of the arrival of weather shocks. Weather realizations cause both direct effects and can induce ex post adaptation. Forecasts affect ex ante adaptation behavior but they are correlated with realizations. By controlling for realizations, we use the portion of the variation in forecasts that specifically identifies ex ante adaptation.

Ex ante adaptation is important in agricultural settings. Previous literature has identified a range of adaptation actions that occur partly or wholly in advance of weather realizations. These range from relatively short-term actions like timing and application of on-farm inputs such as fertilizer and pest control (Reidsma et al., 2010); crop choice that necessarily happens before the growing season (Rising and Devineni, 2020); labor reallocation (Rosenzweig and Udry, 2014; Colmer, 2021); irrigation decisions which have important intertemporal considerations (Hagerty, 2022); changes in land use patterns including the migration of farms and transition of farmland to other uses (Mendelsohn et al., 1994; **?**); and technical change including the development of new seed varieties or farming techniques (Moscona and Sastry, 2022). Our identification strategy can estimate the marginal benefits and costs of some, but not all, of these adaptation strategies. In the empirical application, we examine the effect of temperature forecasts issued from one to five months ahead of weather that arrives throughout the growing season (so the earliest forecast farmers have access to arrives five months prior to the growing season, and farmers continue to receive forecasts throughout the growings season). Thus, we can identify the benefits and costs of actions that a farmer can take with a few months of advance warning of upcoming weather. Changes in inputs, planting decisions, and labor have all been shown to change within that time window. Fitting dynamic models, we can also estimate how these adaptation actions affect farm profits over subsequent growing seasons to better capture longer-term effects. We cannot directly estimate the effects of adaptation actions that operate over longer horizons, though we do conduct an additional analysis to compare the effectiveness of the adaption actions for farmers that are already experiencing high levels of extreme heat versus those that are not.

Estimates of the marginal costs and benefits of this set of adaptation actions allows us to speak to two important sets of results. First, existing literature highlights the importance of shorter-horizon adaptations like crop switching and input choices as farm responses to climate change. We provide novel estimates of the marginal costs of that class of adaptations. Second, we can use our estimates to conduct a sharp but one-sided test of a widely used sufficient statistic approach to estimating climate damages.

We start the paper by outlining the workhorse model used to study firm-level adaptation to weather in the climate economics literature (e.g. Deschênes and Greenstone (2007), Ortiz-Bobea and Just (2013), Hsiang (2016), Carleton et al. (2022)). This model serves two purposes in our analysis. The first is to illustrate our identification strategy for the consequences of adaptation. The second is to serve as a straw man for the study of costs of adaptation. We use it to discuss the revealed preference approach to the estimation of these costs, and a test for the validity of the static revealed preference procedure which we later implement.

The second part of this paper provides evidence for farm-level adaptation to weather, and several mechanisms for adaptation. Within the period of the weather shock, farmers are able to leverage heat forecasts to generate relatively large revenue gains at small costs. A one standard deviation increase in forecasted heating degree days leads to a 1% increase in revenue, while costs have no statistically significant response. In comparison, farms respond to more moderate temperature forecasts by increasing the scale of their production, at an equivalent marginal cost. Both these results hold across alternative specifications, different aggregations of temperature variables, as well as robustness tests. Focusing on heat forecasts, we

decompose results by the sign of the forecast error (w.r.t to the heat realization). Relative to an on-point forecast, under-forecasts are harmful to farm profit, while over-forecasts have no additional positive effect. This suggests that adaptation benefits are not monotonically increasing in adaptation actions, but rather rely on matching the expected heat (Shrader et al., 2023). Making use of additional field level data, we show that adaptation to heat shocks takes in part the form of ex ante timing changes affecting different steps of the growing season. We further provide evidence for crop switching. We identify two potential patterns of switching from high heat-sensitive crops towards less sensitive ones: away from sunflower and towards colza, away from corn and towards wheat. We also observe increases in the area allocated to peas and beans in line with the results of Aragón et al. (2021).

The differences in the cost and revenue consequences of ex ante farm adaptation naturally translate into net profit gains of ex ante adaptation. We investigate the farm-level profit responses to forecasted and realized weather. We further discuss the implications of these net positive profit responses for the applicability of the static envelope theorem to our context. In a climate change damages setting, the envelope theorem states that the cost and benefit consequences of re-optimizing production in response to marginal weather variation should offset each other. In our empirical context, this correspondence does not hold as the 95% confidence interval for the static profit effect of a marginal change in forecasted HDDs is [.7k, 3.4k] euros. The static envelope theorem has been invoked by recent studies of climate change to provide estimates of adaptation costs (Deschênes and Greenstone, 2007; Hsiang, 2016; Deryugina and Hsiang, 2017; Merel et al., 2022; Carleton et al., 2022).[4] In particular, the envelope theorem allows one to equate marginal benefits and marginal costs of adaptation, and hence to infer a measure of the costs of adaptation.

We discuss the potential reasons for the breach of the static envelope theorem. A first element is that growing and heating degree days temperature aggregations bring together marginal and non-marginal variations in weather. Over our period of study, growing degree days are mostly composed of marginal variations,[5] while heating degree days are mostly composed of non-marginal variation. This source of variation tracks with our observed, statistically indistinguishable cost and benefit responses to forecasted growing degree days (GDDs), and net-positive profit responses to forecasted heating degree days (HDDs). A second element of answer is the difference in nature of the response to forecasted growing and heating degree days. Responses to forecasted GDDs imply some timing changes, as well as input responses, while responses to forecasted HDDs are uniquely driven by significant timing and cropping changes. These changes are likely to drive non-continuous changes in farm-level profits, and are hence less suited to the assumptions of usual envelope theorems.

Focusing on the within-period consequences of ex ante adaptation does not, however, provide a complete picture of the consequences of adaptation. In a third part, we study these farm-level consequences of adaptation. We investigate for this the profit and cost response to ex ante adaptation over time. While the period of the forecast shock is marked by net profit gains, the following two periods show significant profit losses. The cost movement parallels the evolution of profit, with increases in costs in future periods. Overall, the net discounted profit consequences of ex ante adaptation have a 95% confidence interval of $[-3.7k, 1.4k]$ euros, which hence overlaps zero. This implies that while we can reject the static envelop theorem, we cannot reject that the dynamic envelope theorem holds. In our context, the present discounted value of costs of adaptation amounts to 1k€, or 95% of the point estimate for the total profit effect of forecasted heat. Because we identify crop switching as one of the main adaptation strategies used by farmers, the presence of these future costs suggests that changes to one's crop mix entail switching costs. The presence of such costs has been central to the growing dynamic land use literature (Livingston et al., 2008; Scott, 2013; Burlig et al., 2024) but we are one of the first papers to recover evidence for them in a reduced form setting. We

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

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

further show these cost increases are mostly driven by future increased fertilizer and pesticide use, as well as labor hours. These are in lines with agro-science interpretations that deviations from optimal crop-rotations will hurt field nutrient stocks, and raise exposure to pest (Livingston et al., 2008).

In a last section, we investigate whether current farm responses to forecasted HDDs are good indicators of future adaptation strategies to climate change in France. This exercise first requires to understand how climate change will impact average French weather, and second how farm adaptation will shift with it. On the first point, we use CMIP6 climate change predictions until the end of the century. Holding cost responses to weather constant, we extrapolate the costs of adaptation. Under these conditions, climate change is expected to increase the yearly costs of adaptation from 2.6k€ to 21k€ by the end of the century, or a change from 3% to 27% of average current-day yearly profits. The underlying assumption to this projection is that farms will not change their adaptation behavior as climate change unfolds. We finally discuss the extent to which we can expect farms to shift their adaptation behavior. We rely for this on a combination of the Ricardian and panel approaches to climate change (Auffhammer, 2018). Specifically, we look at costs of adaptation heterogeneity across hotter and colder areas of France. This approach is cross-sectional in nature, and provides suggestive evidence for decreasing returns to adaptation as average temperature increases. These imply that adaptation practices are unlikely to remain similarly successful in the future.

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

The literature on climate damages has grown rapidly over the last decades, crossing from Ricardian approaches into global dynamic general equilibrium models and fine scale panel studies of context-specific dose-response functions.[6] Our paper stands out in this literature by the unusual precision of the data directly measured at the farm-level, and within a repeated panel format. While studies of adaptation to climate change have previously relied on farm-level data, for example Aragón et al. (2021) with repeated cross sections of smallholder farms in Peru, our panel structure and the depth of the data collected is to the best of our knowledge unprecedented. In the context of developed economies, most research has relied on aggregated data for which yields or profits are available (Schlenker and Roberts, 2009, for example). In France, previous research has been conducted at the department-level and over a longer period of time by Gammans et al. (2017), but has focused on yields specifically, while research that has been done at the establishment-level has focused on a small subset of farms (Bareille and Chakir, 2023). We build on these papers by distinguishing cost and revenue adaptive responses. The granularity of our data also implies that we are identifying agent-level responses to heat and climate change, a level where data is likely better measured and identification more robust.

Our paper also builds on the subset of the literature which has approached adaptation with a focus on the definition of farmers' climate beliefs. Such papers have highlighted how the nature of beliefs about the climate process structures what can be interpreted as adaptive behavior in response to changes in climate. Kelly et al. (2005) defined two sets of costs related to environmental changes: transition costs and equilibrium costs. The first ones being incurred as the economy moves along the transition path towards a new equilibrium where the new environmental context is fully taken into account. Burke and Emerick (2016) illustrate how the definition of beliefs drives our understanding of adaptation. Kala (2019) extends these approaches by comparing different learning models for the timing of the monsoon in India. She makes the point that recovering farmers' learning behavior can depend on our modeling of the objective that they maximize. What we measure as the extent of their adaption to changing climate patterns relies on the behavioral model and the objective function assigned to them. Shrader (2023) offers a way to use weather forecasts to disentangle adaptation effects from direct climate damages. Our paper builds on the arguments

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

from this paper. Lemoine (2024) extends this argument to a richer dynamic setting and, relevant to our results, shows that adaptation can have first-order consequences for an agent's static profits even if adaptation has no net effect dynamically. This is consistent with our empirical findings. Related to this literature, we also 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, our work relates to two papers which have approached climate damages without relying on the static envelope theorem applied to differentiable objectives. 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 traditional envelope conditions apply, and discrete actions for which they do not. They derive a model where marginal variation in weather can lead to arbitrarily large impacts on profit by inducing discrete changes in the type of tree species planted by foresters. Empirically, they find that extensive margin (discrete) adaptive actions can have large profit responses, but typically induce a net-zero profit response. Similarly, Bareille and Chakir (2023) also rely on structural methods to disentangle adaptation effects from direct weather effects, but their method requires a complete enumeration and modeling of adaptation methods, as well as instruments separately identifying adaptation effects from the weather impacts. Mirroring the discussions in the sufficient statistic literature, by-passing the envelope theorem is a difficult task which requires a structural setting and implies trading some assumptions for others. We see our work as shedding additional light on the trade-offs between structural and envelope theorem-free, and reduced form and envelope theorem-dependent methods.

The rest of the paper proceeds as follows. In section 2 describes the central challenge of identifying adaptation costs and describes the sufficient statistic approach used either to justify ignoring or to bound adaptation costs in the literature. 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 present the main results on the marginal benefits and costs of adaptation to weather shocks. section 6 examines the effects on profit, and section 7 examines limits to the adaptation found in the prior section. Finally, section 8 concludes the paper.

2 Model

This model section serves a dual purpose. First, it formalizes the costs of adaptation to weather variation, providing the rationale for our identification strategy and precisely defining the costs we estimate. Second, it demonstrates how our empirical framework can be used to test the empirical applicability of the envelope theorem, which is a central tool of the climate economics literature. We generally keep a structure close to models used in the literature to discuss adaptation to climate change (Deschênes and Greenstone (2007); Ortiz-Bobea and Just (2013); Hsiang (2016); Carleton et al. (2022)). In a later part of the paper, we will provide evidence that a dynamic envelope theorem is likely more appropriate to discuss adaptation (Scheinkman and Schechtman (1983)).

Empirical Framework: We model a static, price-taking, profit-maximizing agricultural firm that makes decisions in a two-stage production process under uncertainty.

A production season runs from time t to t+1. At the beginning of the season, time t, the firm receives

a forecast f_t about the weather that will be realized later in the season. Based on this forecast, the firm chooses an initial vector of inputs, which we denote the ex ante actions x_t^1. Later in the season, at time t+b, with $b \in (0,1)$ denoting forecast lead, the weather shock w_{t+b} is realized. We assume forecasts are unbiased, such that $f_t = w_{t+b} + \varepsilon_{t+b}$, with $\varepsilon_{t+b} \sim \mathcal{N}(0, \sigma^2)$. After observing the realized weather w_{t+b}, the firm can take a second set of actions, the ex post actions, x_{t+b}^2. Production then occurs and profits are realized, conditional on all actions and the realized weather.

The firm's realized profit is given by $\pi_{t+b} = p_{t+b}q_{t+b}(x_t^1, x_{t+b}^2, w_{t+b}) - c(x_t^1, x_{t+b}^2, w_{t+b})$. The firm chooses its actions sequentially. The ex post action x_{t+b}^2 is chosen to maximize profit with full knowledge of w_{t+b} and conditional on the chosen x_t^1. The ex ante action x_t^1 is itself chosen to maximize expected profit conditional on the forecast f_t, anticipating the future optimal choice of x_{t+b}^2. Throughout the analysis, we assume that these functions are continuously differentiable and satisfy the necessary regularity conditions for the dominated convergence theorem to apply.

The firm's initial problem can be written as solving for the ex ante static value function, $\overline{v}(f_t)$:

$$\begin{aligned}\overline{v}(f_t) &= \underset{x_t^1}{max}\ \mathbb{E}\big[\underset{x_{t+b}^2}{max}\ \pi_{t+b}(x_t^1, x_{t+b}^2, w_{t+b})|f_t\big] \\ &= E\big[\pi_{t+b}(x_t^{1,\star}, x_{t+b}^{2,\star}, w_{t+b})|f_t\big],\end{aligned}$$

with $\{x_t^{1,\star}, x_{t+b}^{2,\star}\}$ the set of expected profit-maximizing actions.

We are interested in the marginal cost of adaptation to a change in the weather forecast. In order to isolate this cost, we start by decomposing the effect of a change in forecasts on expected costs:

$$\begin{aligned}\frac{d\mathbb{E}[c_{t+b}|f_t]}{df_t} =& \underbrace{\mathbb{E}\left[\frac{\partial c_{t+b}}{\partial w_{t+b}}\middle| f_t\right]\frac{dw_{t+b}}{df_t}}_{\text{Direct Effect}} + \\ & \underbrace{\mathbb{E}\left[\frac{\partial c_{t+b}}{\partial x_t^1} + \frac{\partial c_{t+b}}{\partial x_{t+b}^2}\frac{\partial x_{t+b}^{2,\star}}{\partial x_t^1}\middle| f_t\right]\frac{dx_t^{1,\star}}{df_t}}_{\text{Ex Ante Adaptation}} + \\ & \underbrace{\mathbb{E}\left[\frac{\partial c_{t+b}}{\partial x_{t+b}^2}\frac{\partial x_{t+b}^{2,\star}}{\partial w_{t+b}}\middle| f_t\right]\frac{dw_{t+b}}{df_t}}_{\text{Ex Post Adaptation}}.\end{aligned} \tag{1}$$

The first term is the direct effect. This term captures how expected costs change because the forecast is correlated with realized weather, which itself can affect costs even holding actions constant (e.g., by increasing wages to compensate for occupational disamenities). Consistent with the literature (Hsiang, 2016), we define adaptation as the firm's costly response to a change in its beliefs. As the direct effect is not mediated by such a response, we distinguish it from the costs of adaptation. The remaining two terms represent the costs of adaptation, which we decompose into ex ante and ex post adaptation. The second term captures ex ante adaptation, how the forecast directly alters the firm's optimal ex ante action, which in turn changes expected costs. This term accounts for the fact that ex ante actions can constrain future ex post actions. This reflects adaptation to the forecast signal itself. The third term reflects ex post adaptation via a more indirect channel. A change in the forecast will empirically be correlated with a change in the realized weather, which prompts a change in the optimal ex post action. This is the cost of adaptation to the new information learned from the weather realization itself.

Our empirical strategy aims to isolate the marginal costs of adaptation. We achieve this by estimating the marginal cost response to forecasts, while holding realized weather constant. This sets $dw_{t+b} = 0$, and

eliminates both the direct effect and the ex post adaptation terms, leaving our empirical estimand as an estimate for costs of ex ante adaptation:

$$\frac{d\mathbb{E}[c_{t+b}|f_t]}{df_t}|_{dw_{t+b}=0} = \mathbb{E}\left[\frac{\partial c_{t+b}}{\partial x^1_t} + \frac{\partial c_{t+b}}{\partial x^2_{t+b}}\frac{\partial x^{2,\star}_{t+b}}{\partial x^1_t}\middle| f_t\right]\frac{dx^{1,\star}_t}{df_t}. \tag{2}$$

Rationale for Using Forecasts: Our empirical strategy uses variation in weather forecasts, conditional on realized weather, to identify the costs of ex ante adaptation. As shown above, this estimate is both an unbiased estimate of costs, but also a partial one focusing on ex ante costs. A natural alternative to our strategy would be to estimate the total derivative of costs with respect to realized weather $\frac{dc_{t+b}}{dw_{t+b}}$. As shown in Equation (2), this captures the sum of the direct effect of weather and the costs of adaptation.

$$\begin{aligned}\frac{dc_{t+b}}{dw_{t+b}} =& \frac{\partial c_{t+b}}{\partial w_{t+b}}\frac{dw_{t+b}}{dw_{t+b}} + \\ & \left(\frac{\partial c_{t+b}}{\partial x^1_t} + \frac{\partial c_{t+b}}{\partial x^2_{t+b}}\frac{\partial x^2_{t+b}}{\partial x^1_t}\right)\frac{\partial x^1_t}{\partial f_t}\frac{df_t}{dw_{t+b}} + \\ & \frac{\partial c_{t+b}}{\partial x^2_{t+b}}\frac{\partial x^2_{t+b}}{\partial w_{t+b}}\frac{dw_{t+b}}{dw_{t+b}}\end{aligned}$$

However, this approach faces two fundamental challenges. First is an identification problem: the direct effect of weather on costs is empirically inseparable from the costs of ex post adaptation to realized weather (the third term). We can only observe their sum. Second, and perhaps more importantly, is a measurement and relevance problem. Realized weather may be a poor measure of the information that drives key adaptation decisions. Many crucial adaptation strategies, such as crop switching, involve high adjustment costs and must be decided before the weather is realized. For these margins, realized weather is not the relevant belief shifter. In Equation 2, these are captured by the second term and only imperfectly captured through the correlation between weather and forecasts. Regressing outcomes on weather would then fail to properly capture these important adaptation costs, leading to attenuated estimates.

In effect, our strategy isolates the costs of ex ante adaptation by using the variation in farmer beliefs generated by forecast errors, holding realized weather constant. This approach is analogous to a reduced form instrumental variable (IV) strategy that simultaneously addresses the two challenges outlined above. It resolves the endogeneity problem by design, as it nets out any direct effects of realized weather on costs. Furthermore, it mitigates measurement error by using forecasts as a more relevant belief shifter for ex ante decisions than the imperfectly observed, contemporaneous weather shock (Shrader (2023)). Consequently, our estimate has a Local Average Treatment Effect (LATE) interpretation: it captures the marginal costs for the subset of adaptation actions that are responsive to forecast information.

Envelope Theorem: A second objective of this model is to provide a framework for testing the applicability of the envelope theorem, a central tool in the climate economics literature. To fix ideas, we first outline the theorem in the context of our model. To see how a change in forecasts affects the maximized

expected value of profit, we take the total derivative of the value function with respect to f_t:

$$\begin{aligned}\frac{d\overline{v}(f_t)}{df_t} =&\mathbb{E}\big[\frac{\partial \pi_{t+b}}{\partial w_{t+b}}|f_t\big]\frac{dw_{t+b}}{df_{t+b}}+\\ &\mathbb{E}\big[(\frac{\partial \pi_{t+b}}{\partial x^1_t}+\frac{\partial \pi_{t+b}}{\partial x^2_{t+b}}\frac{\partial x^2_{t+b}}{\partial x^1_t})|f_t\big]\frac{dx^1_t}{df_t}+\\ &\mathbb{E}\big[\frac{\partial \pi_{t+b}}{\partial x^2_{t+b}}\frac{\partial x^2_{t+b}}{\partial w_{t+b}}|f_t\big]\frac{dw_{t+b}}{df_t}\end{aligned}$$

This expression states that a change in forecasts affects expected profit directly (the first line), and indirectly by inducing the firm to change its optimal actions (second and third terms). The envelope theorem states that these second and third terms are equal to zero. The economic intuition is that the marginal costs and benefits of re-optimizing are already in perfect balance, so small changes have no first-order effect on the maximized value of ex ante profit.

The envelope theorem provides the theoretical foundation for two common empirical exercises in climate economics. The first application allows researchers to infer unobservable costs of adaptation from observable adaptation benefits (Carleton et al. (2022)). The theorem's result stems from the firm's first-order condition, which ensures that the marginal costs and benefits of any adaptive action are equal at the optimum. This implies that in an optimized system, if one can estimate the marginal benefits of an action (e.g., increased revenue), those benefits can serve as a proxy for the marginal costs of that action, even if those costs are not directly observable or hard to define. Some studies then use this principle to extrapolate total costs of adaptation to non-marginal changes in climate from these marginal estimates.

The second, and more widespread, application is the use of short-run weather variation to estimate the damages from long-run climate change (Hsiang (2016)). As outlined by Auffhammer (2018), this approach has been central to the empirical literature that estimates the economic consequences of climate change. The argument rests on the envelope theorem's core result: the effect of a change in an environmental parameter on an agent's welfare is equal to its direct impact, as the first-order effects of re-optimization are zero. This implies that the marginal damages of both weather and climate variation reduce to their respective direct effects. Under the assumption that this direct damage function is the same for both types of shocks, the observable marginal damages from weather can serve as a direct proxy for the unobservable marginal damages from climate change.

Testing the Applicability of the Envelope Theorem: Our framework provides a direct method for testing the applicability of the envelope theorem in a climate economics context. To see this, consider how a change in the forecast affects the firm's maximized profit, holding realized weather constant. Drawing from our previous derivation:

$$\frac{d\overline{v}(f_t)}{df_t}|_{dw_{t+b}=0} = \mathbb{E}\big[(\frac{\partial \pi_{t+b}}{\partial x^1_t}+\frac{\partial \pi_{t+b}}{\partial x^2_{t+b}}\frac{\partial x^2_{t+b}}{\partial x^1_t})|f_t\big]\frac{dx^1_t}{df_t}$$

This expression isolates the change in profit resulting only from ex ante re-optimization in response to forecasts. The envelope theorem's general prediction—that the full consequences of agent re-optimization in response to marginal variation in weather are zero—stems from the premise that each action is individually optimized, including the specific ones driven by forecasts. This provides a testable hypothesis. While our empirical strategy only identifies the profit effect of re-optimiztaion along the forecast margin, and not the ex post one following realized weather variation, its implications are general. Finding that this specific channel has a non-zero effect on profit would constitute strong evidence against envelope theorem as applied in the

literature. We therefore test the null hypothesis implied by the theorem:

$$\frac{d\overline{v}(f_t)}{df_t}|_{dw_{t+b}=0} = 0$$

A rejection of this null suggests that the marginal effect of this type of adaptation is, in fact, non-zero, casting doubt on the assumption that agent reoptimization in general has a zero effect.

3 Empirical Context and Data

3.1 Exposure to Weather in French Agriculture

We test our model using data from French agriculture, a setting that offers several distinct advantages for identifying the costs of ex ante adaptation. We focus on cereal, oil, and protein crops, which are sensitive to weather variation and major contributors to France's agricultural output. As the largest grain producer in the European Union, and one of the world's largest cereal exporters, understanding adaptation in France has direct implications for food security.

Importantly for our identification strategy, this sector has very low exposure to irrigation; only 6% of agricultural land was irrigated in 2015 (Colas-Belcour et al. (2015)). This general absence of irrigation minimizes concerns that unobserved irrigation decisions could endogenously bias our estimates of weather impacts, a challenge noted by Braun and Schlenker (2023). It also aligns with recommendations to study weather shocks in areas where irrigation is not heavily subsidized ((Schlenker et al., 2005)).

French climate features significant weather variability, providing the necessary shocks for our analysis ((Canal (2015))), and significant heterogeneity with a Mediterranean region more heavily exposed to heat and water stress over the summer. Figure A21 shows splines describing the evolution over time of unconditional and conditional growing season mean temperature realizations at the department level. These are helpful to characterize the average climate in France. On average, temperatures remain around 10°C, with little change over our period of study. On average, below 0°C temperature are not very negative, and on average extreme heat temperature remains around 30°C. Figure A6 in the annex further shows that while there is more dispersion across French department in extreme heat events, the general dispersion of average conditional temperatures remains moderate. We also provide maps showing the geographic dispersion in growing and heating degree days in France over our time period.

To characterize weather sensitivity in our sample, we present two sets of descriptive evidence. First, Table 1 shows the land allocation for the four main crops in our sample, which account for over 85% of the agricultural area in the data. We associate each crop with its mean growing season temperature, and with a temperature threshold for its growth taken from the literature.[7] Wheat is the dominant crop, and as noted in the literature, it is relatively more heat tolerant than corn or sunflower. This suggests the average farm in our sample has a baseline resilience to some heat shocks.

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

Table 1: Crop mix composition

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

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

Second, to confirm the impact of temperature on farm output, we estimate a flexible regression model following Schlenker and Roberts (2009). We use a restricted cubic spline to model the non-linear relation between exposure to temperature and outcomes, and include farm fixed effects and region-specific quadratic time trends. Figure 1 shows the crop-specific effect of temperature, while subsubsection A.1.4 looks at the aggregate effect at the farm level. These results show that crop yields have heterogeneous responses to heat. Corn and sunflower yields decline sharply at temperatures above 25°C, whereas wheat and colza exhibit flatter damage functions in this range. These heterogeneous sensitivities highlight the importance of ex ante crop choice as a key adaptation margin.

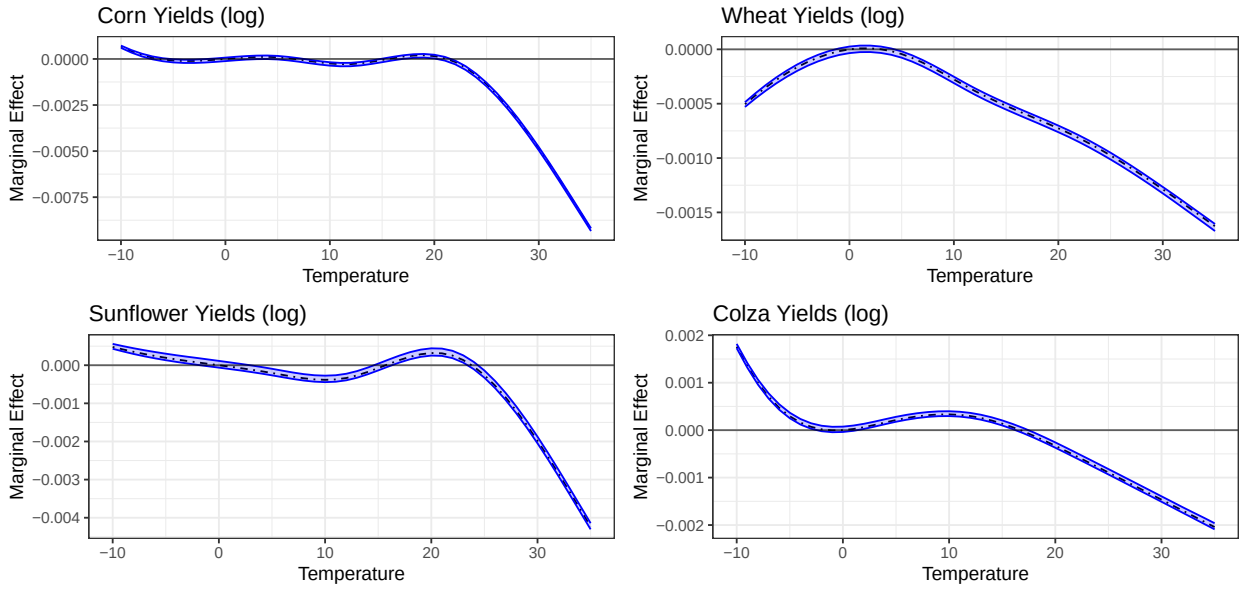


Figure 1: Temperature Effects on Crop-Specific Yields

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

Our strategy relies on farmers responding to weather forecasts. High-quality coupled atmosphere-ocean models started to be used for seasonal forecasts by Météo-France in the mid-to-late 1990s, allowing forecasts

to move past the four months ahead lead (Canal (2015)). For our analysis, we use hindcasts produced by the European Center for Medium-Range Weather Forecasts (ECMWF), which are available from 1994. While our data includes forecast leads of up to four months, our main analysis focuses on one-month-ahead forecasts, which represent a credible and widely available information source for farmers' in-season decisions during our study period. Today, different companies offer an access to month-ahead forecasts throughout the season, and the EU's Joint Research Center has been providing real-time cereal yield predictions based on the Crop Yield Forecasting System with its MARS bulletin since 2007. Meteo France, the French weather agency, also provide three-month ahead weather bulletins which are made public on its website.

Finally, we address two key institutional features of French agriculture that could influence profit-based measures of adaptation: subsidies from the Common Agricultural Policy (CAP) and the national crop insurance scheme. During our period of interest, CAP subsidies were primarily distributed per hectare and thus should be largely insensitive to the weather forecast shocks we study. Agricultural insurance, while available, has a low uptake rate, with only 17% of total agricultural area being insured as of 2023.[8]

To formally account for the potential influence of these programs, we will compare results using different profit measures. The first is value added, capturing profits from production before subsidies or insurance transfers. The second is a broader measure, which includes these transfers. Comparing the results across these two measures allows us to test whether our estimates are driven by production decisions rather than by the institutional environment. We provide precise definitions for all variables in Appendix B.1.

3.2 Agricultural Data

3.2.1 Farm Data

Our analysis relies on farm-level panel data from the French section of the Farm Accountancy Data Network (FADN), an annual survey of approximately 7,000 farms. The FADN is designed to be statistically representative of French commercial agriculture, which accounts for the vast majority (95% in 2000) of the country's total agricultural production.[9] From this dataset, we select a subsample of farms specializing in cereal and oil crops, defined as those for whom these products constitute at least 50% of total annual sales. This focus is justified as these crops dominate large-scale field farming in France, and excludes distinct agricultural sectors like viticulture or animal husbandry.

The richness of the FADN is crucial for our analysis. For the selected farms, it collects data on farm profit, costs and revenue, which are crucial for our analysis of adaptation. The panel further contains data on input-specific expenses and crop-specific production, prices and land use, enabling us to explore specific adaptation margins.

We geocode the farms at the department level, and match them with department-level weather data.[10] While farms locations at the commune-level are available from 2000 onward, the forecast data is too coarse for such granularity. We choose to measure realized and forecasted weather at a similar level. We perform a robustness analysis over the 2000-2018 period, where realized weather is computed at the commune-level and forecasts kept at the department-one.

[8] In France, the main insurance scheme is one of harvest coverage, where extreme weather events leading to substantial losses in harvest lead to compensation. See https://agriculture.gouv.fr/animation-la-reforme-de-lassurance-recolte.

[9] The definition of a commercial farm changed in 2010. This change only affected rules for replacing farms leaving the panel and did not alter the composition of the existing sample, ensuring continuity for our analysis. Before 2010, a commercial farm was defined as a farm with a unique manager, selling more than half of their production, and whose manager's working hours correspond to at least 75% of their total annual work hours. Finally, farms with less than 5ha of land were removed from the targeted population if they were not specialized. In 2000, there were 380,000 such farms recorded in the Agricultural Census out of 663,800, but together they accounted for 95% of the country's total agricultural production. From 2010 onward, the working hours requirement was removed, and the 5ha threshold was replaced by a requirement that farms have a production capacity of at least €25,000.

[10] There are 101 departments in France, which make for slightly larger entities than US counties.

3.2.2 Input Price Data

To properly measure the costs of adaptation, our model requires accurate, location-specific prices for key agricultural inputs. We construct these measures using two official French data sources. Prices for land are derived from the Land Market Value survey (Valeur Venale des Terres).[11] Prices for seeds, fertilizers, and pesticides are sourced from the EPCIA survey, which is the official survey mandated by the European Commission to construct price indices for French agricultural goods.[12] We localise the sale points at the department level, and match them with department-level weather data.

We also use the Laspeyres price indices derived by the INSEE from the EPCIA in order to deflate the FADN farm-level input bills for seeds, fertilizers and pesticides. Specifically we use the Ipampa price index series, from 1994 to 2020. Figure A14 compares input price indices from Ipampa with a time series for nitrogen-based fertilizers producers prices from the Federal Reserve Economic Data, and shows the strong correlations in fluctuations. This confirms that the indices properly track variations in prices. Due to lack of data on water prices, we deflate irrigation expenses using a regular CPI index.

3.2.3 Plot Level Data

While our main analysis uses farm-level panel data to estimate the overall costs of adaptation, this approach cannot reveal adjustments in terms of timing, i.e. when specific actions are undertaken within the season. To investigate this channel, we leverage the "Pratiques Cultures sur les Grandes Cultures" (Agricultural Practices for Field Crops) survey, a plot-level dataset. We use the survey waves from 1994, 2001, and 2006.[13] It is important to note that this dataset is a repeated cross-section, not a panel. Consequently, our analysis of this data will not include farm or plot fixed effects and will rely on regional characteristics to control for confounding factors.

3.3 Climate Data

3.3.1 Weather Data

Our weather data is sourced from the European Centre for Medium-Range Weather Forecasts' (ECMWF) ERA5 reanalysis product, which provides hourly estimates of key climate variables on a 0.25° x 0.25° grid. We use hourly temperature and precipitation, aggregating the gridded data to the French department level by calculating an area-weighted average.

Using the time separability assumption common in the literature on agricultural climate impacts, we measure temperature exposure using growing degree days (GDD) and heating degree days (HDD). GDDs capture moderate temperature exposure beneficial for growth, while HDDs capture exposure to extreme heat. Crucially, to ensure consistency between our realized weather data and the forecast data used in our model, we construct these variables using only the four daily temperature measurements available in the

[11] The survey 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. It is then complemented by data provided by local notaries, and several local administrations. The data was digitized from scanned data magazines for the first years of the series.

[12] The EPCIA ensures representativeness by sampling goods in proportion to their category's sales and by sampling prices from firms in proportion to their market share for each good.

[13] 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.

forecast dataset (midnight, 6:00, 12:00, and 18:00). The resulting daily integrals are scaled to be equivalent to degree days calculated from 24-hourly data.[14]

We define a single, expansive growing season from October of the previous year to July of the current year. This choice is driven by two factors. First, as shown in Figure A13 and Table A1, farms in our sample are predominantly multi-product operations, often growing multiple crops with different sensitivity to heat. A broad temporal window is necessary to capture all potentially relevant weather shocks that affect a farm's decisions and profit. Second, this October-to-July window aligns perfectly with the growing season for winter wheat, the dominant crop in our sample (Gammans et al., 2017). We show in Figure A9 and Figure A10 the cross-sectional and cross-temporal variations in growing seasons for resp. wheat and corn in France.

We set our extreme heat threshold for HDDs at 30°C. This threshold is chosen to balance the heterogeneous heat tolerances of the dominant crops in our sample. While winter wheat suffers significant yield losses primarily at temperatures above 33°C (Gammans et al., 2017), corn—which accounts for 15% of the agricultural area we study—is known to be damaged by temperatures above 29°C (Schlenker and Roberts, 2009). Our 30°C threshold captures the onset of heat stress for a significant portion of the crops grown by farms in our data. We confirm the robustness of our results to this choice by testing an alternative 28°C threshold. GDDs are computed over the $[4°, 30°]$ degree interval.[15]

3.3.2 Seasonal Weather Forecasts

Our identification strategy requires measures of the farmer's information set at various points before the weather is realized. We aim to construct growing-season weather forecasts for different lead times (e.g., one-month ahead, two-months ahead). The raw forecast data, however, presents a structural challenge: new forecasts are issued only on the first day of each calendar month. This means that a forecast with a true, constant one-month lead (e.g., exactly 30 days) is not available for every day of the growing season.

We address this by defining our forecast variables based on their issuance month. We define a "one-month-ahead" forecast for any given day as the forecast that was issued on the first day of the prior month. Consequently, this variable represents the complete forecast information available to a farmer at a horizon of approximately one to two months. For example, the "one-month-ahead" forecast for both May 5th and May 25th is the one issued on April 1st. Similarly, a "two-month-ahead" forecast bundles all predictions issued between two and three months in advance. All one-month ahead forecasts are then aggregated across the growing season, to produce the forecast equivalent to the observed GDD and HDD measures, and the same procedure is applied to other leads. This approach ensures that we are using a consistent information set for each lead-time category.

The underlying forecast data comes from ECMWF's SEAS5 seasonal forecasting system (System 8, from Météo France). On the first of each month, an ensemble of 25 forecasts is produced for the subsequent seven months, which we average to create a single measure. The temporal granularity of this raw data—four temperature readings and one precipitation total per day—is used to construct the seasonal department-level values within our forecast aggregates, ensuring consistency with our realized weather variables.

We will also investigate the role of forecasts received when farmers are making planting decisions, for different crops, e.g. what farmers know about incoming weather on October 1st when they are supposed to sow wheat, or on May 1st when they should sow corn.

[14] As a robustness check, we replicate our analysis using realized degree days computed from the full set of 24-hourly ERA5 data.

[15] A robustness check includes freezing degree day to our analysis, in order to obtain a more comprehensive measure of realized weather over the season.

3.3.3 Forecast Accuracy

A necessary condition for our empirical strategy is that the seasonal forecasts contain meaningful information that can influence farmers' decisions. We confirm this in two ways. First, as shown in the calibration plots in Figure A4 and Figure A5, forecasted GDDs and HDDs are highly correlated with their realized values. For GDDs, the relationship is strong at both one- and two-month horizons, with observations clustering tightly around the 45-degree line, indicating that the forecasts are well-calibrated. While forecasted HDDs are also highly correlated with realizations, the plots show a tendency to underestimate the most extreme heat events.

Having established that the forecasts are informative, we now characterize their error distributions. As shown in Figure A1 and Figure A2, the distributions are generally centered around zero. For GDDs, we observe a slight negative bias and a long left tail, both corresponding to over-predictions. To ensure our estimates are not driven by tail outliers, we trim 120 observations where the forecast error is less than -500 GDDs.[16] The error distribution for HDDs exhibits a slight positive skew, consistent with the calibration plots showing that forecasts tend to under-predict extreme heat. Rainfall forecasts show a slight upward bias but are otherwise largely symmetric around zero. Overall, this analysis confirms that the forecasts provide meaningful information about future weather.

4 Empirical Strategy

To estimate the costs and revenue effects of ex ante adaptation, as well as test for the empirical validity of the envelope theorem, we use an estimating equation that regresses profits, total costs, or total revenue on both realizations and forecasts of weather. In particular, we use the following estimating equation.

$$\begin{aligned} y_{jt} =& \beta_1^w \text{GDD}_{d(j)t} + \beta_2^w \text{HDD}_{d(j)t} + \beta_1^f \text{FGDD}_{d(j)t} + \beta_2^f \text{FHDD}_{d(j)t} + \\ & g(P_{d(j)t}) + g_2(FP_{d(j)t}) + \gamma_j + \eta_t + \zeta^1_{r(j)} t + \zeta^2_{r(j)} t^2 + \varepsilon_{jt} \end{aligned} \tag{3}$$

The outcome variable is either profit, costs or revenues for farm j in growing season t. In later results, we also explore effects with different outcomes including farm inputs, planting decisions, and crop-specific land use. The main right-hand-side variables are realizations of temperature (GDD and HDD) and forecasts of temperature (FGDD and FHDD) experienced by farms located in department d during the growing season. Given that weather and forecasts vary at the department and year level, we use two-way standard errors at the department and year level.[17]

We focus on the marginal effects of forecasts, and do so for two reasons. First, conditional on weather realizations, variation in forecasts should cleanly identify changes in information available to farmers, and through that the consequences of their actions. When looking at outcomes such as revenue, rather than farm decisions, coefficients recovered off weather realizations will not only capture the effect of farmer action, but also of weather realization. Controlling for realizations, and using forecast variation isolates farmer decisions.

Second, forecasts can also better capture the information available to farmers regarding weather, and more cleanly isolate the one relevant for their decisions. Forecasts are a way to avoid the attenuation bias caused by recovering cost and revenue effects off weather realizations—which, to a minimum, will contain both predicted and unpredicted weather (farmers knowing and responding only to the known part of the realization). Additionally, if the farmer faces adjustment costs when choosing actions, then they have an incentive to choose actions prior to the arrival of weather. In such a case, forecasts provide more powerful identification of the effect of temperature on farmer actions than looking at realizations of temperature. In a

[16] Our main results are robust to the choice of trimming strategy, including trimming based on absolute error values. These results are available upon request.

[17] We later also show results with weather realizations measured at the municipality level.

farm setting, adjustment costs are likely high given that many actions need to be taken prior to the growing season (e.g., the choice of which crops to plant, total cropped area) or prior to weather arrival during the growing season (e.g., fertilizer application, defense of crops against freezing).

We focus, in the initial results, on one-month-ahead forecasts. In cases with convex adjustment costs, marginal value of information falls as forecast horizon increases. Thus, short-horizon forecasts should again improve power to detect effects. In additional results, we examine forecasts with longer horizons.

The estimating equation also includes controls for the level and square of realized and forecasted precipitation over the growing season to account for effects of precipitation on farm outcomes. We write these as $g(P_{d(j)t}) + g_2(FP_{d(j)t})$. We include region specific time trends in the form of $\zeta^1_{r(j)}t$ and $\zeta^2_{r(j)}t^2$. These account for potential sub-national trends that would correlate with weather and our outcomes of interest.[18] Finally, farm fixed effects, γ_j, and year fixed effects, η_t, mean that effects are identified from within-farm variation in weather over time, while accounting for national time series patterns in both weather and agricultural costs or revenues, as well as region-level quadratic trends. The identification assumption is that the remaining error term, ε_{jt} is uncorrelated with the temperature forecast variables. The control set is similar to prior work on the effects of climate on agriculture, with one important difference: we are able to use farm fixed effects rather than geographic area fixed effects (e.g., many studies in the U.S. include county fixed effects). This more granular cross-sectional control should alleviate concerns about confounding farm-level characteristics like geographic features that determine crop suitability and weather patterns.

5 Results: What kind of Adaptation?

In this section, we present evidence of costless revenue-increasing ex ante adaptation by farms in response to forecasted heat shocks. We first document the farm-level consequences of this adaptation on costs, revenues, and profits, and highlight a heterogeneous response to forecasts depending on the sign of the forecast error. We finally analyze the mechanisms driving this adaptation, and highlight the consequences our results bear for the implementation of the static envelope theorem.

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

5.1 Farm-level Consequences of Adaptation

Table 2: Farm-Level Reactions to Forecasts (1 month lead)

Dependent Variables:	Profit	Revenue	Costs
Model:	(1)	(2)	(3)
Variables			
GDD	-8.32	5.10	1.21
	(13.16)	(11.84)	(7.81)
GDD (F)	3.26	37.4	30.8**
	(43.1)	(34.7)	(14.5)
HDD	185	-158	-192
	(277)	(332)	(157)
HDD (F)	2,067***	1,192***	-56.7
	(708)	(403)	(178.5)
Mean	86,695	155,386	123,249
Unique Farms	2,603	2,603	2,603
Fixed-effects			
Farm	Yes	Yes	Yes
Year	Yes	Yes	Yes
Fit statistics			
Observations	18,917	18,917	18,917
R^2	0.84	0.89	0.94

Notes: Estimates are based on Equation 3, using the baseline sample. Realized and forecasted rainfall in levels and squares, as well as quadratic region-specific time trends, are included as controls in addition to the indicated fixed effects. Observations are weighted using the sample weights provided in the FADN. Two-way department-by-year standard-errors in parentheses. Stars indicate estimate is significantly different from zero: *$p < .10$, **$p < .05$, ***$p < .01$.

Consequences of Adaptation: We first show the effect of forecasted moderate and extreme temperature on farm profits, revenues and costs, controlling for realizations of temperature. Table 2 displays the results. The outcomes are profits, revenues (measured by total sales) and a measure of the costs of production. The measure of costs is broad and includes the cost for intermediate inputs, social contributions to workers, personnel expenses, taxes, and insurance. This broad measure is less likely to miss potential adaptation costs. The profit variable corresponds to the farm's gross operating income, and includes value added, subsidies, expenses for insurance, and insurance indemnities.[19]

GDDs, as the name implies, help crops grow and can thus be interpreted as positive productivity shocks. HDDs, in contrast, are temperatures so extreme that they cause crop losses. We expect that forecasts of GDDs are useful for taking advantage of better growing conditions, while forecasts of extreme heat are useful to either cut production costs, to increase input usage to compensate for adverse conditions, to better target input usage, or to modify decisions such as the timing of harvest or one's crop mix composition. In a simple model where farms optimize profit, and local adverse weather (either lower GDD or higher HDD)

[19]We perform a test and run the regressions using only expenses for intermediate inputs, as a check for potential mismeasurement, and find similar results. Results are available upon request. See subsection B.1 for the definition of intermediary inputs. We also consider a measure of value added as an alternative to profit, defined as the difference between sales and costs of production. Results are shown in Table A8.

corresponds to a negative TFP shock—but where output prices are not responsive to local conditions—we would expect forecasts of worse conditions to lead to a reduction in the scale of production, leading in turn to reduced revenues and costs. The expected negative TFP shock shifts the marginal cost curve upwards, and under decreasing returns to scale, induces farms to reduce their scale of production until it matches the market price again. On the other hand, changes in the timing of harvest could allow for a positive revenue response without substantial changes in costs.

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

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

The effect of forecasted HDDs shows a different pattern. Profits and sales increase while costs do not show substantial or significant responses. An increase by one degree-day in HDD forecasted one month in advance leads to an increase in profit by 2k€, and an increase in sales of 1k€. Given typical sales per farm of about 150k€ and a standard deviation of forecasted HDD of just over 1 (see Table A2), this coefficient indicates that a typical change in forecasted HDD causes sales to change by about 1%, on average. Similarly, given an average yearly profit of 87k€, a typical change in forecasted HDD leads to an increase in profits of 2%. In contrast, forecasts of extreme heat have no statistically significant effect on production costs. The difference in magnitude between costs and revenue responses is in line with the observed net increase in profits. These results indicate that within the period of the shock, there are large benefits to adaptation and near-zero costs of adaptation, on average, for these farmers.[20]

Forecast Error and Adaptation: It is useful to decompose these results by the sign of the forecast error (relative to the heat realization): whether forecasts under- or over-estimate the realized heat shocks of the season. We run a specification similar in all points to the one used for Table 2, but for the decomposition of the forecasted HDD variable—forecasted HDDs located in distance bins from the realized shock. We take a bin of $[-.1, .1]$ around the realized HDD shock, and use it as our omitted category in the regression. We then identify the median negative error (realization is larger than forecast), and median positive error, and construct bins (as denominated on the figure): $bin = 1$ for the below median negative errors (below the median value of negative errors), $bin = 2$ for above median negative errors (errors closer to the realization), $bin = 3$ for our omitted category, and $\{4, 5\}$ work similarly for positive forecast errors. These four dummies replace our $HDD(forecast)$ variable in our main regression, which remains otherwise identical. In the graph below we show the coefficients for these four variables, which then indicate the marginal profit effects

[20] Table A7 compares results to one-month and two-months ahead forecasts. Figure A16 compares the results for all available leads, each coefficient coming from a separate regression. Table A12 gives the profit results for one-month and two-months ahead leads, and compares results using the profit and the value added variables.

of one extra forecasted HDD in the specific bin. Complete cost, revenue and profit responses are shown in Table A33.

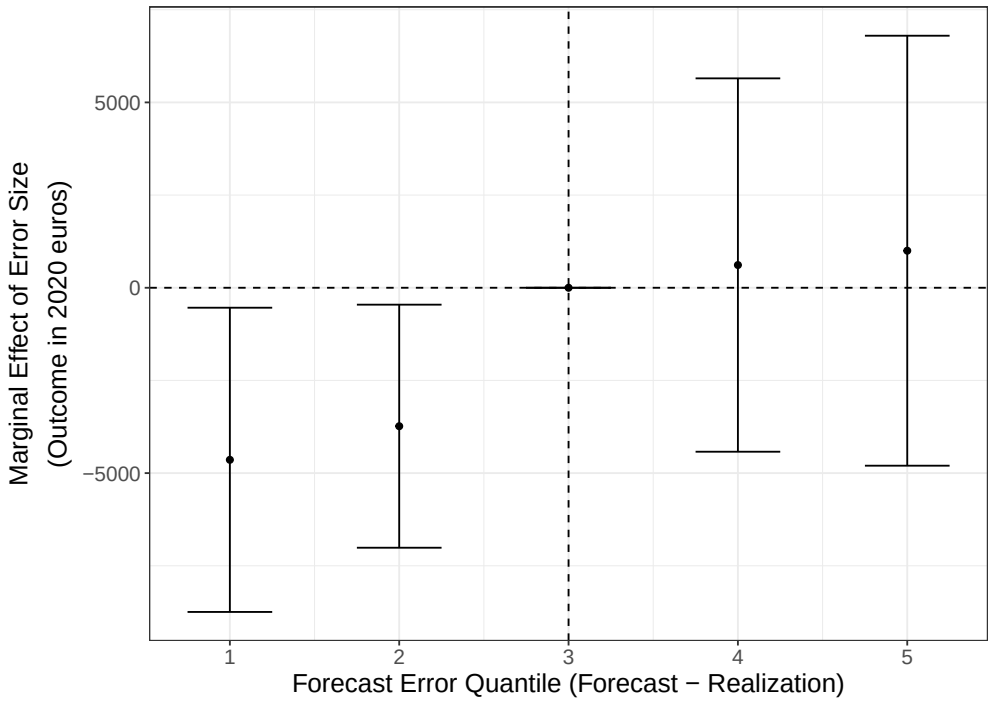


Figure 2: Profit & Sign of the Forecast Error

We see an interesting an intuitive pattern: negative forecast errors are statistically significant and have negative values, while positive ones are small in magnitude and not significant. Negative forecast errors are costly to farmers relative to on-point forecasts—they likely lead farmers to under-forecast future heat shocks and suffer from the full un-adapted impact of them. On the contrary, over-forecasts lead to no statistically significant gains in profit, relative to on-point forecasts. It seems that heat forecasts induce farmers to take protective actions which limit the negative consequences of realized heat on production. Adaptive actions in excess of the realized heat then matter little for the realized profit. Coming back to our main results in Table 2, an increase in forecasted heat—holding realized heat constant—will reduce the risk of a negative forecast error, and allow farmers to take the necessary protective actions and shield their profit from potential negative shocks.

5.2 What is Driving Farm-Level Adaptation?

5.2.1 Timing of Adaptation

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

$$y_{jt} = \beta_1^w \text{GDD}_{d(j)t} + \beta_2^w \text{HDD}_{d(j)t} + \sum_{\ell=1}^{5} \left(\beta_{1,\ell}^f \text{FGDD}_{d(j)t}^{\ell} + \beta_{2,\ell}^f \text{FHDD}_{d(j)t}^{\ell} + g_{2,\ell}(FP_{d(j)t}^{\ell}) \right) + \quad (4)$$
$$g(P_{d(j)t}) + \zeta_{d(j)}^1 t + \zeta_{d(j)}^2 t^2 + \gamma_j + \eta_t + \varepsilon_{jt}$$

where all variables are the same as in Equation 3 except we have added forecasts for each horizon, as indicated by the variables FGDD^{ℓ}, FHDD^{ℓ} and $FP^{\ell}_{d(j)t}$. This regression identifies the precise timing of information arrival, and is useful to understand what type of adaptation our analysis is recovering, i.e. which of these different information shocks drives our results on profit, costs and revenue. When running regressions with a unique set of forecasts, this set will capture the information received at other leads, via their non-controlled-for correlation with these other forecasts. Essentially, not including all the leads available creates a form of omitted variable bias where the included forecast captures a composite of the effects of all forecast horizons, with the composite effect being determined by the autocorrelation of forecasts across horizons. As long as we control for realized weather, this omitted variable will not be an issue for identifying benefits or costs of adaptation, because the interpretation of the forecast coefficients is still that it causes changes in the agent's belief, and through it, action. However, it is useful to include all the possible forecast leads in order to understand which one is most useful to farmers—in the sense that it generates the largest response.[21] It also puts the forecast and realization effects on similar footing in the sense that both are then identified by shocks: surprising realizations in the case of the realized temperature and news shocks in the case of all forecasts for horizons less than five months ahead.

Figure 3 plots the coefficients associated with heating degree days from estimating two different versions of Equation 4. In particular, the figure shows β^{f}_{ℓ} for $\ell \in [1, 4]$ and shows β^{w} for the lead value of 0 for regressions with costs and sales on the left hand side. We exclude the coefficients for lead 5, given that these are less cleanly identified, accounting for all the information received more than four months in advance.[22]

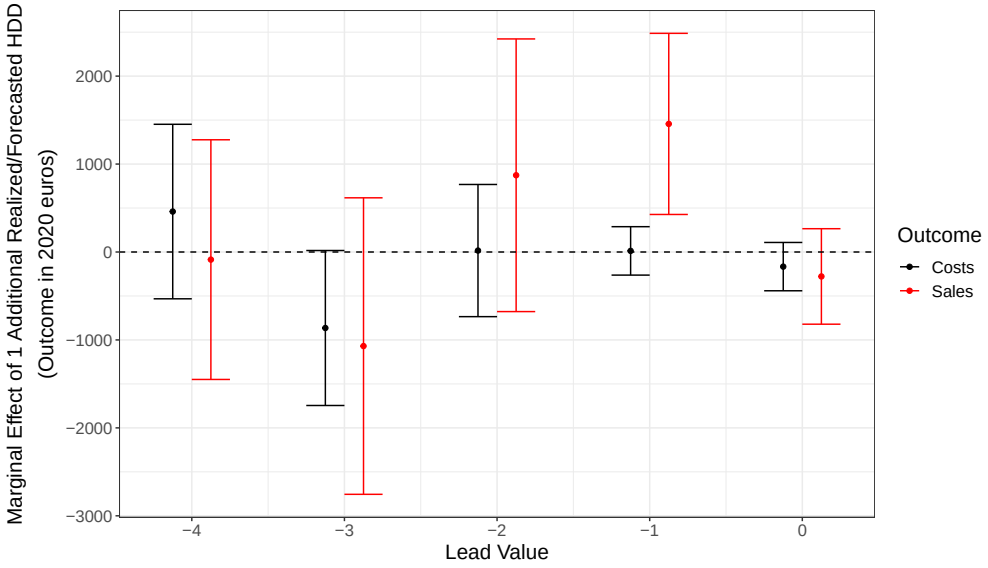


Figure 3: Costs and Revenue Reactions

We see that the the effect on sales associated with a one-month-ahead forecast is statistically significant and of a similar magnitude to the effect found in Table 2. The cost effect of a one-month-ahead forecast of

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

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

HDDs remains centered around zero. Revenue responses are not statistically significant at the 5% level for horizons longer than one month, but the point estimates indicate that information close to the shock allows for revenue-enhancing adaptation while information farther from the shock leads to a decrease in revenues. Costs are near zero for all horizons aside from three months ahead. The fact that costs do not move more strongly at short horizons than at longer horizons is inconsistent with adjustment costs driving the results that we find (Downey et al., 2023). The significant three-month-ahead forecast could indicate that there are adjustments that are uniquely available at a quarterly frequency that allow farms to reduce the scale of their operation in anticipation of adverse weather conditions.[23] We also note that the coefficients associated with realizations become more negative and almost statistically significant at the 5% level. This might mean that including more leads allows for more powerful estimation of the direct effect of extreme heat on agricultural revenues and costs. Overall, there results indicate that our analysis will be targeting the effects of up to three months-ahead ex ante adaptation to variation in expected heat. The same regression looking at profits confirm these results: the one month ahead revenue gain and zero cost response is associated with a positive significant coefficent, while lead 3 had cost and revenue responses of similar magnitude, and no significant profit response (Figure A15).

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

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

Before investigating the mechanisms by which firms might be adapting to temperature, we first rule out some alternative explanations of the results (Figure 4).

[23] As the French meteorological agency delivers every month a publicly available document with forecasts up to three month ahead, the three months ahead forecast is also the first easily observable by farmers, which could explain that farmers do not respond to four months ahead forecasts. See for an example of these publicly available and processed forecasts. https://meteofrance.fr/actualite/publications/les-tendances-climatiques-trois-mois.

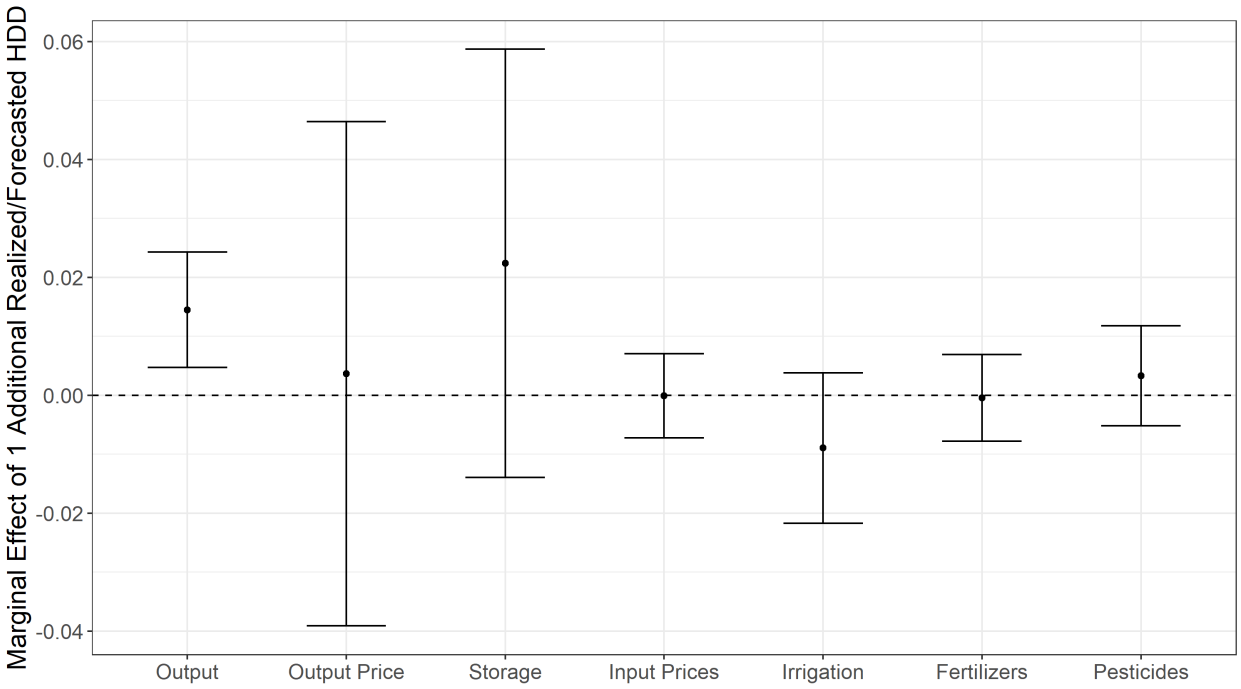


Figure 4: Prices and Quantities for (log) Outputs and Inputs

Offsetting Prices and Quantities: A first issue is that input quantities and prices responses might offset each other. On this, we show a separate regression focusing on input prices in Figure 4, and we further decompose the total farm input bill into input-specific ones. We consider the inputs easiest to adapt to one-month ahead information shocks: irrigation, fertilizers and pesticides.[24] The table corresponding to the figure is in Table A10. We use input prices as recorded in an agricultural input price survey used to build nationally representative input-specific Laspeyres indices. The precision of this survey is useful, in that we can run the regressions including store and product fixed effects in addition to year fixed effects. We show the results pooling together all categories of input prices here, and outline the price responses per input category in Table A15.[25]

We show in Table A13 results for a larger set of inputs, including land and labor directly measured in quantity, in both Figure A17 and Figure A18. Store prices are unresponsive to HDD forecasts, as are inputs. In Table A13, we see that none of land, labor, fertilizers, phytosanitary products, seeds and irrigation respond to forecasted HDDs.[26] Labor does respond to forecasted GDDs, matching the increase in total costs of production see in Table 2.

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

Measurement Error in Costs: A second issue is that the total costs variable could be subject to measurement error, leading to attenuation of the effect of forecasted HDDs. Total costs are measured as

[24] Fertilizers and phytosanitary products are measured as deflated input bills, using France-level input-specific Laspeyres price indices which translate the deflated input bills into 2020 euros. See subsection B.1 for detailed definitions of the different variables.

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

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

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

We further decompose land and labor responses per sub-categories of land ownership and labor contracts. Decomposing land responses in Figure A17, we see that this does not seem to be driven by an aggregation of different land categories with heterogeneous adjustment costs. Both the response of rented land with short rental contracts, and of owned and used land are non-statistically significant. Additionally, fallow land does not respond as well, whereas we could have considered it as a potential source of extensive margin adaptation. When decomposing labor responses in Figure A18, we note that the two most flexible types of labor respond negatively to expected hot weather, with non-regular salaried workers dropping in a statistically significant way by about 2.83 hours per additional forecasted HDD. This response remains small, given that labor is measured in total hours of work per year, and that the standard deviation in forecasted HDD is of one.

Substitution across Periods vs. Output Increases: Next, we ask whether the gains in revenue observed come from changes in storage practices at the farm level, or from actual increases in produced output. We focus on similar regressions for the quantity of output (simply the sum of all produced quantities), storage (sum of all stored quantities), and an output price index, also outlined in Figure 4 and Table A10.[27][28] First, within-period output does increase in a statistically significant fashion in response to forecasted HDDs, while output prices and storage are unresponsive. This confirms that farms are indeed undertaking adaptation actions which raise their production.

In Table A17, we decompose output effects at the crop level. These show that the gains are positive and significant for both wheat (one forecasted HDD increases output by 1.6%) and beets (.7%), while also positive but not significant for colza (1.4%). On the contrary, increases in forecasted HDDs decrease the amount produced in corn and sunflower, albeit in a non-significant way. We further note that while aggregate output does not respond positively to forecasted GDDs, the crop-specific decompositions performed in Table A17 show both wheat, corn and colza benefit from them, even though the coefficients are noisy. As expected, realized HDDs hurt production. In Table A9, we show the effect of two-months ahead forecasts.[29]

Further Robustness: Finally, we examine the robustness of the results to changes in the specification

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

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

[29] In our context, wheat shows no statistically significant relation with heat. In comparison, and using department-level time series of yields, Gammans et al. (2017) find that an additional one-day exposure to temperature above 32°C will decrease wheat production by about 2.5ppt. The comparison between these two results is not directly straightforward. Our analysis is done at the farm level, while theirs is at the department one, and it is not guaranteed that the elasticities of production to heat would be the same at the two levels. Our heating degree day cut-off is also at 30°C, in order to account for the impacts such temperature can already have on corn. We also account for weather realizations and forecasts over the entire possible growing season, allowing for farmers to both switch their crop mix composition (and hence the relative density of their land exposed to summer or winter growing seasons), and switch their growing season at the crop level (moving their wheat or corn specific growing seasons earlier or later for example). We also note that they focus on weather accrued during the warm part of the wheat growing season–although warm weather accrued over the total growing season, and the warm months should be comparable quantities in the French context. As such, it makes sense that they find a larger effect for unavoidable hotter weather, while we find that farmers face lower consequences form their endowed weather which might not necessarily be the effective weather they end up facing.

of the estimating equation including the addition of lags of realized weather or removing department-specific time trends. We vary the cutoffs used to define a GDD versus an HDD and add a measure of freezing degree days. We also use a disaggregated measure of realized weather, measured at the municipality level, combined with the same department-level measure of forecasts. In this case, realized weather is also computed using the entire range of hourly realizations, and not only four observations per day. And we investigate heterogeneity in the response to forecasts across farms. In all cases, we find results that are in line with the baseline estimates. These robustness checks are reported in Appendix D.

5.2.3 What Is Driving the Results? Heterogeneous Weather Variability, Endogenous Exposure to Weather and Crop Switching

What explains the different farm-level responses to forecasted GDDs and HDDs? In this section, we argue that the answer lies in two related mechanisms.

First, we show that the statistical properties of GDD and HDD shocks differ fundamentally. GDDs typically represent marginal variations in temperature, while HDDs represent non-marginal, tail-risk events. We posit that this distinction drives different types of adaptive behavior. In response to marginal GDD changes, farmers appear to make continuous adjustments to input use, leading to the observed effects on costs and revenues. In contrast, the threat of a non-marginal HDD shock prompts "lumpy" strategic responses—such as shifting the timing of harvest—which primarily involve opportunity costs rather than direct input costs, explaining the muted effect on our measured cost variable.

Second, we provide direct evidence that farmers use forecasts to endogenously manage their exposure to these shocks. We explore two key channels: (1) shifts in the timing of key agricultural activities within the growing season, and (2) changes in crop choice to alter the farm's fundamental sensitivity to heat.

Heterogeneous Variability along the Temperature Distribution: To test the hypothesis that the nature of the shock drives the nature of the response, we first decompose all temperature variations into "marginal" and "non-marginal" events. We define for this a non-marginal shock as any hourly temperature realization that deviates by more than 1.96 standard deviations from its historical hour-day-month mean.[30] Applying this definition to our data reveals a fundamental difference between our two temperature variables. As shown in Table A21, HDDs are overwhelmingly non-marginal events (85% of occurrences), while GDDs are overwhelmingly marginal (only 13% of GDD variation is classified as non-marginal).

This decomposition is the key to understanding our main results. We re-estimate our baseline specification, separating both forecasts and realizations into their marginal (m) and non-marginal (nm) components. The results, shown in Table A18, reveal a crucial distinction for forecasted extreme heat. The profitable adaptation response is driven entirely by non-marginal HDD forecasts. In contrast, the much rarer, marginal HDD forecasts lead to a statistically significant reduction in sales. This is precisely the pattern our hypothesis predicts: lumpy adaptations are only triggered by the threat of a significant, non-marginal shock. When the forecasted shock is marginal, firms engage in the more standard response of scaling back production. Given that the support for marginal HDD is small, our average effect for forecasted HDDs is mostly driven by the response to non-marginal changes in forecasted HDDs

In contrast to the sharp distinction we find for HDDs, the decomposition for GDDs is less conclusive. Costs also respond to non-marginal GDD forecasts, while they do not respond in a significant way to marginal GDDs. This result is likely well identified given the significant marginal component of GDDs in our data. We see however that the revenue and cost responses to marginal GDDs overlap, and that the responses to marginal and non-marginal forecasted GDDs cannot be strongly distinguished statistically.

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

Timing Decisions: The difference in variation underlying forecasted GDDs and HDDs suggests that farmers might also be taking different types of actions in response to forecasts of these different measures. In our second step, we investigate whether farmers change the timing of production actions in response to the forecasts. Such adjustments could explain the positive revenue and profit effects we find, as they can improve crop outcomes without large changes in input expenditures. To test this, we use plot-level data on the dates of ploughing, sowing, irrigation, and harvest for wheat from the PKGC survey.[31] As this is a repeated cross-section, we cannot include farm fixed effects and instead use department fixed effects along with region-by-year-clustered standard errors. Our specification is otherwise identical to the baseline in Equation 3.[32]

The results, shown in Table A19, confirm that farmers adjust their production schedule in response to heat forecasts. An anticipated HDD shock leads farmers to bring ploughing and sowing dates forward, though the effect on sowing is not statistically significant. In contrast, the first irrigation is delayed, perhaps to conserve water that is expected to become scarce. While the magnitudes of these shifts are modest, they demonstrate that timing is an active adaptation margin. Here, changes in the timing of irrigation could shift a farm's sensitivity to heat, while keeping exposure constant. On the contrary, changes in sowing and harvesting could change exposure, while keeping sensitivity constant.

The effect of these timing changes on weather exposure in an empirical question. We quantify it using a simple simulation. First, we compute the baseline distribution of temperatures experienced by farms over a fixed October-July growing season. We then use the coefficients from Table A19 to predict how the start and end dates of the season shift in response to each department-year's specific HDD forecast. Finally, we compute a new "endogenous" temperature distribution based on these adjusted season boundaries. Figure A20 compares the cumulative distribution functions of the endogenous and baseline seasons. The results show that the timing adjustments, while small, systematically shift the distribution in the expected direction. Farmers alter their effective growing season to increase exposure to cooler temperatures and reduce exposure to the hottest temperatures.

Overall, these timing responses represent a plausible channel for the "no-cost" adaptation we observe. They allow farmers to mitigate potential damages and secure higher revenues without significantly increasing input costs. However, the modest size of these shifts suggests they are unlikely to be the sole explanation for the large profit effects in our main results.

Crop Switching: We finally explore crop switching as a potential adaptive strategy. Crop switching has been identified as crucial to reduce the damages from heat exposure and climate change in agriculture (Rising and Devineni (2020)). Like timing, crop switching can also play along the sensitivity and exposure margins. On sensitivity, one can switch towards heat resilient crops when expecting hotter weather, hence reducing the losses faced when actually encountering bad weather. Corn and wheat, colza and sunflower have different heat tolerance, as highlighted by the specifications in Figure 1—and moving from corn to wheat and sunflower to colza might reduce one's sensitivity to heat shocks. Switching crops can also change one's exposure to weather. We show in Figure A9 and Figure A10 that typically, wheat is sown around November, and harvested before August in France, while corn is sown around May and harvested around October. Sowing wheat when a hot summer is expected can allow you to avoid August heat shocks, hereby reducing the effective amount of HDDs faced.

[31] See subsubsection 3.2.3 for a description of the survey.

[32] We cluster the standard errors at the region and year levels to account for the coarser sampling procedure used to build this data.

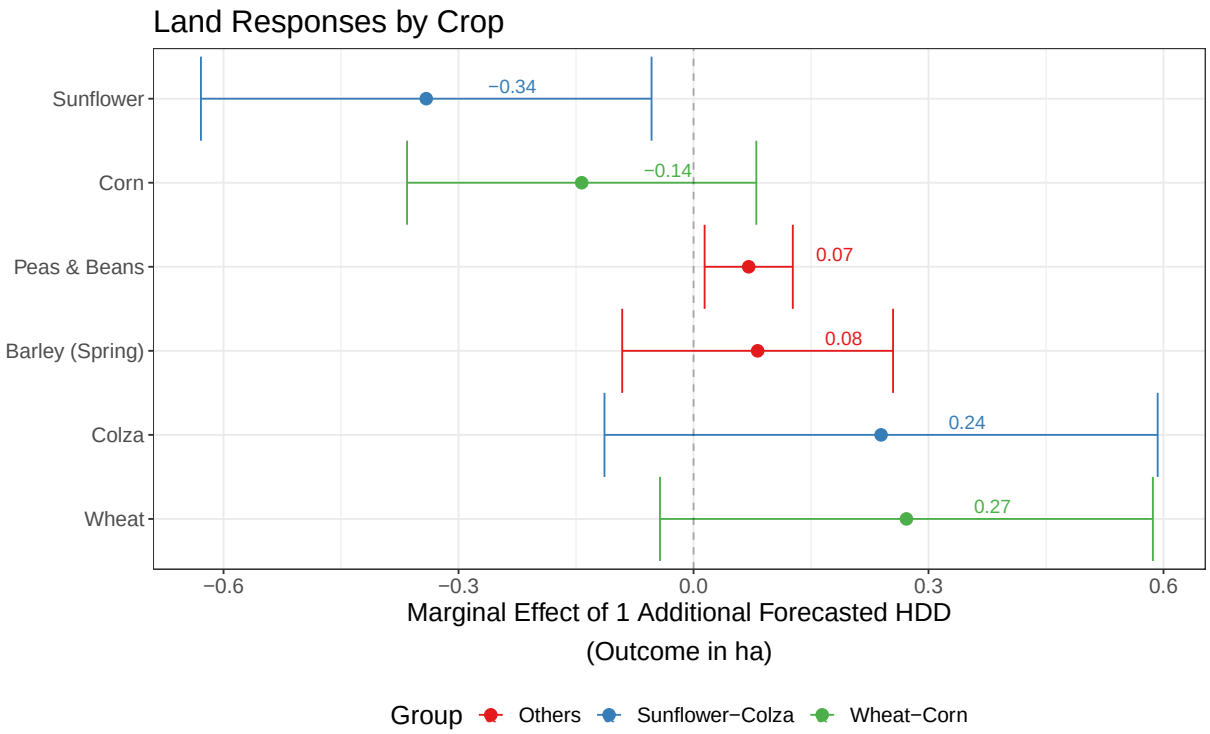


Figure 5: Decomposition of Cropland Responses

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

To test for this channel, we estimate our baseline model using the land area dedicated to specific crops as the outcome variable. Figure A17 plots the coefficients for one-month-ahead HDD forecasts. The results show a clear pattern consistent with adaptation: in response to forecasted heat, land allocated to heat-sensitive crops like corn and sunflower decreases, while land allocated to more heat-resilient crops like wheat, colza, barley, and legumes increases. Land use patterns are in line with output results we discussed earlier (Table A17). The increase in land dedicated to peas and beans, for example, matches cropping patterns observed by Aragón et al. (2021).

To formally test these substitution patterns, we use a system of OLS regressions and perform Wald tests on the coefficients for key pairs. The results, shown in Table A24, provide clear evidence of switching between oilseed crops; we can reject the null of an equal response for colza and sunflower at the 5% level (p=0.048). For cereals, the pattern is suggestive but less precise. While the point estimates in Figure A17 show a clear reallocation away from corn and towards wheat, a formal test fails to reject the null that their responses are equal (p=0.29). This suggests that while farmers are actively reallocating land away from heat-sensitive crops in response to forecasts, the substitution between winter wheat and summer corn is either less pronounced or measured with more noise in our data than the substitution between colza and sunflower.

5.3 Implications for the Static Envelope Theorem

Framework and Implications: Our results have direct and significant implications for the application of the static envelope theorem in climate economics. As shown in Table 2, forecasts of extreme heat (HDDs) have

a significant positive effect on farm profit (confidence interval of [.68k, 3.5k]), while forecasts for moderate heat (GDDs) have a null effect (confidence interval of [-81,88]). This finding—that adaptation to forecasted extreme weather has a non-zero, positive effect on profits, is a direct rejection of the theorem's prediction in this context. The results are robust to using value added as an alternative profit measure (Table A8).

The implication is that standard methods for estimating climate damages may be biased. Regressions of profits on realized weather implicitly assume that any adaptive responses have, at the margin, a zero effect on profits. Our findings show this is not the case for extreme heat. Instead, the coefficient on realized heat in a standard regression will capture the sum of the true, direct damage and the offsetting profit gain from adaptation. This omitted variable bias will lead to a significant underestimation of the true damages from extreme heat. We demonstrate this bias in Table A12, where the coefficient on realized HDDs changes significantly when forecast variables are included in the specification.

Conditions for the Application of the Envelope Theorem: Why does the static envelope theorem fail here? The theorem's validity rests on two key assumptions: (1) the firm's optimization problem is differentiable, implying smooth, continuous adjustments; and (2) the shocks are marginal. Our previous results suggest neither assumption is met.

First, the adaptation mechanisms we identify are inherently "lumpy". As shown in subsubsection 5.2.3, farmers respond to heat forecasts by making discrete decisions about harvest timing and crop choice. These actions involve opportunity costs but do not necessarily appear as continuous changes in input expenditures, violating the differentiability assumption. They are also discrete in nature, as a choice among a set of days or a set of crops.

Second, the shocks themselves are not marginal. Our decomposition analysis revealed that HDDs are overwhelmingly non-marginal. The envelope theorem is not expected to hold for such large shocks. This suggests that decomposing weather variation into its marginal and non-marginal components could be a crucial robustness check for studies in this literature, as a divergence in the estimated effects could signal a failure of the theorem's underlying assumptions.

A third possibility is that the costs and benefits of adaptation are distributed unevenly over time, meaning the theorem may hold dynamically but not statically. We explore this possibility in the next section.

6 What are the Consequences of Adaptation?

Our analysis has so far revealed a puzzle: static, within-period adaptation to forecasted extreme heat (HDDs) is significantly profitable, a direct contradiction of the static envelope theorem. This raises a crucial question: are these profit gains real, or are they offset by costs in subsequent periods? In this section, we investigate this dynamic trade-off, asking whether a dynamic version of the envelope theorem holds where the static theorem fails.

To do so, we examine the inter-temporal consequences of adaptation, estimating its effect on the stream of profits and costs over time. We find compelling evidence that the initial profit gains from adapting to heat forecasts are not persistent. Instead, they are followed by significant increases in production costs in subsequent periods. These delayed costs entirely offset the initial gains, rendering the net effect on the present discounted value of farm profits statistically indistinguishable from zero. This suggests that while adaptation to extreme events appears profitable in the short run, its overall effect is neutral, consistent with the predictions of a dynamic envelope theorem. We then detail the sources of these delayed costs and discuss their implications for crop switching costs.

6.1 Framework

To test for the dynamic effects of adaptation, we estimate the impact of a weather forecast issued in year t on farm outcomes in the years before, during, and after the forecast, from $t-2$ to $t+3$.

The primary challenge in this dynamic setting is to isolate the causal effect of the forecast at time t on outcomes at time $t+h$. This requires ensuring the forecast is not correlated with other factors that determine outcomes in the surrounding years. The main threats are potential correlations between the forecast at t and weather shocks or forecasts in other periods. Under the assumption that the forecast is unbiased and informationally redundant once the realized weather at time t is known, controlling for current weather will remove the correlation between the time t forecast and both future weather or forecast values.

However, controlling for realized weather at time t is not sufficient to remove the potential correlation between the time t forecast and past weather events, which could in turn bias our estimates. Therefore, to isolate the effect of the new information arriving at time t, we must explicitly control for past weather realizations. Our empirical design implements this strategy. Our estimating equation for outcome y for farm j in year $t+h$, where $h \in [-2, 3]$, is:

$$\begin{aligned} y_{j,t+h} = \sum_{k=-2}^{0} & \beta^{w}_{1,h,k} GDD_{d(j),t+k} + \beta^{w}_{2,h,k} HDD_{d(j),t+k} + g^{h,k}\big(P_{d(j),t+k}\big) + \\ & \beta^{f}_{1,h} FGDD_{d(j),t} + \beta^{f}_{2,h} FHDD_{d(j),t} + g^{h}\big(FP_{d(j),t}\big) + \\ & \zeta^{1}_{d(j)} t + \zeta^{2}_{d(j)} t^2 + \gamma_j + \eta_t + \varepsilon_{jt}, \end{aligned} \tag{5}$$

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

We provide evidence to support our identification strategy with a falsification test. If our controls are sufficient to absorb confounding information from other periods, then the HDD forecast at time t should not predict forecasts for other years. We test this by regressing lead and lagged forecasts on the time t forecast, conditional on our full set of controls from Equation 5. The results in Table A22 show that the coefficient on the time t forecast is statistically indistinguishable from zero in all cases. This demonstrates that, conditional on our controls, the forecasts are not serially correlated, alleviating concerns that our estimates of $\{\beta^{f}_{2,h}\}_h$ are biased by unobserved weather information from other periods.

6.2 Profits

We now present the main results of our dynamic analysis, focusing on the effect of ex ante adaptation on the stream of farm profits over time. Figure 6 plots the dynamic response of profits to a forecasted HDD shock at time t. The table of results is available in Table A23, and coefficients reveal a clear pattern of intertemporal trade-offs. We find a positive and significant contemporaneous profit effect of $+1.9k$€. This initial gain, however, is followed by statistically significant profit losses in the subsequent two periods, with effects of $-1.8k$€ in year t+1 and $-1.4k$€ in year $t+2$. The effect in year $t+3$ is small and not statistically significant.

This dynamic pattern allows for a direct test of the dynamic envelope theorem. We calculate the net present value (NPV) of the profit stream over the $[t, t+3]$ horizon, using a 10% discount factor. The resulting point estimate for the total effect of adaptation is $-1.1k$€, and its 95% confidence interval of $[-3.7k, 1.4k]$ includes zero. Relative to the average discounted profit over this period ($298k$€), this represents a negligible

change. Therefore, we cannot reject the null hypothesis that the dynamic envelope theorem holds. While adaptation is profitable within a single period, its net effect on the present value of profits is zero.

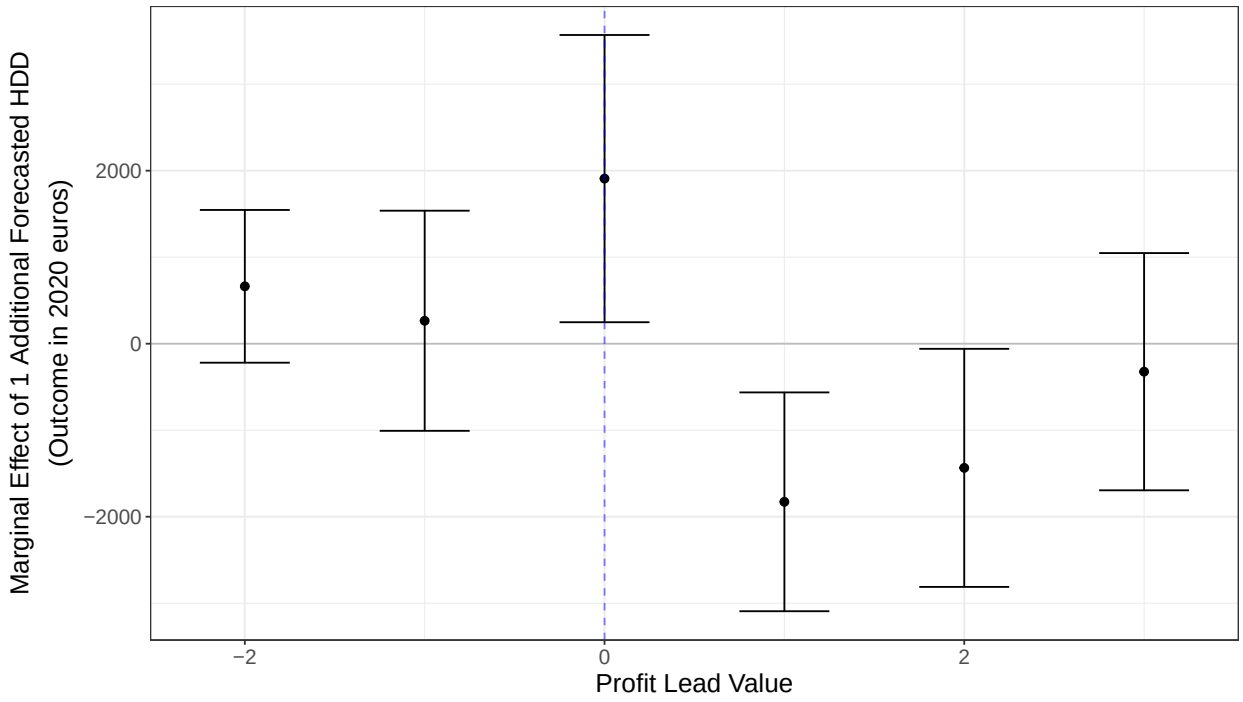


Figure 6: Dynamic Effects of ex ante Adaptation on Profit

This finding suggests that farms face a trade-off. Assuming farmers are rational, the value of adapting must be greater than or equal to the value of facing heat without adaptation, implying that the counterfactual of foregoing adaptation would have been costly, and likely leading to a loss in the NPV of profit. The dynamic profit profile is consistent with adaptation strategies like crop switching, which we identified as a key mechanism. Switching away from a farm's long-run optimal cropping pattern can generate immediate gains by avoiding heat damage, but likely incurs future costs associated with soil health, input requirements, or disease exposure, leading to the subsequent profit losses we observe. We explore these delayed costs in the next section.

6.3 Evidence for Dynamic Costs of Crop Switching

We next examine the dynamic pattern of production costs to understand the source of the intertemporal profit trade-offs. The results are shown graphically in Figure 7 and in table form in Table A25.

The cost results closely mirror the profit findings. Consistent with our static results, there is no significant change in production costs in the year of the forecast (t). However, this is followed by a statistically significant cost increase of $1k€$ in year $t+1$. This pattern of delayed cost increases strongly suggests that adaptation strategies like crop switching, while profitable upfront, have future financial consequences. The benefits (avoided heat damage) and costs (disrupted production patterns) are realized at different times.

Quantitatively, the present discounted value of this cost stream is significant. Summing the coefficients over the $[t, t+3]$ horizon with a 10% discount rate, we estimate the NPV of the future costs of adapting to a one-HDD forecasted shock to be $1.1k€$.

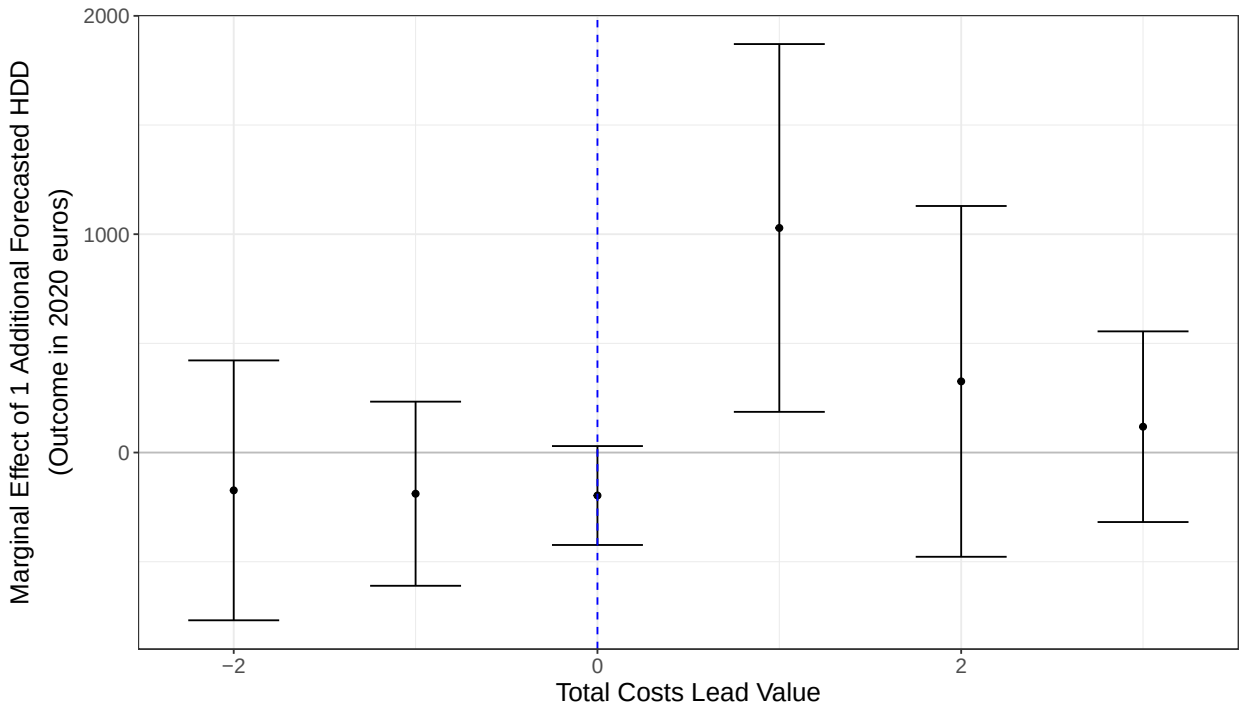


Figure 7: Dynamic Effects of ex ante Adaptation on Production Costs

These dynamic costs almost perfectly explain the dynamic profit effects found in the previous section. The Net Present Value of the profit stream was $-1.13k€$. The Net Present Value of the cost stream is $1.08k€$. Therefore, the delayed increase in production costs accounts for over 95% of the total present-value profit loss. In our context, the net welfare consequence of a forecasted heat shock arises not from direct, un-adapted damages, but almost entirely from the future costs of the adaptive actions taken to avoid those damages. This implies that farms in our sample are well-adapted to the direct effects of heat, but this adaptation is not costless.

This interpretation, that adaptation via crop switching incurs delayed costs, is well-supported by the literature on dynamic land use. Foundational models by Livingston et al. (2008), Scott (2013), and Burlig et al. (2024) show that deviating from an optimal crop rotation has future productivity consequences. More specifically, such disruptions can decrease soil nutrients and increase exposure to pests (Vanino et al. (2019)). This suggests a clear, testable hypothesis: the aggregate cost increases we observe should be driven by future increases in specific inputs like fertilizers, pesticides, and the labor required to apply them. We test this directly in our final analysis by decomposing the dynamic response by input type, using the same specification as in Equation 5.

The results, shown in Figure 8, confirm this hypothesis and reveal a clear pattern of delayed input intensification. We find no significant change in input use in the year of the forecast (t). However, in year $t+1$, we observe statistically significant increases in expenditures on both labor and pesticides. The effect on pesticide use persists into year $t+2$, and the increase in labor re-emerges in year $t+3$. We also find a positive, though not statistically significant, increase in fertilizer use at $t+1$.

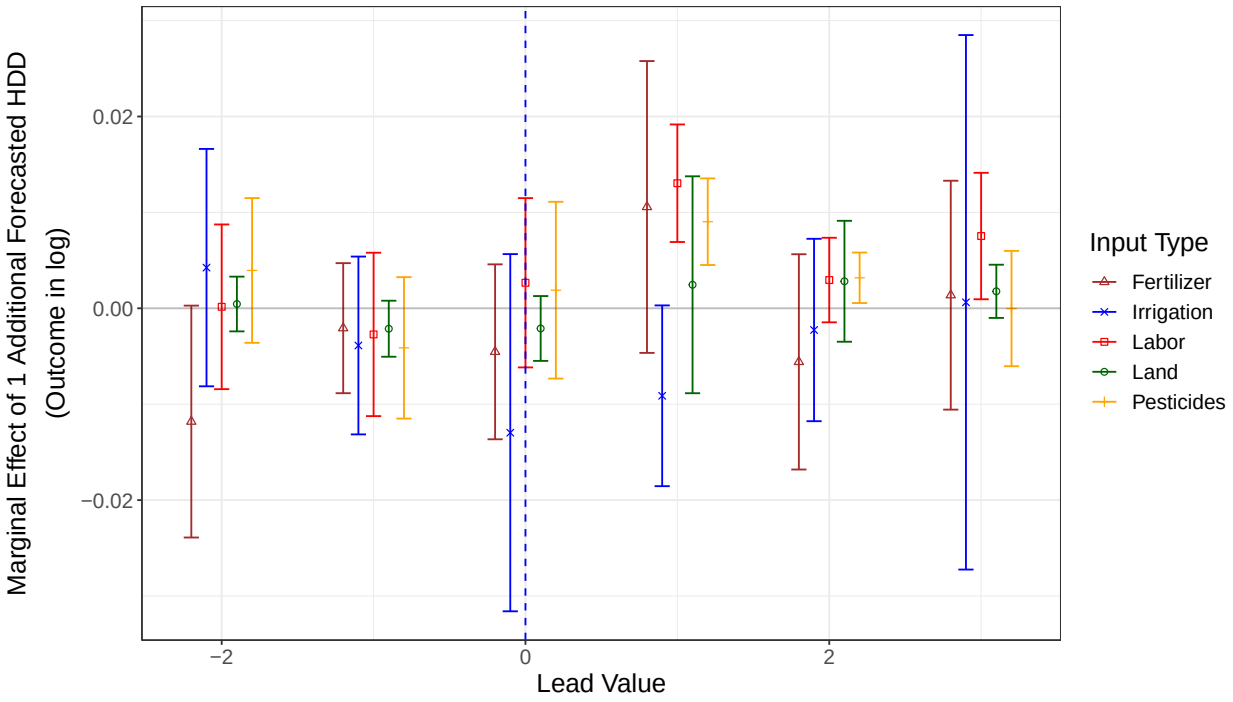


Figure 8: Dynamic Effects of ex ante Adaptation on (log) Input Use

These findings provide direct, reduced-form evidence for the mechanisms highlighted in the dynamic land-use literature. The crop switching that is profitable in the short run appears to be associated with a future decrease in land productivity. Farms then compensate for this deviation from their optimal cropping pattern by increasing their use of pesticides, fertilizers, and the labor required to apply them in subsequent years.

7 Implications for Climate Change

In this final section, we explore the implications of our findings for the long-run costs of adapting to climate change. The total cost of adaptation to climate change can be defined as the difference between the cost of adapting to a future climate versus the cost of adapting to the current one. Estimating this is challenging for two main reasons. First, future climate conditions are themselves projections. Second, and more critically, the cost function for adaptation is unlikely to be stable over time. As the climate changes, and more crucially as farmer expectations about climate change, they may adopt entirely new strategies, and the costs of adaptation strategies may evolve. To provide a credible, bounded estimate in the face of these challenges, we conduct two distinct exercises.

First, we perform a straightforward projection under a constant-cost assumption. We use projections of future heat exposure for France and apply our empirically estimated marginal cost of adaptation. This exercise calculates the future costs of adaptation under the strong assumption that the way farmers respond to heat, and the cost of that response, remains the same as it is today. This provides a valuable, albeit partial, baseline estimate.

Of course, it is unlikely that adaptation costs will remain constant. Our second exercise tests this assumption by using spatial variation in local climates as a proxy for temporal change in climate, following the Ricardian cross-sectional tradition (Auffhammer (2018), Carleton et al. (2022)). We compare the marginal

cost responses to heat across French departments that currently have different climates (i.e., higher vs. lower average temperatures). This cross-sectional comparison provides suggestive evidence of how adaptation costs might evolve as cooler regions become hotter. We stress that, as is common with this approach, these results are not causally identified, as regions with different climates likely differ along other unobserved dimensions. Nevertheless, they provide a useful lens through which to consider the potential state-dependence of adaptation costs.

7.1 Costs of Adaptation to Heat under Climate Change and Constant Costs

In section 6, we estimated the average farm-level cost of adapting to a one-HDD increase to be 1.1k€. Given a baseline average of 2.4 HDDs per year in our sample period, this implies that current adaptation to extreme heat costs farms approximately 2.6k€ annually, or 2.9% of their yearly profit. We now use this empirically-estimated marginal cost to project the future costs of adaptation under a climate change scenario, holding the cost function itself constant.

Our methodology involves several steps. First, we use climate projections from the ECMWF CMIP6 model under the SSP3-7.0 scenario, a middle-range pathway. Second, we correct these projections for systematic biases by comparing them to ERA5 historical data on their 2015-2024 overlapping period (Fisher et al. (2012)).[33] The distribution of these biases is shown in Figure A12. Finally, we use the method of Snyder (1985) to interpolate daily minimum and maximum projection values into hourly realizations, from which we calculate our projected HDD variable.

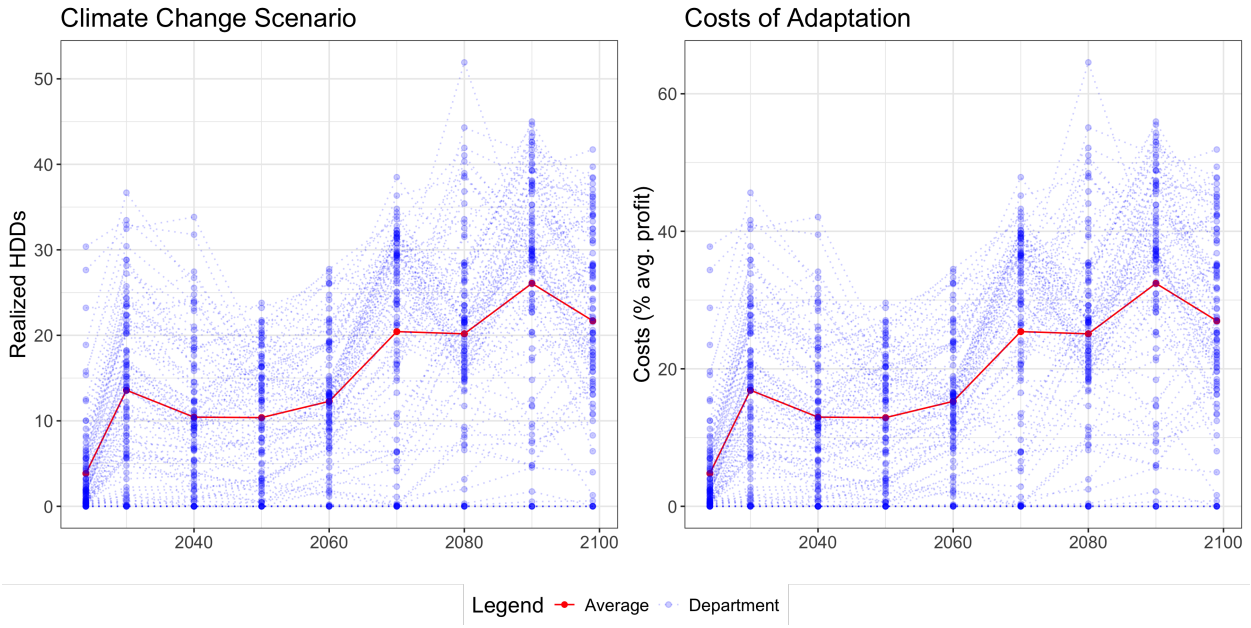


Figure 9: Costs of Adaptation Under Climate Change

The left panel of Figure 9 plots the projected decadal evolution of average HDDs in France. The number of HDDs is projected to rise from a baseline of 2.4 to approximately 22 by the end of the century. To calculate the cost of adapting to this change, we multiply the projected increase in HDDs (19.6) by our constant marginal cost of 1.1k€. This implies that by the end of the century, the annual cost of adapting to extreme heat will increase by approximately 21.6k€ per farm. This projection, while reliant on the strong

[33] See the discussion on debiasing in https://climate.copernicus.eu/sites/default/files/2021-01/infosheet7.pdf. Here we compute the average bias at the minimum and maximum day-of-the-month level,

constant-cost assumption, highlights the substantial financial burden that adaptation to future heat extremes may impose on the agricultural sector.

7.2 Suggestive Evidence for Long-Term Adaptation

A central limitation to the exercise conducted in the previous section is that costs of adaptation are unlikely to remain constant as climate change keeps unfolding. We could expect both that current adaptation channels will become inaccessible to farmers in the future—for example a farmer only growing wheat and colza cannot further switch its crop mix composition. On the contrary, one can also imagine that costs of adaptation to changes in climate will be lower than the costs of adaptation to weather variation today. The presence of learning-by-doing could reduce costs in the future, for example improvements in irrigation techniques.

Here we use a cross-sectional analysis frequently used in the climate change literature to test the sensitivity of costs of adaptation to variations in climate[34]. Specifically, we compute the department-level average temperature over 1994-2018. We then run a specification similar to the one in Equation 3, to which we add the interaction of all temperature related variables with this average temperature.

We are specifically interested by the coefficient for the interaction between forecasted HDDs and average temperature, or how the marginal response rate of farm profit, revenue and costs, changes along with the average climate faced by these farms, This analysis is correlational in nature, but suggestive of how farms might change their adaptation strategies with climate.

Results are shown in Table A26. Including these interactions has two consequences. The first one is that the coefficients associated to forecasted HDDs in levels become much larger in value for both the profit and the revenue regressions. Second, the coefficients associated to the interacted terms is negative and also statistically significant. These suggest that there are potentially decreasing returns to adaptation as the average temperature increase, or that the current adaptation strategies implemented in French farming are mostly relevant in moderate climate conditions. For example, it might be the case that farms in hot regions of France already grow little corn and colza, and have little room for adapting their crop mix to expected heat variations.

8 Conclusion

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

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

[34] See a discussion of this method described as a combination of the Ricardian and panel approaches in Auffhammer (2018).

This data also allows us to underline how the presence of these net profit gains from additional weather information contradict the assumptions of the static envelope theorem frequently used in the climate literature. Indeed, in our context, farmers' response to small variations in forecasted heating degree days have net positive profit impacts. We decompose the variation in weather used to build growing and heating degree days, and show how heating degree days are mostly composed of non-marginal variation ill-suited to the envelope theorem. Furthermore, the presence of no-cost adaptation strategies like timing, imply that some types of adaptation responses are likely to have non-continuous effects of profits, again going against the usual assumptions of the envelope theorem. We show however that the dynamic envelope theorem seems to hold: while farms face profit gains in the period of adaptation, the next periods are marked by significant decreases in profit—leading the long-term present discounted impact of adaptation to converge to zero. Overall, this paper demonstrates that understanding adaptation requires a dynamic perspective, and to pay attention to the intertemporal trade-offs undertaken by farms when they adapt. This time structure is driven by costs of adaptation which are faced by farms in future periods, and seem to correspond to the long-term costs of crop switching.

In a last section, we investigate how the costs of adaptation are likely to change with climate change. Under constant costs of adaptation, costs are expected to be multiplied by about 9 by the end of the century. We conclude by providing suggestive evidence that the effectiveness of current adaptation strategies will decline as temperature rises.

References

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

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

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

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

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

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

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

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

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

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

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

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

Colmer, J. (2021): "Temperature, Labor Reallocation, and Industrial Production: Evidence From India," *American Economic Journal: Applied Economics*, 13, 101–124.

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., M. Greenstone, and J. S. Shapiro (2017): "Defensive investments and the demand for air quality: Evidence from the NOx budget program," *American Economic Review*, 107, 2958–2989.

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

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

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

Fisher, A. C., W. M. Hanemann, M. J. Roberts, and W. Schlenker (2012): "The Economic Impacts of Climate Change: Evidence From Agricultural Output and Random Fluctuations in Weather: Comment," *American Economic Review*, 102, 3749–3760.

Freeman, A. M., J. A. Herriges, and C. L. Kling (2014): *The Measurement of Environmental and Resource Values: Theory and Methods*, New York: Routledge, 3 ed.

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

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

Hagerty, N. (2022): "Adaptation to surface water scarcity in irrigated agriculture," *Unpublished, Working Paper*.

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

Hultgren, A., T. Carleton, M. Delgado, D. R. Gergel, M. Greenstone, T. Houser, S. Hsiang, A. Jina, R. E. Kopp, S. B. Malevich, et al. (2022): "Estimating global impacts to agriculture from climate change accounting for adaptation," *Available at SSRN 4222020*.

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*.

Lemoine, D. (2024): "Estimating the consequences of climate change from variation in weather," *NBER Working Paper*, 25008, 41.

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

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

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

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

Ortiz-Bobea, A. and R. Just (2013): "Modeling the Structure of Adaptation in Climate Change Impact Assessment," *American Journal of Agricultural Economics*.

Reidsma, P., F. Ewert, A. O. Lansink, and R. Leemans (2010): "Adaptation to climate change and climate variability in European agriculture: the importance of farm level responses," *European journal of agronomy*, 32, 91–102.

Rising, J. and N. Devineni (2020): "Crop Switching Reduces Agricultural Losses from Climate Change in the United States by Half under RCP 8.5," *Nature Communications*.

Rosenzweig, M. R. and C. Udry (2014): "Rainfall Forecasts, Weather, and Wages Over the Agricultural Production Cycle," *American Economic Review Papers & Proceedings*, 104, 278–283.

Scheinkman, J. A. and J. Schechtman (1983): "A Single Competitive Model with Production and Storage," *The Review of Economic Studies*.

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

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

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

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

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

Snyder, R. (1985): "Hand Calculating Degree Days," *Agricultural and Forest Meteorology*.

Vanino, S., R. Farina, R. Hüppi, B. Pennelli, and J. Six (2019): "A Toolkit of Indicators for Characterising Productivity and Sustainability," *Report*.

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

A Additional Figures and Tables

A.1 Figures

A.1.1 Precision of Forecasts

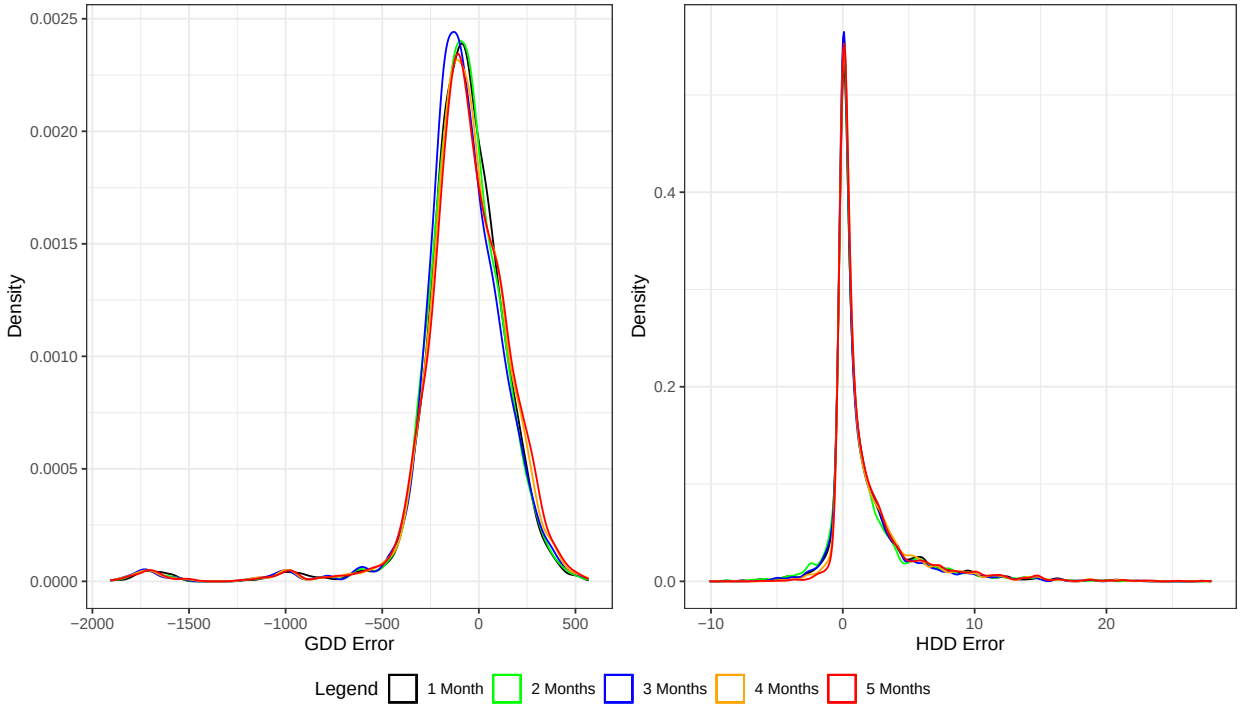


Figure A1: Distribution of Temperature Forecast Error

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

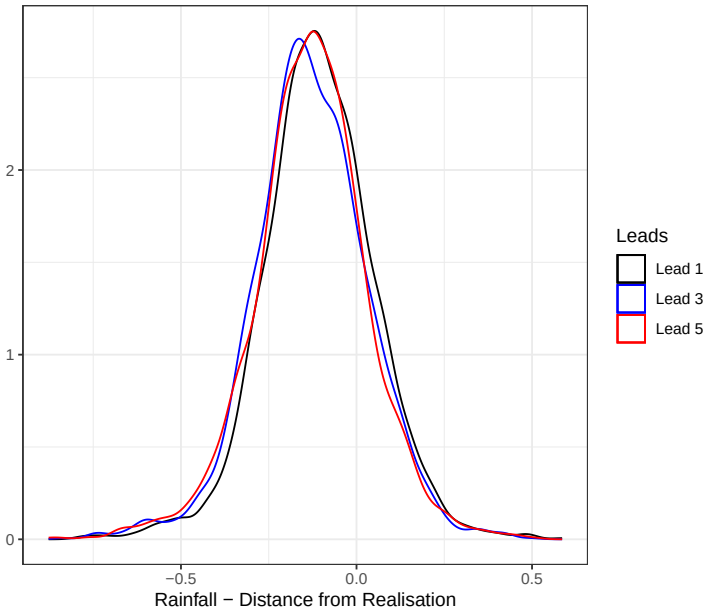


Figure A2: Distribution of Rainfall Forecast Error

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

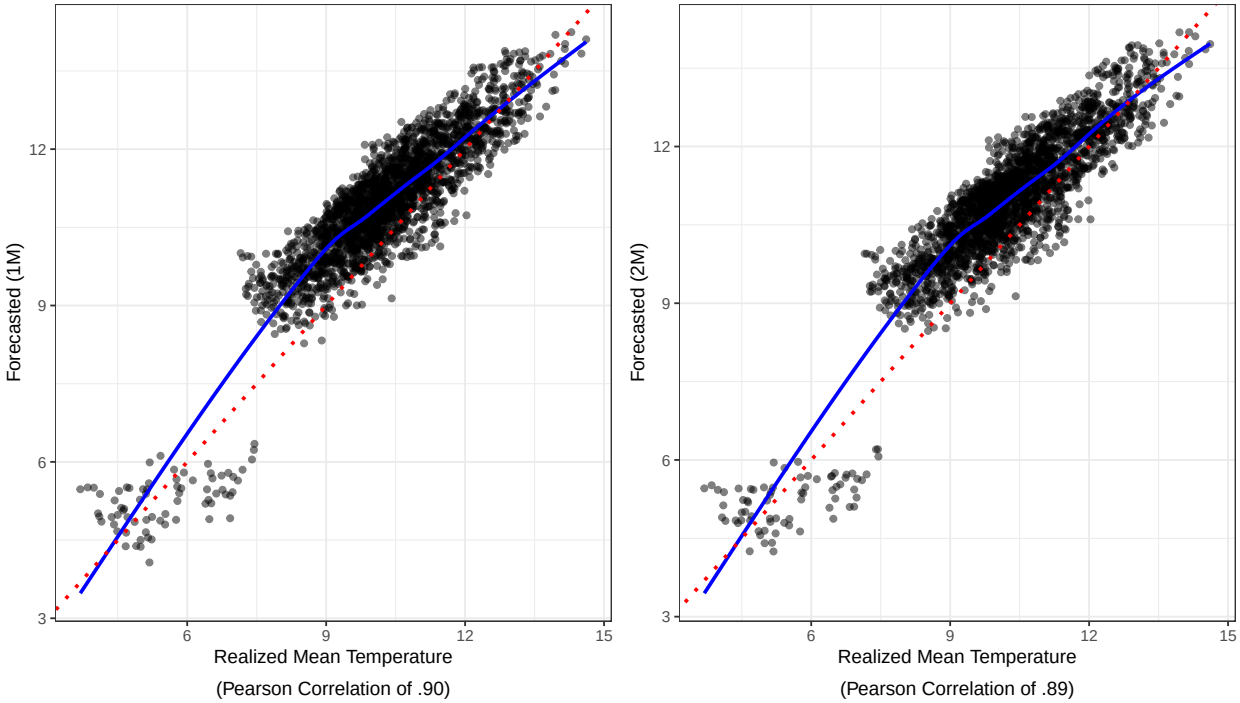


Figure A3: Correlation Realized and Forecasted Mean Temperature

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

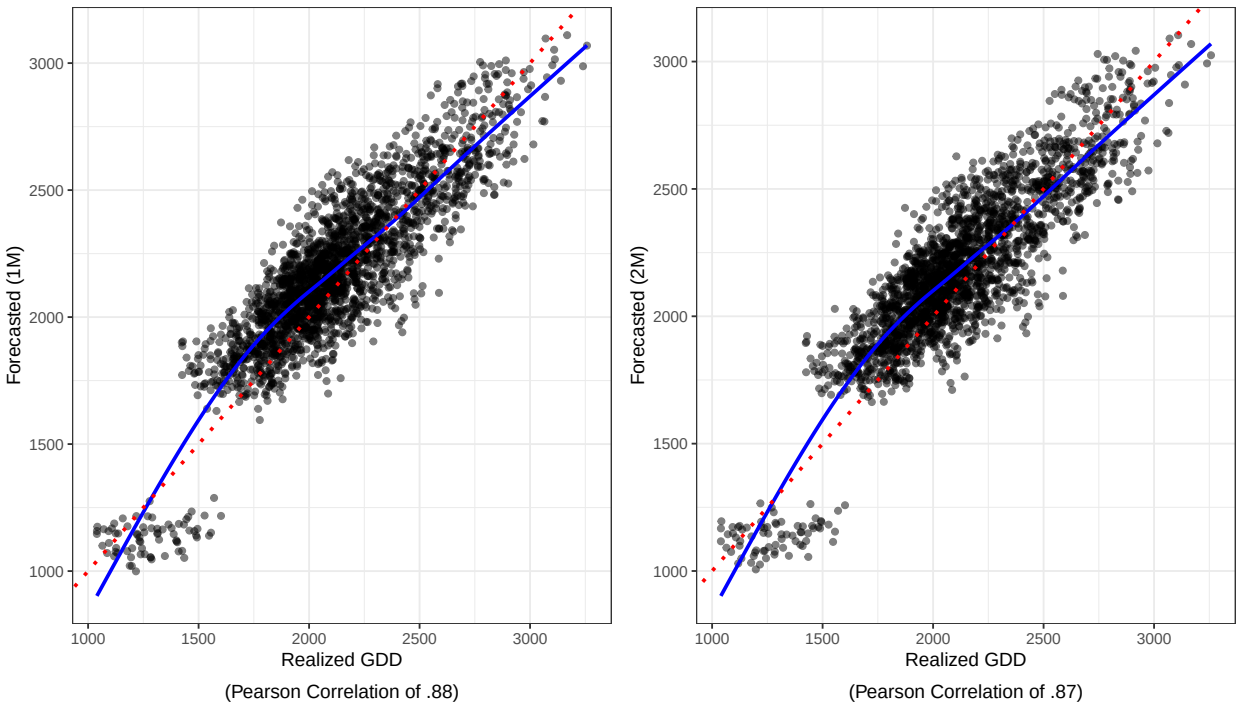


Figure A4: Correlation Realized and Forecasted GDDs

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

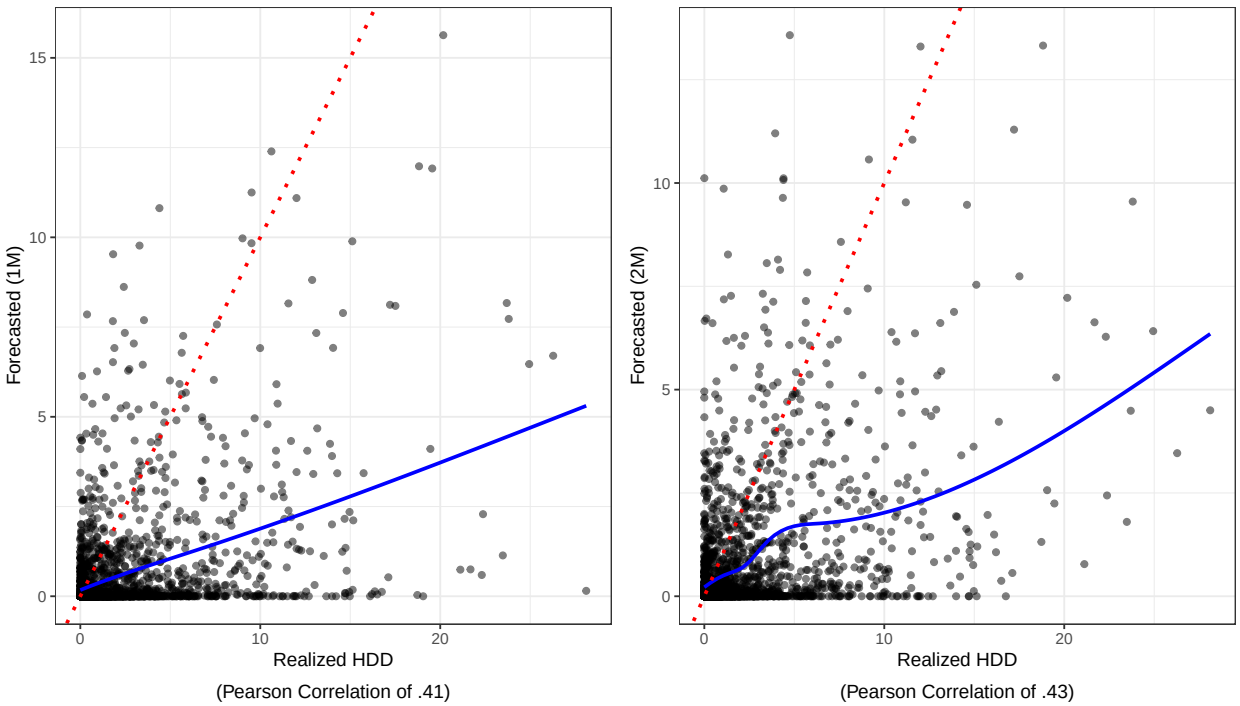


Figure A5: Correlation Realized and Forecasted HDDs

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

A.1.2 Climate in France

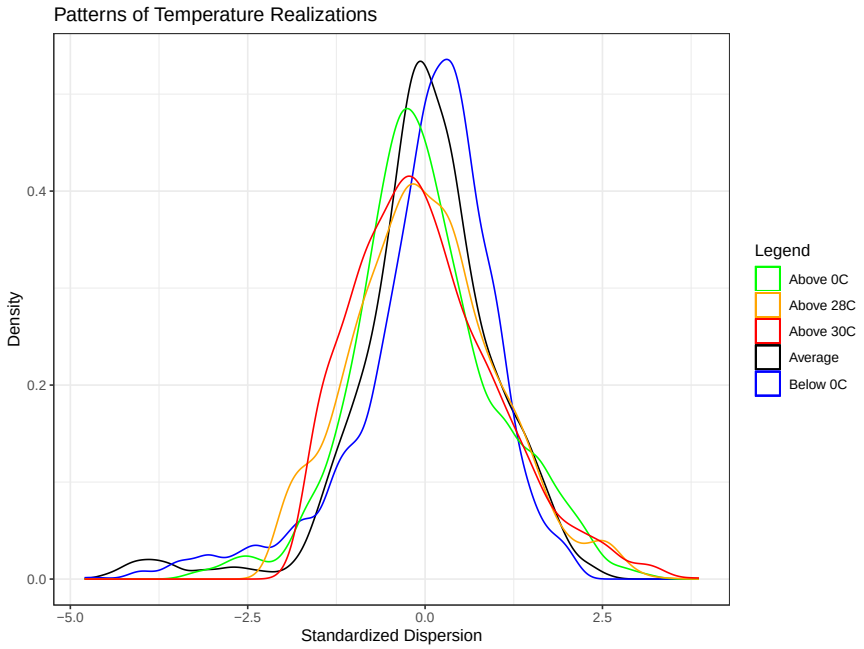


Figure A6: Distribution of the Standardized Conditional and Unconditional Mean Temperature

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

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

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

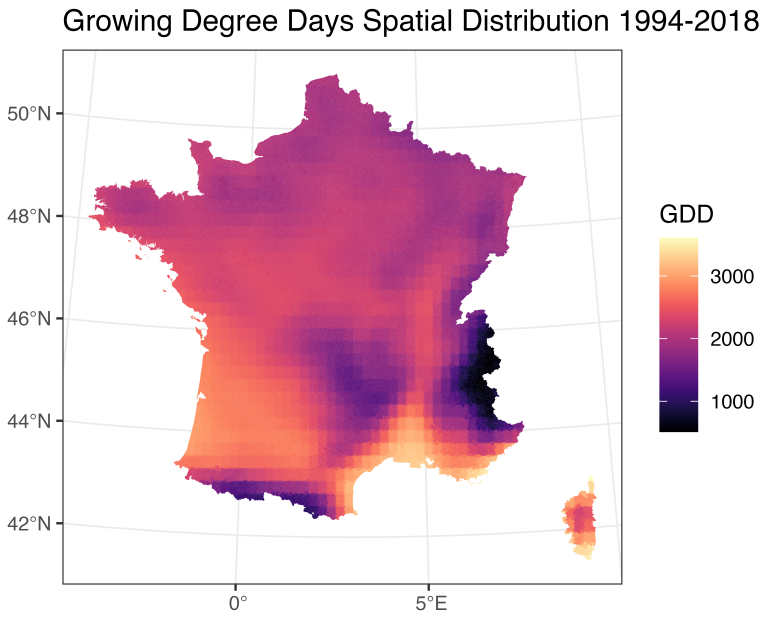


Figure A7: Distribution of Growing Degree Days

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

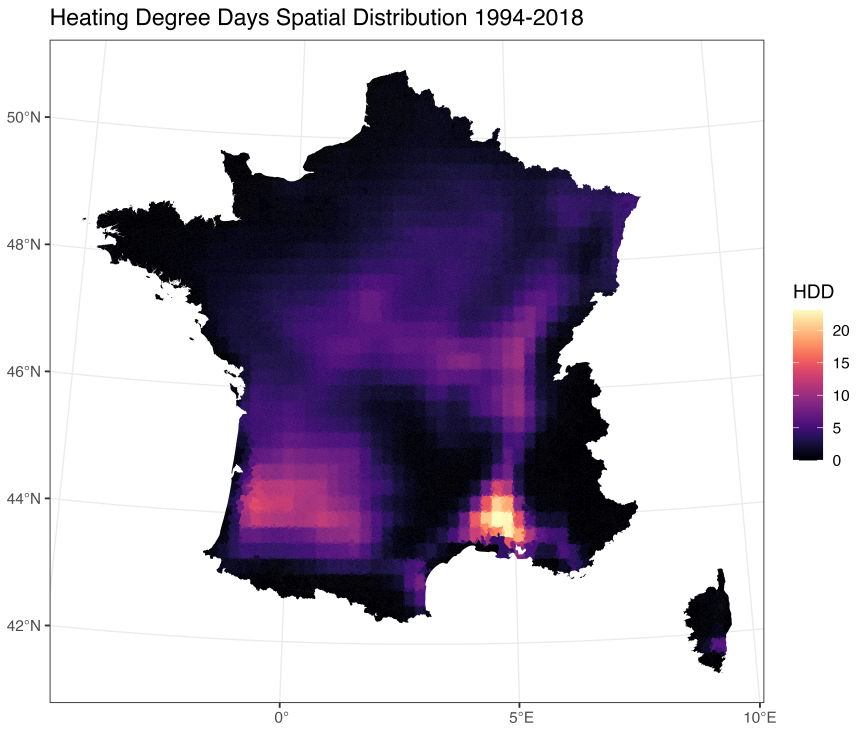


Figure A8: Distribution of Heating Degree Days

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

A.1.3 Growing Seasons

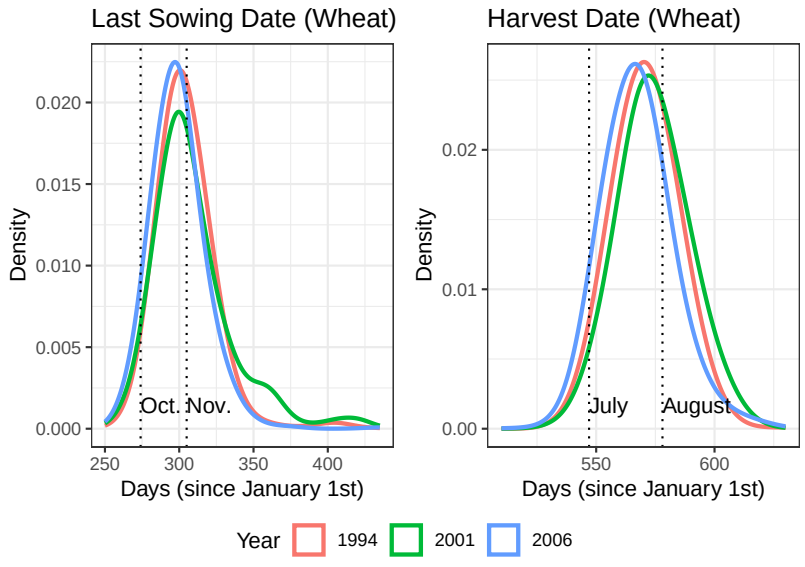


Figure A9: Wheat Growing Season in France

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

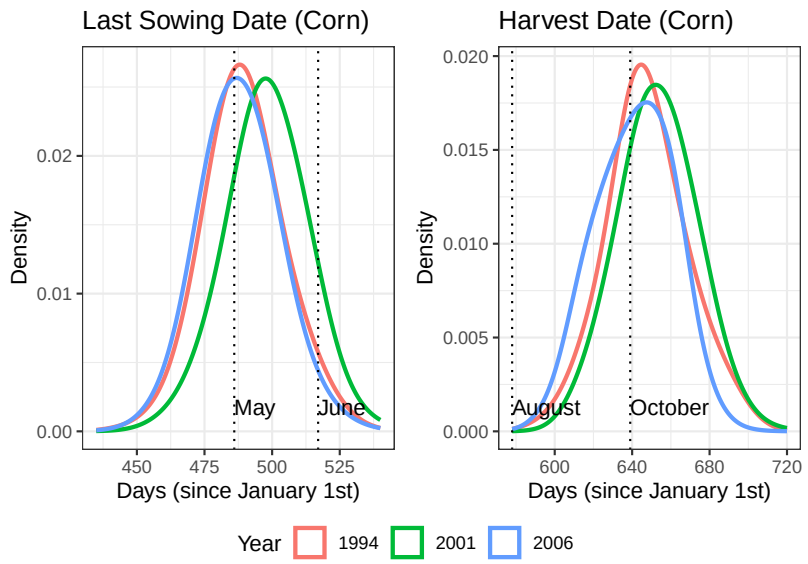


Figure A10: Corn Growing Season in France

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

A.1.4 Yield and Output Effects of Weather Variation

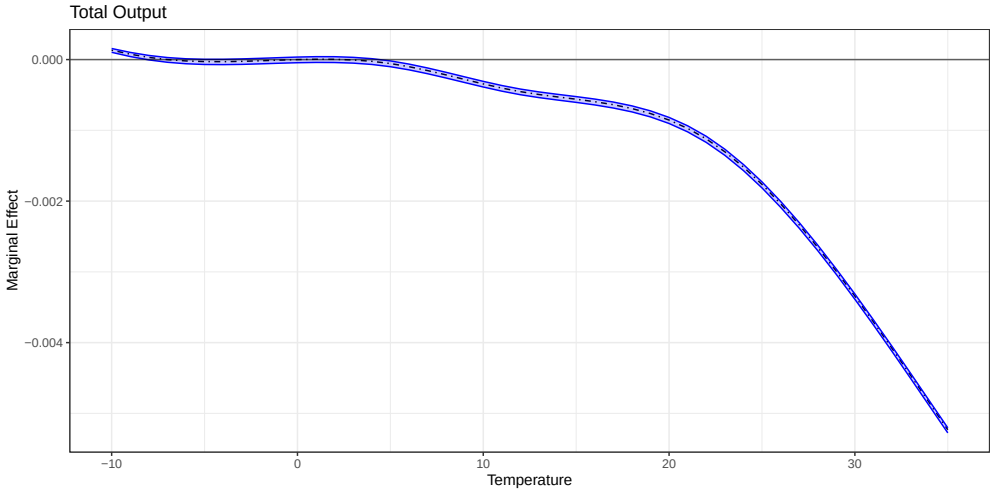


Figure A11: Temperature Effects on Farm Total Output

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

A.1.5 Bias of Climate Projections

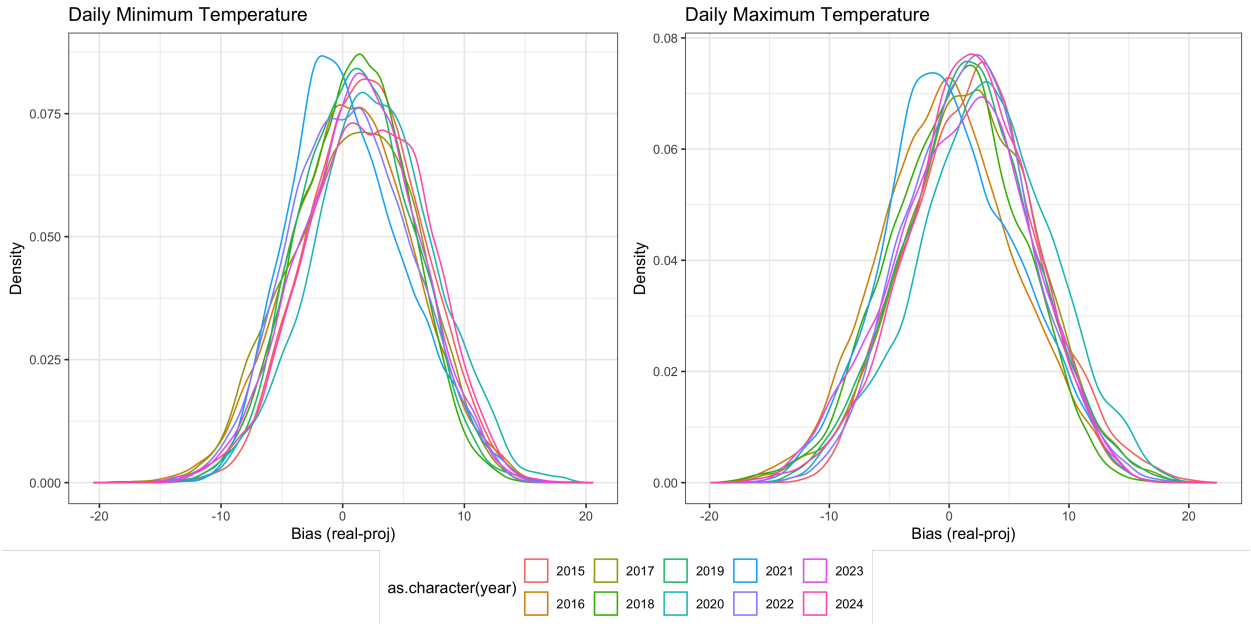


Figure A12: Distribution of Biases between CMIP6 and ERA5

Notes: This figure shows the distribution of the projection bias relative to observed temperature realizations. Observations are at the department-year-month-day level. This data is used to debias the climate temperature projections.

A.1.6 Multi-Product Farms

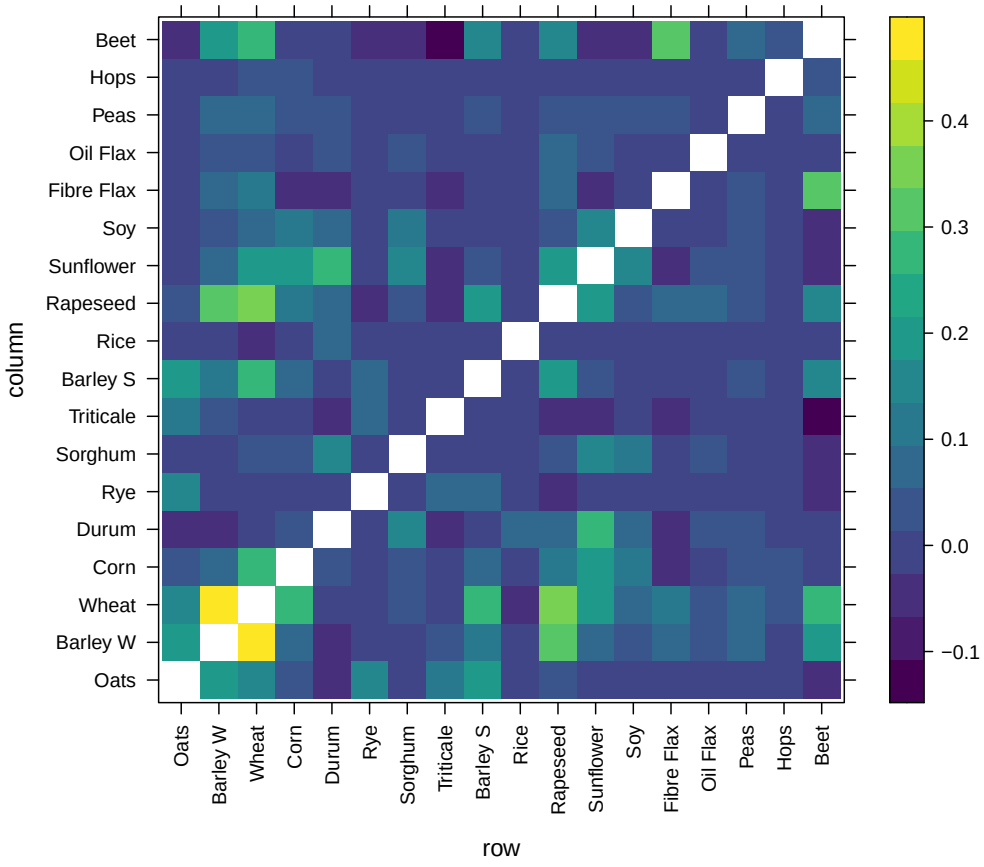


Figure A13: Conditional Probabilities to Grow Crop Pairs

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

Table A1: Descriptive Statistics - Farm-Level Crop Mix

	Cereals	Oil-Protein	Industrial
Cereals	1	0.33	0.09
Oil-Protein	0.99	1	0.13
Industrial	0.99	0.44	1

A.1.7 Input Price Indices

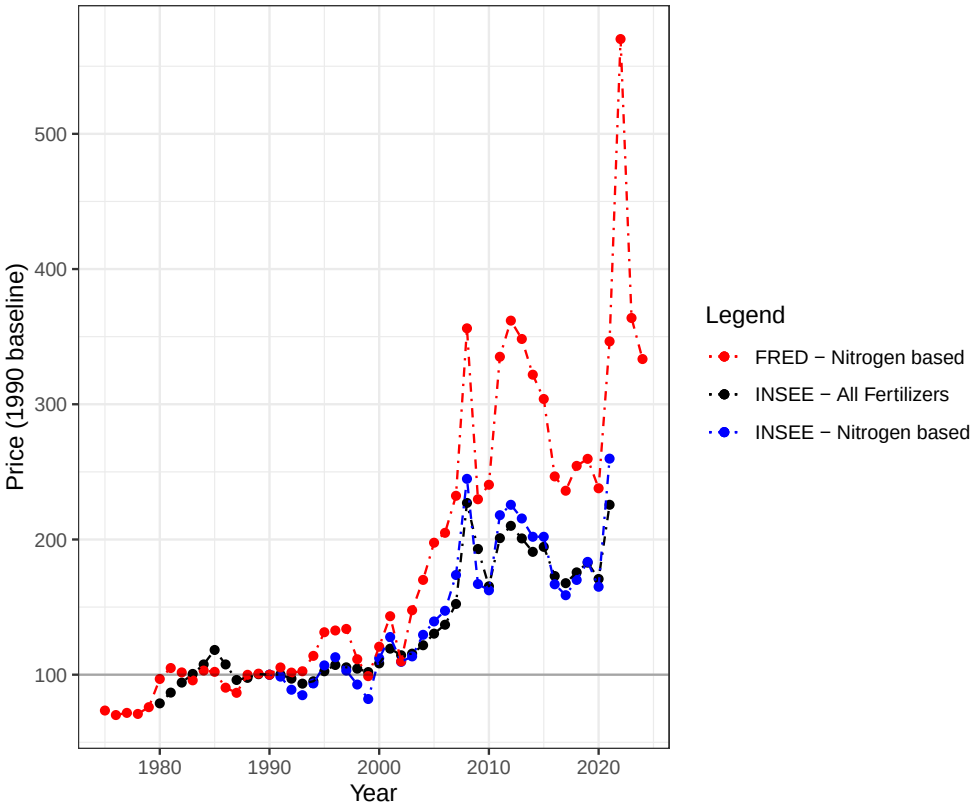


Figure A14: Comparison of Input Price Indices Series

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

A.1.8 Additional Results

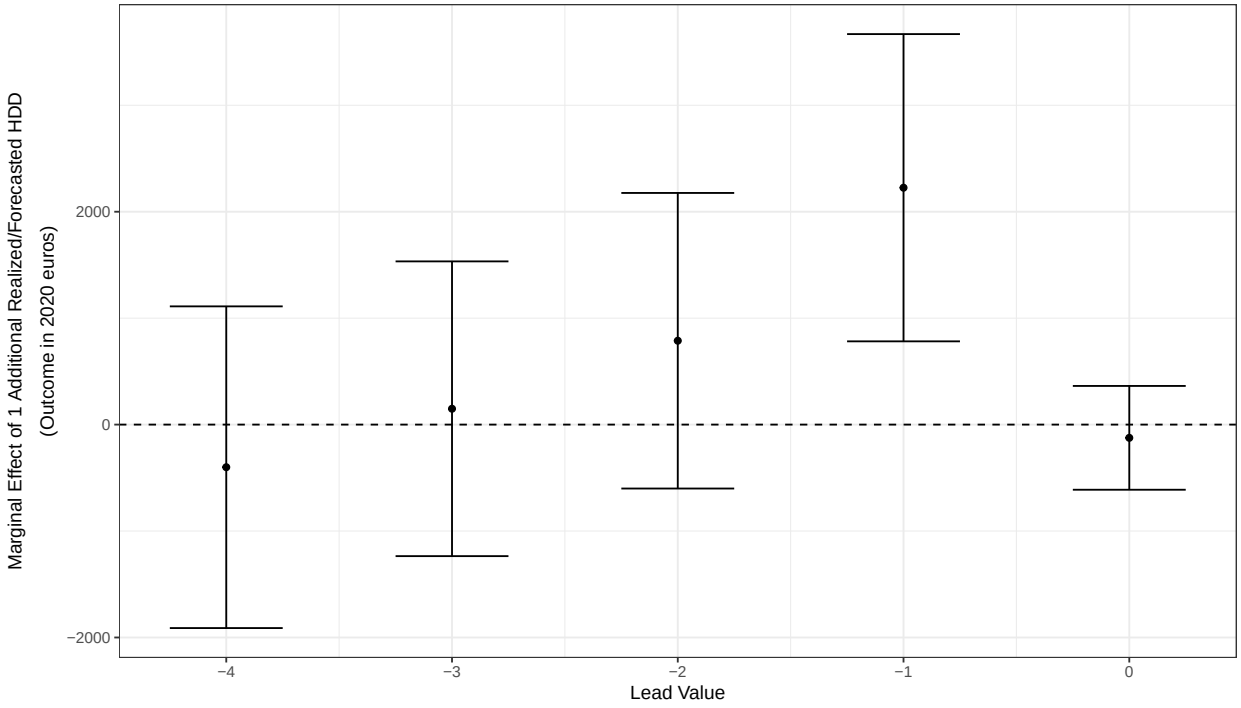


Figure A15: Profit Reactions

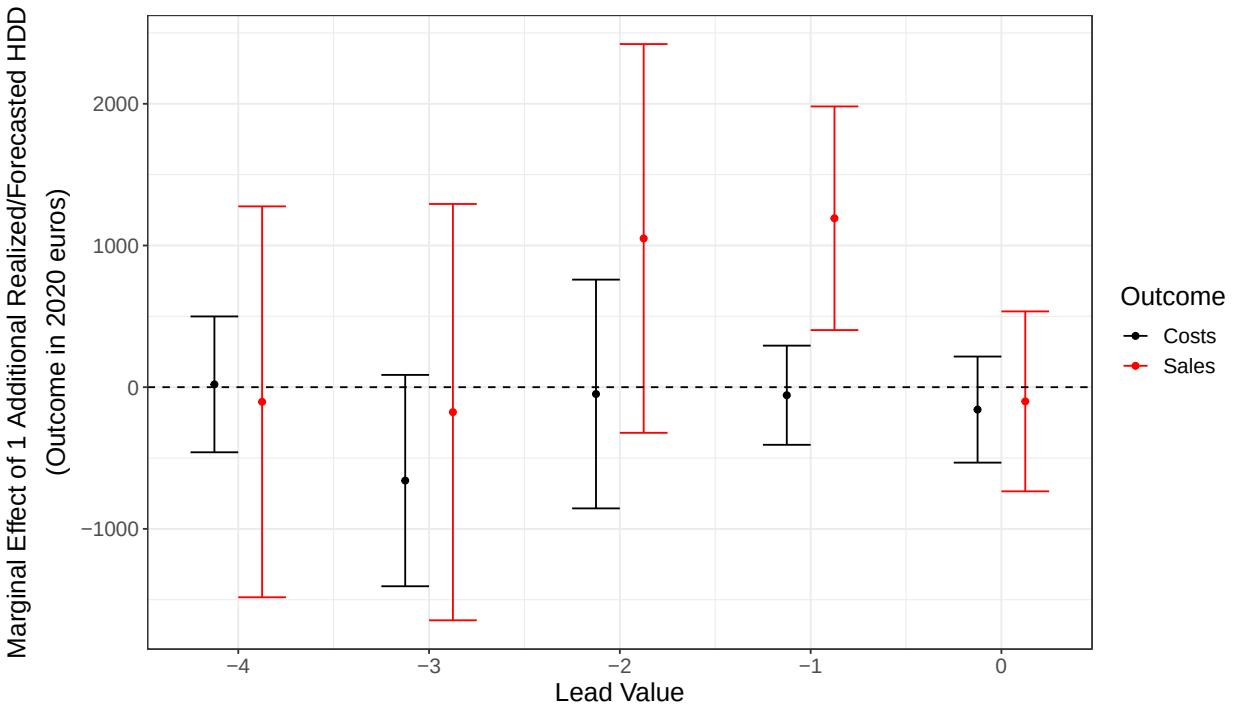


Figure A16: Effect of Forecated HDDs

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

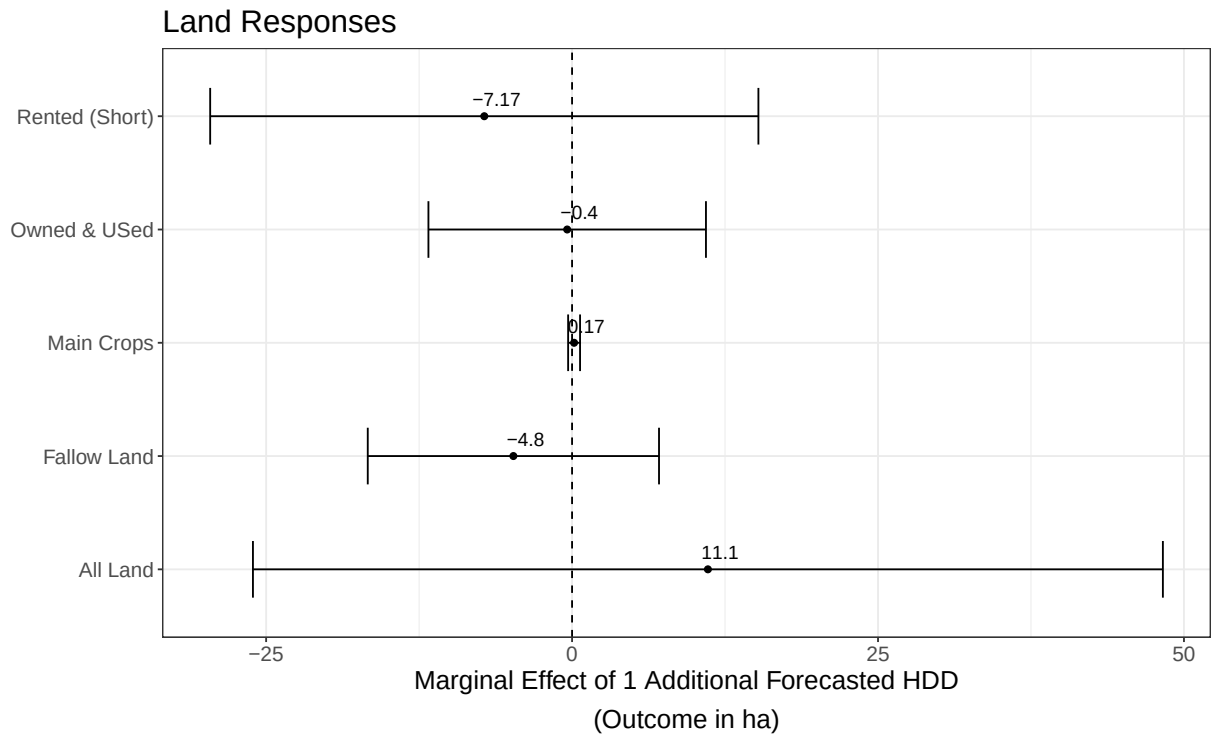


Figure A17: Decomposition of the Land responses

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

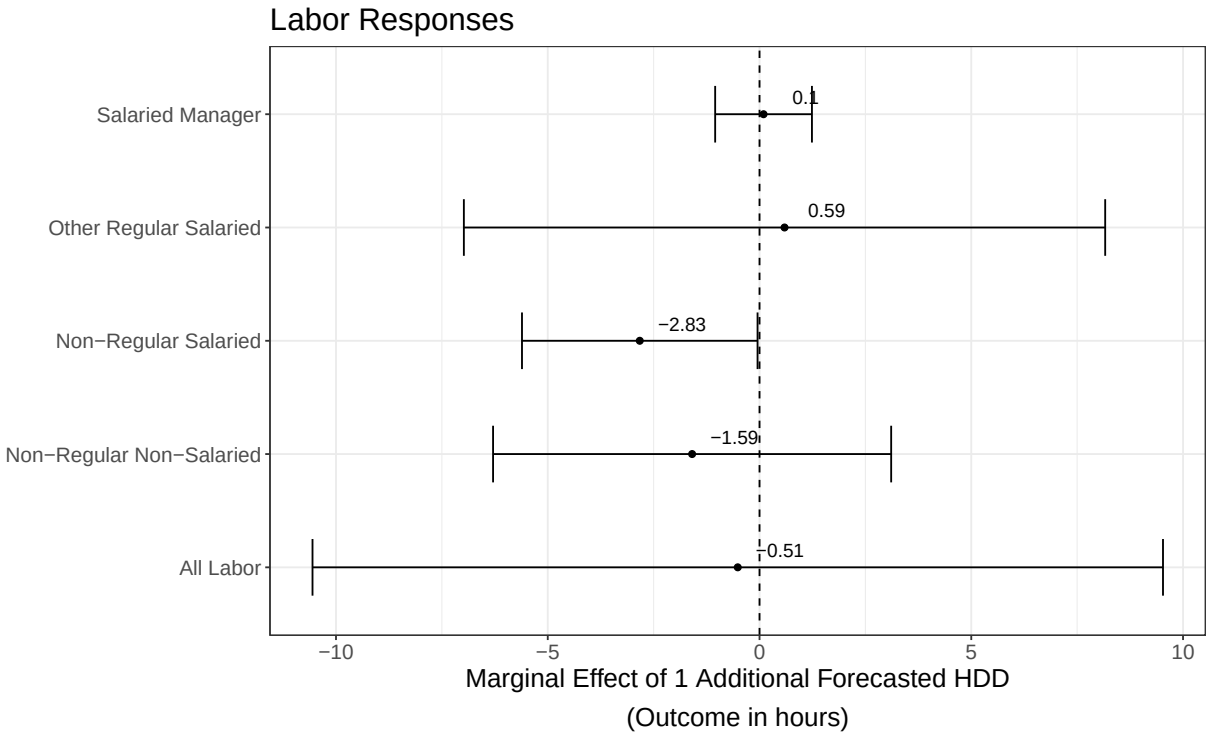


Figure A18: Decomposition of Labor Responses

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

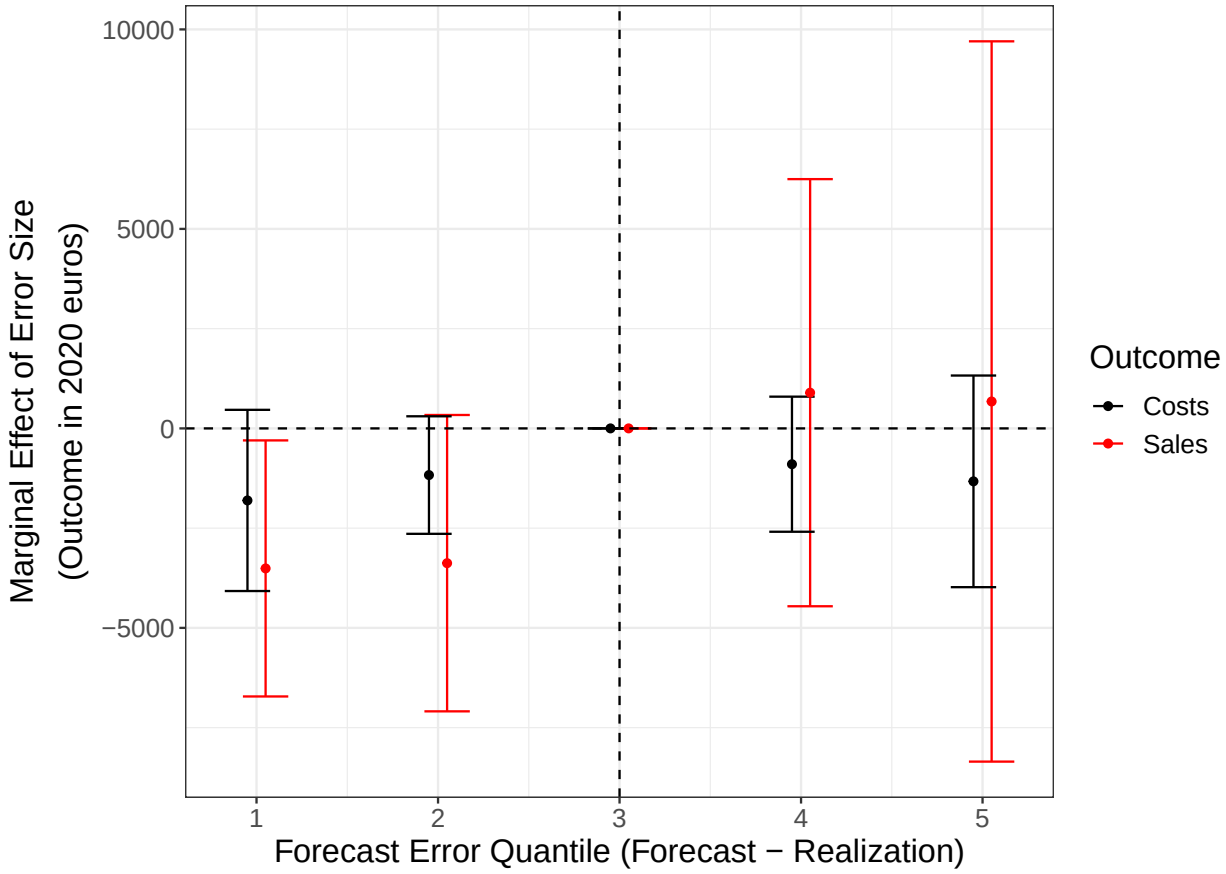


Figure A19: Cost and Revenue & Sign of the Forecast Error

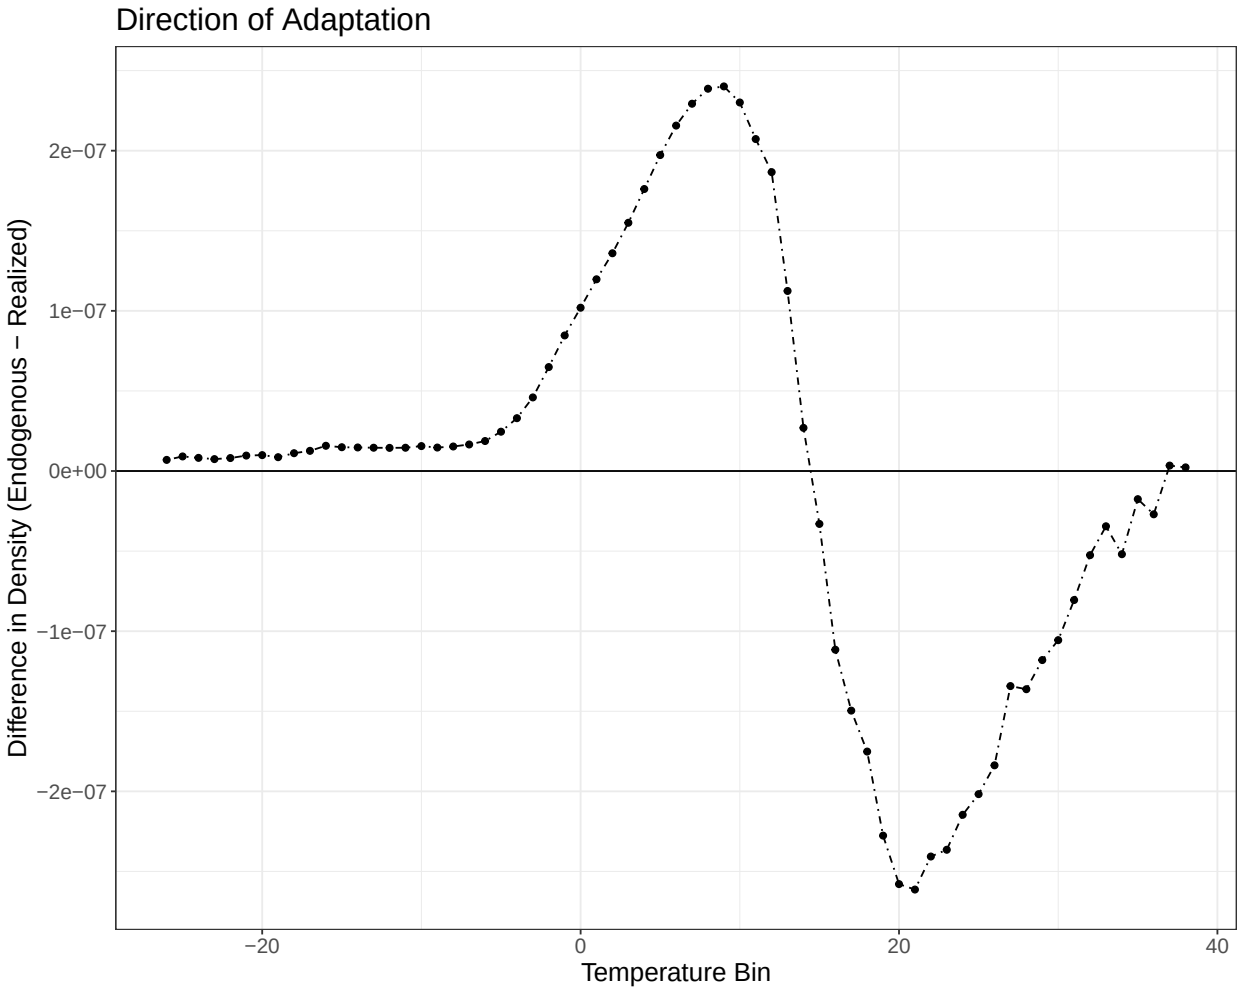


Figure A20: Difference in Temperature Distributions

A.2 Tables

Table A2: Descriptive Statistics - Farm-Level Dataset

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max
Rainfall	18,917	0.75	0.17	0.37	0.63	0.85	1.63
Rainfall (F1)	18,917	0.85	0.13	0.61	0.76	0.92	1.75
Rainfall (F2)	18,917	0.87	0.13	0.65	0.78	0.94	1.81
GDD	18,917	2,140	264	1,133	1,957	2,292	3,240
GDD (F1)	18,917	2,182	238	1,054	2,022	2,305	3,015
HDD	18,917	2.36	3.2	0	0.24	3.13	23.81
HDD (F1)	18,917	0.44	1.12	0	0	0.4	16
Sales	18,917	155,386	113,740	1,106	78,739	199,717	1,333,809
Total Costs	18,917	123,249	85,427	6,955	66,903	156,468	1,035,695
Intermediate Inputs	18,917	103,749	67,149	5,861	59,074	132,589	841,067
Value Added	18,917	58,357	62,858	−159,167	18,306	80,436	875,181
Profit	18,917	86,695	70,268	−135,057	40,066	114,771	973,561
Price Index	18,917	263	296.4	1.46	140.5	230.6	11,323
Storage (sum)	18,917	88.7	1,877	−28,569	−280	500	25,085
Storage (index)	16,075	65.4	983.7	−19,787	−67.9	176	25,085
Output (sum)	18,917	7,793	5,033	146	4,252	10,170	58,374
Output (index)	18,917	2,401	2,926	27.6	769	2,847	58,374
Production (corn)	10,346	3,464	4,046	0	877	4,682	58,374
Price (corn)	9,958	150	39.1	36	122.9	172.8	1,332
Sales (corn)	10,346	53,508	69,739	0	11,163	68,978	1,155,906
Quantity Sold (corn)	10,346	3,437	4,133	0	817	4,617	68,524
Production (wheat)	17,336	3,840	3,005	0	1,665	5,248	31,452
Price (wheat)	17,181	163.1	37.95	16.8	136.4	184	1,769
Sales (wheat)	17,336	61,538	53,986	0	24,345	82,922	660,068
Quantity Sold (wheat)	17,336	3,799	3,162	0	1,560	5,166	41,140
Irrigation	18,917	683.1	2,665	0	0	0	46,053
Labor	18,917	2,684	1,330	160	1,600	3,200	13,200
Phytosanitary	18,917	23,072	16,499	0	11,388	30,872	189,311
Fertilizer	18,917	31,777	20,727	0	17,456	41,061	217,676
Land	18,917	14,504	8,619	948	8,379	18,7	79,5
Seeds	18,917	12,231	9,686	0	5,847	15,882	146,059

Table A3: Descriptive Statistics - Store-Level Dataset

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max
Pesticide Prices	2,126	615.5	718.3	23.8	258	697	9,974
Fertilizer Prices	3,098	3,254	2,959	3.5	2,450	3,783	125,018
Seed Prices	1,831	343.7	1,777	28.5	95.7	171.8	48,892

Table A4: Descriptive Statistics - Plot-Level Data

Statistic	N	Mean	St. Dev.	Min	Max
Ploughing	32,263	337.4	83.5	15	570
Sowing	39,988	378.6	99.2	15	570
Irrigation	3,775	532.1	24.3	170	630
Harvest	38,818	594.3	41.5	510	720

Table A5: Descriptive Statistics - Land Price Data

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max
Land Prices	3,405	5,754	3,085	448	3,717	6,834	30,097

Table A6: Descriptive Statistics - Weather Outcomes within Samples

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max
Panel A	Entire Sample						
Rainfall	18,926	0.76	0.17	0.37	0.63	0.85	1.63
Rainfall (F1)	18,926	0.85	0.13	0.61	0.76	0.92	1.75
Rainfall (F2)	18,926	0.87	0.13	0.65	0.78	0.94	1.81
GDD	18,926	2,140	264	1,133	1,957	2,292	3,240
GDD (F1)	18,926	2,182	237.7	1,054	2,022	2,306	3,015
HDD	18,926	2.361	3.2	0	0.24	3.13	23.8
HDD (F1)	18,926	0.44	1.12	0	0	0.4	16
Panel B	Worst Region-Years						
Rainfall	1,425	0.68	0.14	0.37	0.56	0.75	1.15
Rainfall (F1)	1,425	0.88	0.11	0.61	0.79	0.95	1.2
Rainfall (F2)	1,425	0.9	0.12	0.69	0.81	0.97	1.19
GDD	1,425	2,418	255.7	1,928	2,196	2,587	3,139
GDD (F1)	1,425	2,356	276.1	1,866	2,136	2,548	3,004
HDD	1,425	10.4	3.51	6.13	7.35	12	23.8
HDD (F1)	1,425	1.41	2.07	0	0.02	2.1	16

Table A7: Cost and Revenue Reactions to HDD

Dependent Variables:	Sales	Costs	Sales	Costs
Model:	(1)	(2)	(3)	(4)
Variables				
GDD	5.1	1.21	10.2	2.95
	(11.8)	(7.81)	(10.9)	(8.43)
GDD (F)	37.4	30.8**		
	(34.7)	(14.5)		
GDD (F2)			55.5	38.4**
			(38.6)	(17.5)
HDD	-158	-192	-225	-197
	(332)	(157)	(312)	(194)
HDD (F)	1,192***	-56.7		
	(403)	(178)		
HDD (F2)			1,049	-48.2
			(699)	(412)
Mean	155,386	123,249	155,386	123,249
Unique Farms	2,603	2,603	2,603	2,603
Fixed-effects				
Farm	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Fit statistics				
Observations	18,917	18,917	18,917	18,917
R^2	0.89	0.94	0.89	0.94

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

Table A8: Profit and Value Added Reactions to HDD

Dependent Variables:	Value Added	Profit	Value Added	Profit
Model:	(1)	(2)	(3)	(4)
Variables				
GDD	2.59	-8.32	8.68	-3.42
	(12.7)	(13.2)	(13.4)	(13.3)
GDD (F1)	16.3	3.26		
	(43.1)	(43.1)		
GDD (F2)			64	42.8
			(46.1)	(46.4)
HDD	53.4	185	-13.7	120
	(292)	(277)	(261)	(240)
HDD (F1)	2,429***	2,067***		
	(809)	(708)		
HDD (F2)			1,023	1,076
			(721)	(665)
Mean	58,357	86,695	58,357	86,695
Unique Farms	2,603	2,603	2,603	2,603
Fixed-effects				
Farm	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Fit statistics				
Observations	18,917	18,917	18,917	18,917
R^2	0.81	0.84	0.8	0.84

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

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

Table A9: Price and Quantity Reactions to Forecasts

Dependent Variables:	Price	Output	Storage	Price	Output	Storage
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
GDD	-0.0275	0.0000	-0.0018**	-0.025	0.0000	-0.0017**
	(0.029)	(0.0001)	(0.0007)	(0.029)	(0.0001)	(0.0007)
GDD (F)	0.077	0.0002	-0.0002			
	(0.059)	(0.0004)	(0.0015)			
GDD (F2)				0.12**	0.0005	0.0004
				(0.053)	(0.0005)	(0.0016)
HDD	-0.23	-0.0045*	-0.0021	-0.14	-0.0049**	-0.0012
	(0.32)	(0.0022)	(0.015)	(0.36)	(0.0022)	(0.014)
HDD (F)	0.26	0.0089*	0.025			
	(1.56)	(0.0046)	(0.021)			
HDD (F2)				-1.74	0.0024	-0.0027
				(1.146)	(0.0056)	(0.027)
Mean	205	7,793	88.7	205	7,793.0	88.6
Unique Farms	2,603	2,603	2,183	2,603	2,603	2,183
Fixed-effects						
Farm	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
Observations	18,917	18,917	9,602	18,917	18,917	9,602
R^2	0.73	0.95	0.66	0.73	0.94670	0.66557

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

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

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

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

Table A11: Cost and Revenue Reactions to HDD—Comparison

Dependent Variables:	Revenue		Costs	
Model:	(1)	(2)	(3)	(4)
Variables				
GDD	5.1	9.88	1.21	2.48
	(11.8)	(11.7)	(7.81)	(8.44)
GDD (F)	37.4		30.8**	
	(34.7)		(14.5)	
HDD	-158	-100	-192	-158
	(332)	(324)	(157)	(191)
HDD (F)	1,192***		-56.7	
	(401)		(170)	
Mean	155,386	155,386	123,249	123,249
Unique Farms	2,603	2,603	2,603	2,603
Fixed-effects				
Farm	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Fit statistics				
Observations	18,917	18,917	18,917	18,917
R^2	0.89	0.89	0.94	0.94

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

Table A12: Profit and Value Added Reaction to HDD—Comparison

Dependent Variables:	Value Added		Profit	
Model:	(1)	(2)	(3)	(4)
Variables				
GDD	2.59	7.93	-8.32	-3.99
	(12.7)	(14.5)	(13.2)	(14.2)
GDD (F)	16.3		3.257	
	(43.1)		(43.1)	
HDD	53.4	118	185	232
	(292)	(273)	(277)	(253)
HDD (F)	2,429***		2,067***	
	(809)		(708)	
Mean	58,357	58,357	86,695	86,695
Unique Farms	2,603	2,603	2,603	2,603
Fixed-effects				
Farm	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Fit statistics				
Observations	18,917	18,917	18,917	18,917
R^2	0.81	0.8	0.84	0.84

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

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

Dependent Variables:	Land	Labor	Fertilizer	Phytosanitary	Seeds	Irrigation
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
GDD	-0.15	0.066	-2.65	-1.29	0.39	0.44
	(0.23)	(0.14)	(2.66)	(1.26)	(1.18)	(0.27)
GDD (F)	0.096	0.39**	7.64	-0.69	-0.25	-0.49
	(0.74)	(0.18)	(8.72)	(2.42)	(2.27)	(0.67)
HDD	5.88	-0.53	-86.7	-16.1	-17.9	7.51
	(6.06)	(4.22)	(63.4)	(36.7)	(25.7)	(6.73)
HDD (F)	11.1	-0.51	-9.14	54.6	-32.9	-23.9
	(18.9)	(5.12)	(77.5)	(71.5)	(52.6)	(17.4)
Mean	14,504	2,684	31,777	23,072	12,231	683
Unique Farms	2,603	2,603	2,603	2,603	2,603	2,603
Fixed-effects						
Farm	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
Observations	18,917	18,917	18,917	18,917	18,917	18,917
R^2	0.98	0.86	0.89	0.92	0.88	0.86

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

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

Dependent Variables:	Land	Labor	Fertilizer	Phytosanitary	Seeds	Irrigation
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
GDD	-0.09	0.096	-2.21	-1.47	0.31	0.32
	(0.24)	(0.14)	(2.64)	(1.27)	(1.15)	(0.27)
GDD (F2)	-0.48	0.692**	7.58	2.2	-0.71	0.14
	(1.27)	(0.33)	(9.08)	(3.12)	(2.81)	(0.69)
HDD	4.59	-1.46	-91.6	-7.06	-14.9	10.1
	(4.93)	(3.62)	(73.2)	(42.1)	(25.3)	(6.74)
HDD (F2)	40.9	4.45	-65.3	-10	-118**	-33.4
	(39.9)	(10.2)	(119)	(84.2)	(57.3)	(23.7)
Mean	14,504	2,684	31,777	23,072	12,231	683
Unique Farms	2,603	2,603	2,603	2,603	2,603	2,603
Fixed-effects						
Farm	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
Observations	18,917	18,917	18,917	18,917	18,917	18,917
R^2	0.98	0.86	0.89	0.92	0.88	0.86

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

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

Dependent Variables:	All Store Prices	Fertilizers	Pesticides	Seeds	Land
Model:	(1)	(2)	(3)	(4)	(5)
Variables					
GDD	-0.13	-0.12	-0.33	-0.0022	-0.92**
	(0.26)	(0.11)	(0.36)	(0.71)	(0.38)
GDD (F)	-0.51	0.29*	-1.02	-0.54	0.38
	(1.06)	(0.17)	(1.34)	(2.56)	(0.74)
HDD	-15.6	-0.9	-17.4	-15.2	5.28
	(10.3)	(1.97)	(14.1)	(26.5)	(7.41)
HDD (F)	-0.28	-2.03	3.34	3.64	-1.19
	(11.6)	(3.21)	(21.1)	(20.9)	(12.7)
Mean	1,899	517	3,151	311	4,848
Unique Farms	298	203	283	195	2,428
Fixed-effects					
Company	Yes	Yes	Yes	Yes	Year
Product	Yes	Yes	Yes	Yes	
Year	Yes	Yes	Yes	Yes	Yes
Farm					Yes
Fit statistics					
Observations	33,174	9,258	17,883	6,033	17,003
R^2	0.35	0.62	0.22	0.38	0.96

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

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

Dependent Variables:	All Store Prices	Fertilizers	Pesticides	Seeds	Land
Model:	(1)	(2)	(3)	(4)	(5)
Variables					
GDD	-0.14	-0.079	-0.37	-0.028	-0.8877**
	(0.22)	(0.12)	(0.38)	(0.6)	(0.37)
GDD (F2)	0.55	0.59**	1.18	-0.54	0.7
	(0.58)	(0.22)	(1)	(2.46)	(0.61)
HDD	-13.8	-1.01	-15	-14.4	6.4
	(8.47)	(2.2)	(14.5)	(28.4)	(7.86)
HDD (F2)	-13.9	-3.46	-18.3	-32.1	-6.15
	(19.3)	(4.28)	(34.8)	(32.3)	(14.4)
Mean	1,899	517	3,151	311	4,848
Unique Farms	298	203	283	195	2,428
Fixed-effects					
Company	Yes	Yes	Yes	Yes	
Product	Yes	Yes	Yes	Yes	
Year	Yes	Yes	Yes	Yes	Yes
Farm					Yes
Fit statistics					
Observations	33,174	9,258	17,883	6,033	17,003
R^2	0.35	0.62	0.22	0.38	0.96

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

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

Dependent Variables:	Wheat	Corn	Sunflower	Colza	Beets
Model:	(1)	(2)	(3)	(4)	(5)
Variables					
GDD	-0.53	0.39	0.13	-0.41**	1.84
	(0.64)	(0.36)	(0.1)	(0.19)	(2.87)
GDD (F)	0.49	0.95	-0.19	0.51	-6.31
	(1.09)	(1.26)	(0.24)	(0.56)	(4.83)
HDD	1.23	-27.7*	-2.15	9.31**	-24
	(8.99)	(14.6)	(2.61)	(3.57)	(58.4)
HDD (F)	61.7***	-6.59	-1.01	13.3	66.4**
	(19.6)	(9.32)	(5.21)	(12.4)	(29.8)
Mean	3,840	3,464	509	915	9,128
Unique Farms	2,448	1,607	1,357	1,940	346
Fixed-effects					
Farm	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes
Fit statistics					
Observations	17,336	10,346	7,203	12,394	2,267
R^2	0.9	0.93	0.79	0.81	0.93

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

Table A18: Farm-Level Reactions: Decomposed Weather

Dependent Variables:	Revenue	Costs	Revenue	Costs
Model:	(1)	(2)	(3)	(4)
Variables				
GDD (m)	10.2	14.7	9.52	11.2
	(11.4)	(11.2)	(13.5)	(10.8)
GDD (nm)	6.28	3.59	8.04	1.77
	(11.5)	(9.02)	(10.8)	(8.56)
GDD (m) (F)	21.8	24.5		
	(34.1)	(16.9)		
GDD (m) (F2)			76.4**	27.2
			(34.6)	(17.4)
GDD (nm) (F)	1.87	31.3*		
	(49.8)	(15.8)		
GDD (nm) (F2)			44.2	27.5*
			(39.1)	(15.6)
HDD (m)	-281	-226	-308	-198
	(326)	(145)	(279)	(132)
HDD (nm)	-484	168	-218	239
	(1,735)	(809)	(1,357)	(744)
HDD (m) (F)	-1,687**	-78.9		
	(735)	(934)		
HDD (nm) (F)	1,713***	-79.3		
	(438)	(363)		
HDD (m) (F2)			753	-72.9
			(759)	(385)
HDD (nm) (F2)			-641	-637
			(786)	(424)
Mean	155,396	123,265	155,396	123,265
Unique Farms	2,523	2,523	2,523	2,523
Fixed-effects				
Farm	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Fit statistics				
Observations	17,581	17,581	17,581	17,581
R^2	0.89	0.94	0.89	0.94

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

Table A19: Timing Response to Forecasts (Wheat - 1 month lead)

Dependent Variables:	Ploughing	First Irrigation	Sowing	Harvest
Model:	(1)	(2)	(3)	(4)
Variables				
GDD	0.02	0.079	0.021	4.7×10^{-5}
	(0.024)	(0.11)	(0.012)	(0.0028)
GDD (F)	0.07	1.57***	0.14	0.0094
	(0.078)	(0.13)	(0.049)	(0.018)
HDD	0.91	-3.85***	0.47	-0.14
	(0.38)	(0.28)	(0.31)	(0.047)
HDD (F)	-1.17*	2.12*	-0.91	0.36
	(0.28)	(0.66)	(0.52)	(0.21)
Mean	338	532	379	594
Fixed-effects				
Department	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Fit statistics				
Observations	7,313	214	9,987	9,974
R^2	0.37	0.54	0.24	0.53

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

Table A20: Cost and Revenue Reactions to Alternative Weather

Dependent Variables:	Revenue	Costs	Revenue	Costs
Model:	(1)	(2)	(3)	(4)
Variables				
GDD	2.73	-1.26	9.88	-0.29
	(10.6)	(8)	(9.13)	(8.97)
GDD (F)	21.3	35.9**		
	(37.1)	(14.8)		
GDD (F2)			54.9	37.3**
			(37.2)	(17.4)
HDD	-534	-106	-70.3	-85.4
	(178)	(90.8)	(150)	(108)
HDD (F)	628**	-13.6		
	(261)	(128)		
HDD (F2)			30.4	46.1
			(154)	(92.7)
FDD	22.7	-2.68	36.9	-3.24
	(30)	(12.4)	(29.8)	(12)
FDD (F)	-23.4	-96		
	(107)	(84.3)		
FDD (F2)			-7.97	-41.8
			(139)	(78.2)
Mean	155,638	123,304	155,638	123,304
Unique Farms	2,625	2,625	2,625	2,625
Fixed-effects				
Farm	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Fit statistics				
Observations	18,428	18,428	18,428	18,428
R^2	0.89	0.94	0.89	0.94

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

Table A21: Weather Decomposition

	Mean
GDD	2066
GDD marginal	1803
GDD non-marginal	263
GDD - ratio	0.13
HDD	2
HDD marginal	0.25
HDD non-marginal	1.78
HDD - ratio	0.85

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

Table A22: Persistence in Heating Degree Days Forecasts

Dependent Variables:	FHDD (lag 2)	FHDD (lag 1)	FHDD	FHDD (lead 1)	FHDD (lead 2)	FHDD (lead 3)
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
GDD	-0.0024	-0.0018	1.56×10^{-17}	-1.61×10^{-5}	0.0016	-0.0005
	(0.0018)	(0.0022)	(4.29×10^{-10})	(0.00)	(0.0025)	(0.0025)
GDD (F)	-0.0007	-0.0019	2.26×10^{-17}	0.0008	0.0058	0.002
	(0.0045)	(0.0035)	(1.17×10^{-9})	(0.0052)	(0.0067)	(0.0035)
GDD (lag)	-0.0022	0.0009	-3.73×10^{-17}	0.0001	0.0001	-0.0002
	(0.0014)	(0.0016)	(2.1×10^{-10})	(0.0013)	(0.0012)	(0.0015)
GDD (lag 2)	0.0008	-0.0001	-4.94×10^{-17}	0.0004	-0.0006	0.0017
	(0.0015)	(0.0009)	(6.3×10^{-11})	(0.0011)	(0.0010)	(0.0021)
HDD	-0.029	-0.0033	-1.27×10^{-15}	0.049	0.0014	0.053
	(0.053)	(0.026)	(1.5×10^{-8})	(0.068)	(0.046)	(0.09)
HDD (F)	-0.11	-0.11	1^{***}	-0.14	-0.19	-0.23
	(0.14)	(0.13)	(1.75×10^{-8})	(0.14)	(0.11)	(0.33)
HDD (lag)	-0.0014	0.013	-9.05×10^{-16}	0.022	0.03	0.0085
	(0.038)	(0.049)	(1.71×10^{-8})	(0.037)	(0.084)	(0.045)
HDD (lag 2)	-0.006	0.027	5.2×10^{-16}	0.022	0.003	-0.033
	(0.052)	(0.032)	(1.87×10^{-8})	(0.086)	(0.042)	(0.056)
Fixed-effects						
Farm	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
Observations	14,097	14,097	14,097	12,193	10,504	9,019
R^2	0.69	0.73	1	0.71	0.69	0.70

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

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

Table A23: Dynamic Effects on Profit

Dependent Variables:	Profit (lag 2)	Profit (lag)	Profit	Profit (lead)	Profit (lead 2)	Profit (lead 3)
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
GDD	-37.5	3.04	0.86	1.67	6.46	-12.1
	(26.4)	(18.6)	(15.6)	(21.2)	(28.9)	(15.2)
GDD (F)	-28.5	19.2	22.6	35.8	-23	74.2
	(44.2)	(41.4)	(51.6)	(31.8)	(65.7)	(68.7)
GDD (lag)	-20.2	23.1	-8.89	-11.2	24.5	-19.4
	(13.5)	(20)	(10.8)	(12.8)	(15.5)	(23.4)
GDD (lag 2)	33.1	-1.88	-19.3*	17.5	7.01	-4.01
	(20)	(10.6)	(10.9)	(11.4)	(22.9)	(23.4)
HDD	85.8	-646	110	344	295	-271
	(351)	(399)	(339)	(272)	(539)	(537)
HDD (F)	663	265	1,908**	-1,828***	-1,435*	-324
	(450)	(649)	(847)	(645)	(702)	(699)
HDD (lag)	-628	-383	523**	747**	-301	1,425*
	(388)	(500)	(234)	(313)	(533)	(703)
HDD (lag 2)	-810	546**	709**	-239	1,350*	196
	(528)	(225)	(255)	(408)	(774)	(732)
Mean	89,684	87,948	86,695	86,836	86,756	86,848
Unique Farms	1,904	1,904	1,904	1,689	1,485	1,308
Fixed-effects						
Farm	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
Observations	14,097	14,097	14,097	12,193	10,504	9,019
R^2	0.84	0.84	0.84	0.84	0.83	0.83

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

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

Table A24: Crop-Specific Land Movement

Dependent Variable:	Land	
Model:	(1)	(2)
Variables		
HDD (corn)	0.0017	
	(0.093)	
HDD F (corn)	-0.14	
	(0.17)	
HDD (wheat)	-0.051	
	(0.087)	
HDD F (wheat)	0.27	
	(0.32)	
HDD (colza)		0.07
		(0.075)
HDD F (colza)		0.24
		(0.19)
HDD (sunflower)		-0.016
		(0.067)
HDD F (sunflower)		-0.34**
		(0.15)
Unique Farms	2,604	2,604
Wald	1.13	3.89
Fixed-effects		
Farm	Yes	
Year		Yes
Fit statistics		
Observations	37,834	37,834
R^2	0.94	0.84

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

Notes. Two-way department-by-year standard-errors in parentheses. Regressions are run using system OLS, in order to compute the coefficient equivalence Wald tests. Stars indicate estimate is significantly different from zero: $*p < .10$, $**p < .05$, $***p < .01$. Realized and forecasted rainfall in levels and squares are included as controls, as well as quadratic region-specific time trends. We also include GDD realization and forecasts, but only show the HDD specific coefficients as they are the focus of the tests.

Table A25: Dynamic Effects on Costs

Dependent Variables:	Costs (lag 2)	Costs (lag)	Costs	Costs (lead)	Costs (lead 2)	Costs (lead 3)
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
GDD	-3.58	-6.16	-1.82	9.01	-1.59	-0.63
	(2.94)	(7.39)	(6.77)	(7.43)	(6.44)	(15.3)
GDD (F)	4.23	10.8	25.2	1.03	19.5	0.15
	(12.4)	(15.8)	(16.9)	(15.5)	(13.9)	(13.4)
GDD (lag)	-2.73	-1.79	8.35**	-1.98	-7.64	7.86
	(5.28)	(5.15)	(3.89)	(3.21)	(5.51)	(8.93)
GDD (lag 2)	-7.52	1.59	-1.1	-6.29*	3.88	-1.25
	(5.38)	(3.26)	(3.04)	(3.29)	(7.43)	(10.7)
HDD	156	85	-284	-213	-20	164
	(103)	(93.5)	(176)	(223)	(178)	(228)
HDD (F)	-173	-189	-197	1,028**	326	118
	(304)	(215)	(116)	(429)	(409)	(223)
	(127)	(116)	(118)	(145)	(301)	(186)
HDD (lag)	1.88	-243	-88.2	124	218	-152
	(114)	(157)	(175)	(144)	(222)	(314)
HDD (lag 2)	-177	-61.7	56	172	-117	212
Unique Farms	1,904	1,904	1,904	1,689	1,485	1,308
Fixed-effects						
Farm	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
Observations	14,097	14,097	14,097	12,193	10,504	9,019
R^2	0.94	0.94	0.94	0.94	0.94	0.95

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

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

Table A26: Heterogeneity of Farm-Level Results

Dependent Variables:	Profit	Revenue	Costs
Model:	(1)	(2)	(3)
Variables			
GDD	-64.7	-46.6	-22.2*
	(55.9)	(54.5)	(12.5)
GDD (F)	146	39.2	-53.8
	(87.1)	(82.1)	(84.7)
HDD	-1,233	-2,694	-1,106
	(1,311)	(1,998)	(677)
HDD (F)	13,143***	16,512**	4,357
	(4,326)	(6,247)	(3,303)
GDD × Mean Temperature	6.04	5.39	2.16***
	(5.09)	(4.97)	(0.76)
GDD (F) × Mean Temperature	-12.8**	-1.33	6.49
	(6.16)	(6.94)	(7.31)
HDD × Mean Temperature	124	224	80.2
	(130)	(183)	(62.5)
HDD (F) × Mean Temperature	-968**	-1,350**	-394
	(379)	(536)	(289)
Mean	86,695	155,386	123,249
Unique Farms	2,603	2,603	2,603
Fixed-effects			
Farm	Yes	Yes	Yes
Year	Yes	Yes	Yes
Fit statistics			
Observations	18,917	18,917	18,917
R^2	0.84	0.89	0.94

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

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

B Data Details

B.1 Definition of Variables

Definitions for the main variables used:

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

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

  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:

$$q_{jt} = \sum_{c \in \mathbb{C}_{jt}} \frac{l_{jct}}{\sum_{c \in \mathbb{C}_{jt}} l_{jct}} q_{jct} \tag{7}$$

- **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. By definition, this does not account for non-wage work, which are not collected.
- **Intermediary Inputs.** Defined as a deflated bill. The sum of expenses for: fertilizer, seeds, pesticides, animal food, veterinary products, products for animal reproduction, packaging, fuel, maintenance products, supplies, food for workers, raw materials, purchases of services for cultivation, breeding or others, water, gas, electricity, irrigation water, lease installments, material rental, animal rental, maintenance for buildings, lands and material, studies and research, veterinary services, communication and commercials, transportation costs, travel costs, postal services, banking services, other services and costs.
- **Total Costs.** Defined as a deflated bill. Corresponds to the sum of intermediary inputs, social contributions for workers, personnel expenses, taxes, insurance.
- **Value Added.** Defined as the difference between total production value minus animal purchases, minus intermediary inputs. Total production itself corresponds to total production sold, self-consumed

production, immobilised production, stored production, gains from animal boarding, land rentals, other rentals, agricultural tourism.

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

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

B.2 Weather Data

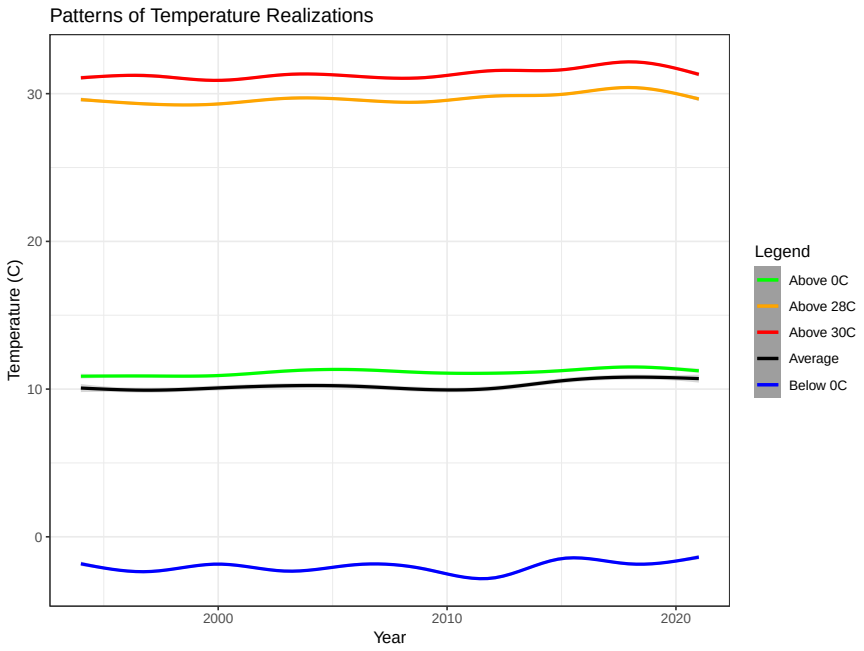


Figure A21: Unconditional and Conditional Average Temperature Realizations

C Validation of Results

Table A27: Corn-Specific Outcomes

Dependent Variables:	Output	Price	Sales	Sales Quantity
Model:	(1)	(2)	(3)	(4)
Variables				
GDD	0.47	-0.0041	8.82	0.66
	(0.37)	(0.0094)	(8.33)	(0.52)
HDD	-30.7*	0.73**	-134	-13.7
	(15.2)	(0.29)	(186)	(11.4)
Mean	3,464	150	53,508	3,437
Unique Farms	1,607	1,581	1,607	1,607
Fixed-effects				
Farm	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Fit statistics				
Observations	10,346	9,958	10,346	10,346
R^2	0.94	0.73	0.88	0.88

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

Table A28: Farm-Level Outcomes

Dependent Variables:	Revenue	Costs	Intermediary Inputs	Value Added	Profit	Output Price Index
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
GDD	9.88	2.48	-3.79	7.93	-3.99	-0.16
	(11.7)	(8.44)	(5.61)	(14.5)	(14.2)	(0.12)
HDD	-100	-158	-127	118	232	0.28
	(324)	(191)	(161)	(273)	(253)	(1.72)
Mean	155,386	123,249	103,749	58,357	86,695	263
Unique Farms	2,603	2,603	2,603	2,603	2,603	2,603
Fixed-effects						
Farm	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
Observations	18,917	18,917	18,917	18,917	18,917	18,917
R^2	0.89	0.94	0.93	0.80	0.84	0.70

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

D Robustness

D.1 Heterogeneous Response to Forecasts

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

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

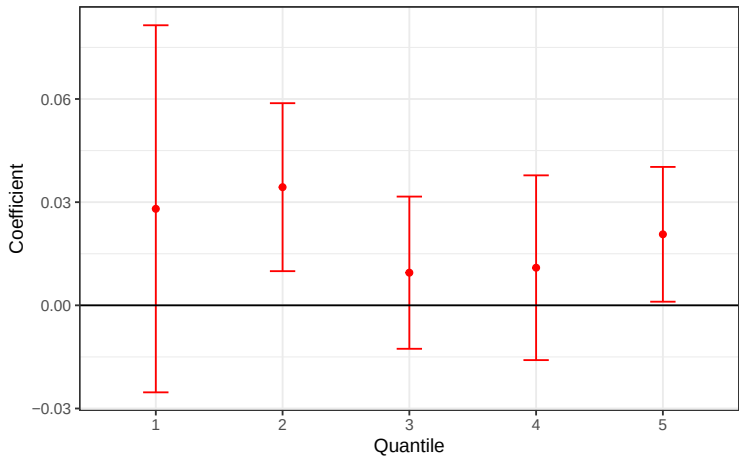


Figure A22: Varying Response to Fixed Lead of HDD

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.

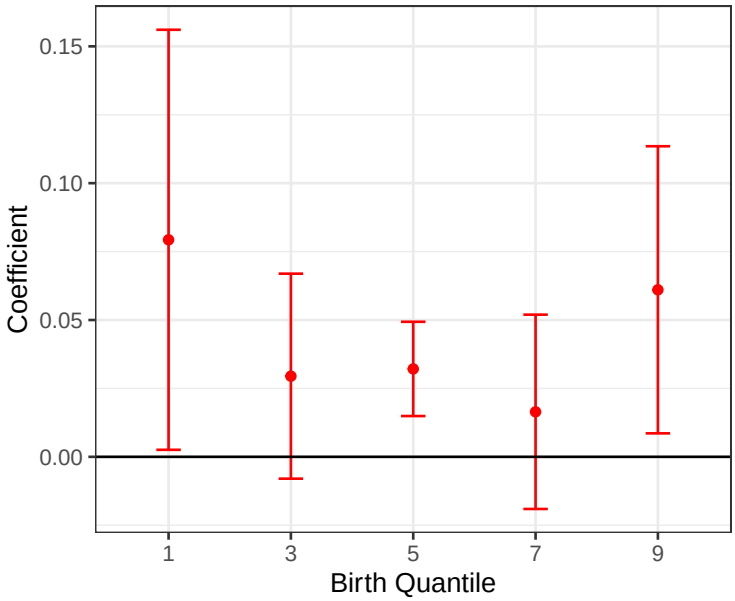


Figure A23: Profit Response to Forecasts across Manager's Age

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

D.2 Including Lags

As a robustness check, we run the same regression as before but including lagged values of weather realization (rainfall in levels and squared, GDD and HDD). Again, the regressions also include farm and year fixed effects, and we cluster standard errors at the department level. Lagged realizations might matter, first because of auto-correlation in weather, for example due to the role of large weather patterns such as the North Atlantic Oscillation. Second, lags might also play a role in the setting of farmers' beliefs about the upcoming weather, as modeled 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 A29: Farm-Level Profit - Lead 1 with Lags

Dependent Variables:	Value Added	Profit	Value Added	Profit
Model:	(1)	(2)	(3)	(4)
Variables				
GDD	7.83	-2.9	16.7	3.51
	(14.8)	(14.4)	(16.1)	(15.4)
GDD (lag)	-3.25	-7.94	-5.78	-9.42
	(9.94)	(9.86)	(9.75)	(9.9)
GDD (F)	24.8	9.61		
	(47.1)	(45.9)		
GDD (F2)			74.1	48.5
			(48.8)	(49.7)
HDD	-22.8	134	-111	61.4
	(333)	(331)	(290)	(280)
HDD (lag)	353*	429*	390*	448*
	(201)	(219)	(221)	(229)
HDD (F)	2,202**	1,781**		
	(849)	(762)		
HDD (F2)			1,058	1,098
			(720)	(644)
Mean	58,357	86,695	58,357	86,695
Unique Farms	2,217	2,217	2,217	2,217
Fixed-effects				
Farm	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Fit statistics				
Observations	16,314	16,314	16,314	16,314
R^2	0.81	0.84	0.81	0.84

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

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

D.3 Alternative Weather Aggregation

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

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

Table A30: Farm-Level Profit - Alternative Weather

Dependent Variables:	Value Added	Profit	Value Added	Profit
Model:	(1)	(2)	(3)	(4)
Variables				
GDD	6.97	-4.1	16.2	2.99
	(11.2)	(11.3)	(13.6)	(13.1)
GDD (F)	-12.4	-20.3		
	(44.7)	(44.6)		
GDD (F2)			54.2	33.2
			(42.6)	(42.9)
HDD	61.3	129	-1.85	70.3
	(177)	(171)	(155)	(143)
HDD (F)	943***	797***		
	(296)	(269)		
HDD (F2)			152	125
			(248)	(238)
FDD	68.2*	64.4	81.9**	76.2*
	(39.1)	(37.7)	(39)	(37.4)
FDD (F)	189	142		
	(158)	(156)		
FDD (F2)			139	176
			(124)	(134)
Mean	58,332	87,330	58,332	87,330
Unique Farms	2,625	2,625	2,625	2,625
Fixed-effects				
Farm	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Fit statistics				
Observations	18,428	18,428	18,428	18,428
R^2	0.81	0.84	0.81	0.84

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

The results are similar to the ones from our main specification. Forecasted HDDs have a positive impact of profit, albeit a smaller one than previously. A one unit increase in forecasted HDD corresponds

here to a smaller increase in temperature over the growing season, and it should be expected that it implies a smaller response. Forecasted freezing days also imply a positive profit response. As such, it seems that farmers adapt both to extremely hot and cold events, in a way that leads to non-marginal changes in their optimal profit.

In Table A20 we show the costs and revenue responses to these alternative weather calculations.

D.4 Removing Time Trends

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

Table A31: Farm-Level Profit - No Time Trend

Dependent Variables:	Sales	Costs	Sales	Costs
Model:	(1)	(2)	(3)	(4)
Variables				
GDD	3.47	-1.27	8.79	0.87
	(10.7)	(7.36)	(10.8)	(7.93)
GDD (F)	47.8	37.7**		
	(34.4)	(14.4)		
GDD (F2)			68.2*	45.5**
			(37.1)	(17.9)
HDD	-212	-184	-263	-183
	(303)	(149)	(309)	(187)
HDD (F)	1,342***	39.3		
	(455)	(270)		
HDD (F2)			1,382*	156
			(703)	(437)
Mean	155,386	123,249	155,386	123,249
Unique Farms	2,603	2,603	2,603	2,603
Fixed-effects				
Farm	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Fit statistics				
Observations	18,917	18,917	18,917	18,917
R^2	0.89	0.94	0.89	0.94

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

D.5 Disaggregated Realized Weather

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

Table A32: Cost and Revenue Reactions to Disaggregated Realized Weather

Dependent Variables:	Revenue	Costs	Revenue	Costs
Model:	(1)	(2)	(3)	(4)
Variables				
GDD	1.51	-2.43	1.51	-2.43
	(16.7)	(5.36)	(16.7)	(5.36)
GDD (F)	42.9	52.7**	42.9	52.7**
	(45.7)	(22.5)	(45.7)	(22.5)
HDD	43.1	39.9	43.1	39.9
	(328)	(114)	(328)	(114)
HDD (F)	1,289***	52.9	1,289***	52.9
	(256)	(37.9)	(256)	(37.9)
Mean	164,532	131,969	164,532	131,969
Unique Farms	1,890	1,890	1,890	1,890
Fixed-effects				
Farm	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Fit statistics				
Observations	12,754	12,754	12,754	12,754
R^2	0.89	0.94	0.89	0.94

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

D.6 Decomposing along the Forecast Error Sign

Table A33: Cost and Revenue Responses - Heterogeneity along the sign of Forecasts Errors: Forecasts - Realization (1 month forecasts)

Dependent Variables:	Sales	Costs	Profit	Value Added
Model:	(1)	(2)	(3)	(4)
Variables				
GDD	7.59	2.33	-3.98	6.73
	(10.5)	(7.56)	(13.2)	(12.8)
GDD (F)	37.1	28.6**	3.8	18.3
	(33.9)	(12.8)	(45.1)	(45.1)
HDD	-84.5	-185	322	179
	(392)	(193)	(291)	(298)
Forecast Bin 1	-3,509**	-1,806	-4,642**	-4,291*
	(1,637)	(1,159)	(2,094)	(2,209)
Forecast Bin 2	-3,378*	-1,169	-3,735**	-3,962**
	(1,895)	(752)	(1,673)	(1,722)
Forecast Bin 4	895	-898	613	340
	(2,731)	(864)	(2,569)	(2,601)
Forecast Bin 5	675	-1,327	999	1,755
	(4,605)	(1,353)	(2,959)	(3,184)
Mean	155,386	123,249	86,695	58,357
Unique Farms	2,603	2,603	2,603	2,603
Fixed-effects				
Farm	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Fit statistics				
Observations	18,917	18,917	18,917	18,917
R^2	0.89	0.94	0.84	0.80

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