The Cost of Adaptation to Heat and Climate Change: Evidence From French Agriculture*

Emma Du Puy[1] and Jeffrey G. Shrader[2]

[1]Sciences Po
[2]Columbia University

August 18, 2026

Abstract

Climate policy requires credible estimates of the costs of adaptation, yet such estimates are scarce. Combining a panel of French farms with realized and forecasted weather, we use forecasts as information shocks to identify the costs and benefits of anticipatory responses to heat. In the year heat occurs, forecasts cause inexpensive responses that generate large benefits, because they reduce the likelihood of harmful under-preparation. Costs arrive later, eventually offsetting these gains. Projected to 2100, adapting to climate change would cost 27% of average profit. Finally, existing estimators in the literature would overstate direct damages of heat and static costs.

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

*We are grateful for comments from Sylvain Chabé-Ferret, Andrew Hultgren, Derek Lemoine, and Wolfram Schlenker as well as numerous participants in seminars at NBER EEE, ASSA annual meeting, Columbia University, the MIT/NOVA School of Business Lisbon Workshop on Climate Change Adaptation, London School of Economics, Paris School of Economics, Toulouse School of Economics, the Transatlantic Kellen Conference Series, the Big Sky Workshop on the Economics of Agriculture and the Environment and the Heartland Environmental and Resource Economics Workshop. 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. The damages from climate change are a function of both the direct effects of climate change and the costs people bear when adapting to a changing climate.[1] Adaptation costs also feed back on the direct effect: the extent to which people pay such costs determines their ultimate exposure to a changed climate. Thus, adaptation costs are fundamental to understanding all aspects of climate change damages. Despite the importance of adaptation costs, there is currently little direct evidence on them, driven by a lack of data. In this paper, we use uniquely rich, farm-level panel data from France to provide novel estimates of the costs of adaptation in agriculture, one of the sectors of the economy most directly affected by the climate.

Because direct cost data are scarce, the standard approach in the literature uses a revealed preference argument to infer adaptation costs indirectly (see, e.g., Carleton et al., 2022; Mérel et al., 2024; Hultgren et al., 2025). Optimizing farmers choose adaptation actions such that the marginal benefit and marginal cost of each action is equal. Researchers exploit this equality to bound the costs of adaptation using adaptation benefits, the latter of which can be estimated from more widely available output data.[2] In contrast to this literature, we use data on the French agriculture sector that contains accounting measures for both revenue and—crucially—costs directly at the farm level. Using these data, we are able to estimate both the marginal costs and marginal benefits of farmer adjustments to forecasted weather shocks while controlling for realized weather shocks. These estimates allow us to test the standard revealed preference argument by testing for equality between these estimates, rather than assume it.

The use of forecast variation while controlling for weather realizations also allows us to overcome an identification challenge. Adaptation is endogenous, so simply regressing an outcome like revenue or costs on weather cannot separate what the weather does to the farmer from what the farmer does in response to the weather. Forecasts let us separate these two effects. To see why, imagine an experiment: tell one group of farmers that heat is coming and let them respond as they normally would, tell another group nothing, and then expose both groups to the exact same heat. The farmers with information will be able to undertake ex ante protective actions. The difference in costs and revenues across groups would isolate the effect of the responses and their effectiveness at mitigating the damages from heat. Forecasts approximate this experiment in the data. Conditional on the forecast, the heat that actually arrives is a surprise, so farmers cannot have responded to it. Conditional on the realized heat, variation in the forecast is variation in who was provided information with which to respond.

In an observational analysis that only uses realized weather, one cannot separate the treatment group (farmers who have responded) and the control group (farmers who have not). A direct implication is that the sole reliance on realized weather variation to identify costs and benefits of actions will carry a bias related to the value of information: because weather is likely correlated with

[1] In other words, adaptation costs are a form of defensive expenditure, and reducing defensive expenditures is a benefit from reducing a negative externality like greenhouse gas pollution (Freeman et al., 2014; Deschênes et al., 2017).

[2] In studies of climate impacts on agriculture, effects on crop yields are commonly estimated, with additional papers studying profits and farm land values. For a review, see(Carter et al., 2018).

expectations of that weather—and therefore the actions farmers take based on their beliefs—the effect of weather on outcomes will also carry the value of being able to better target anticipatory responses.

In our empirical analysis, we estimate the effect of seasonal weather forecasts issued up to four months in advance. Thus, our estimates speak to anticipatory, within-season adaptation such as changes in the timing of planting and harvest, crop mix, and input use. We do not capture durable adaptation like capital investment, farm relocation, or the development of new seed varieties (Costinot et al., 2016; Hagerty, 2022; Moscona and Sastry, 2022; Obolenski, 2025). Transitory, short-run adaptation is nonetheless important for understanding the total economic impact of climate change (Lemoine et al., 2025). It is a type of adaptation with low fixed cost that will likely continue to be relevant as the climate changes and farmers change the frequency with which they engage in it. And prior work has shown that such actions can be important for avoiding climate change damage (Reidsma et al., 2010; Rising and Devineni, 2020; Colmer, 2021).

Relying on a simple model of farm production, we start by deriving an expression for the bias related to the value of information about upcoming weather conditions and use the expression to establish conditions under which it will arise. We then turn to showing that these conditions are met in our specific context. First, we show that farms do undertake actions prior to weather being realized. Here, we show that farms respond to heat forecasts by undertaking anticipatory actions in the form of crop switching, moving away from heat-sensitive crops like sunflower and corn and towards heat-resilient crops such as wheat and colza. The presence of anticipatory actions is a precondition for a bias related to expectational errors to arise.

Second, we show that farm profits are asymmetric in forecast errors. Specifically, relative to an accurate forecast, under-forecasting heat (and the subsequent under-preparation of farms) reduces revenue and profits. Over-forecasting heat has no significant effect on farm profits. As such, holding realized heat constant, an increase in heat forecasts will make under-preparing for heat less likely and should increase farm profits. The underlying asymmetry of the profit function in forecast errors is the second condition generating the aforementioned information-related bias. As a consequence of this bias, regressions of realized profit on realized weather will overstate the direct damages of heat, because they silently absorb the value of having been able to prepare for it.

With these two results both on mechanisms and the effect of forecast errors as context, we turn to estimating the main outcomes of the paper: the marginal benefits and costs of farmer responses to forecasts. In the year that heat arrives, we find that the costs of farmers' anticipatory responses are small and not statistically distinguishable from zero, while the benefits are large and immediate. Bounding those costs using benefits alone, in the style of the revealed preference approach, would put them at around one percent of total annual production costs, substantially overstating adaptation costs in that first year of heat. But looking beyond the year of the shock changes the picture. Over the following two seasons, costs rise (and, correspondingly, profits fall), driven mainly by the crop switching farmers undertook to prepare for the heat: moving away from sunflower and corn and toward wheat and colza likely depletes soil nutrients and raises exposure to pests, which shows up later as higher fertilizer, pesticide, and labor spending. Netted out over four years, the discounted profit effect of the response is close to zero, with cumulative costs reaching about three percent of average annual profit.

Because we estimate both the costs and the benefits of adaptation directly, rather than inferring one from the other, we can test an assumption that the revealed preference literature typically has to impose: that farmers optimize, causing the marginal benefit and marginal cost of adaptation to be equal. We find that in the year of the shock, they are not. Marginal benefits substantially exceed marginal costs, so a static reading of that assumption fails. But once we let costs accumulate over the following seasons, the two converge, and we cannot reject equality in the dynamic version. The static envelope theorem argument, in other words, does not hold, but the dynamic one does, which points toward using a dynamic rather than static version of this approach whenever it is applied elsewhere in the climate damages literature.

In a last section, we use the farm responses to investigate adaptation to climate change over coming decades. This exercise first requires projections of how climate change will impact French weather, and second to understand how farm actions will shift with changes in the climate. On the first point, we use CMIP6 climate change predictions until the end of the century. On the second, we hold the weather exposure-response function constant, we then extrapolate the costs of adaptation. Under these conditions, climate change is expected to increase the yearly costs of adaptation from 3% of average current-day yearly profits to 27% by the end of the century. We finally discuss the extent to which we can expect farms to shift their adaptation behavior. For this exercise, we rely on a combination of Ricardian and panel approaches to climate change damage estimation (Auffhammer, 2018). Specifically, we look at costs and benefits 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, which implies that the adaptation practices identified here might not remain similarly successful in the future. If their efficiency wanes, we might expect the expenditures associated with them to fall, and farms to either turn to newer and costlier strategies, or face higher direct damages from heat.

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).

To understand the effects of climate change on the economy, agriculture is a central part of the story. From the earliest analyses through to the most recent updates to the U.S. Federal Government's estimates of climate change damage, agriculture has been recognized as one of the economic sectors most affected (Schelling, 1992; EPA, 2023). Agriculture is also an ideal setting for the study of climate change, as the physical consequences of temperature are well understood, and representative spatial data tends to be more readily available. Our paper stands out 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 are 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). Global analyses have shown that climate affects agriculture all around the world (Hultgren et al., 2025). We build on these papers by explicitly estimating cost and revenue adaptive responses, and relating those to climate adaptation costs and benefits. These estimates reveal that farmers act on information, leading to an omitted effect in direct damage estimates that rely on realized weather and allowing us to directly test the revealed preference assumption underlying previous studies.

Our paper thus also relates to the broader literature on sufficient statistic approaches (Chetty, 2009; Kleven, 2021). Most directly, recent papers use such an approach to estimate climate adaptation costs by bounding those costs with estimates of adaptation benefits, and use responses to weather variation as an estimate for the direct effect of weather net of adaptation. We complement these papers in two ways. First, by directly estimating marginal costs and benefits of adaptation, we are able to test the two central implications of this approach, on the one hand that regressing realized profits on realized weather identifies the direct effects of exposure to heat (Hsiang, 2016), and on the other hand that the optimality of adaptation actions implies an equality of the marginal benefits and costs of adaptation (Carleton et al., 2022; Mérel et al., 2024; Hultgren et al., 2025). We show that statically, the presence of anticipatory actions combined with an asymmetry of profits in forecasts errors implies that the direct effects of exposure to heat will be combined with a bias related to the value of information. Further, we show that this asymmetry of profits derives from the dynamics of costs, which are only realized over time, while benefits of adaptation are immediate. As such, we find support for the dynamic, rather than static, application of these approaches.[3] Second, we use an alternative identification strategy that has pros and cons relative to other approaches. By using seasonal forecasts, we are able to include fixed effects and other controls that eliminate potential omitted variable bias, along the lines of Deschênes and Greenstone (2007), while still allowing for the estimation of adaptation costs and benefits, though only for a subset of possible adaptation actions as discussed above. Previous literature relies on comparisons across locations with hotter or colder average weather to estimate adaptation benefits, potentially capturing a broader set of adaptation actions at the cost of cross-sectional identification.[4]

Another approach to study adaptation costs and benefits is to construct a structural model that explicitly specifies the possible adaptation mechanisms to be studied. Bareille and Chakir (2023) pursue such an approach in the context of French agriculture. Again, this approach is complementary to our own. We can estimate adaptation benefits and costs without enumerating each possible mechanism while they can quantify the role of each mechanism included in the model.

We find substantial adaptation costs, but only dynamically and due to crop switching. 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

[3] Theoretical analysis of the agricultural sector has previously pointed to the applicability of the dynamic envelope theorem in this context (Scheinkman and Schechtman, 1983).

[4] Other papers which use comparisons between high and low frequency weather, or use lower frequency variation to estimate adaptation benefits include Butler and Huybers (2013); Burke and Emerick (2016); Auffhammer (2022).

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

Finally, 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. For example, Kelly et al. (2005) shows that steady state differences in equilibria do not fully characterize climate damages. Instead, transition dynamics between steady states play an important role. Building on this, Downey et al. (2023) show that even dynamics around steady states can be important to understand climate damages, as agents repeatedly adjust to climate shocks. Both of these results imply that beliefs are an important component to the response of agents to the climate. 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 adaptation 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 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 in our context, using only variation in realized weather induces an over-estimation of the negative profit impacts of extreme heat—extreme heat seems more damaging that it really is. Finally, Burlig et al. (2025) use a randomized control trial to study the effect of seasonal forecasts on agricultural decision making in India. Like ours, their paper shows that farmers update their beliefs based on forecasts, and use them to meaningfully change their production practices.

The rest of the paper proceeds as follows. Section 2 describes the revealed preference approach used either to justify ignoring or to bound adaptation costs in the literature, derives a sufficient statistic for its bias and its interpretation. 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 outline two main features of our framework that will give rise to the value of information related bias. Section 6 provides the static estimates of the consequences of ex ante responses to heat, while section 7 provides their dynamic counterpart. Section 8 discusses how our estimates inform costs of adaptation to climate change. Finally, section 9 concludes the paper.

2 Model

We model a profit-maximizing agricultural firm making production decisions during a growing season. The growing season has two stages: a first where weather is imperfectly observed via a forecast, and a second where weather is both fully realized and observed. At the beginning of the growing season, the firm receives a forecast, f, about the weather that will be realized later in the season. Based on this forecast, the firm chooses an initial vector of inputs, the ex ante actions, denoted x_1. Later in the season, the weather shock w is realized. We assume forecasts are unbiased, such that $f = w + \varepsilon$, with $\varepsilon \sim \mathcal{N}(0, \sigma^2)$. The farmer forms beliefs about weather conditional on the forecast realization, such that $\phi(w|f)$ is the conditional density of a realization w after observing forecast f. After observing the realized weather w, the firm can take a second set of actions, the ex post actions, x_2. Production then occurs and profits are realized, conditional on all actions and the realized weather. The firm's realized profit is given by the following expression:

$$\pi = pq(x_1, x_2, w) - c(x_1, x_2, w)$$

The firm chooses its actions sequentially. The ex post action x_2 is chosen to maximize profit with full knowledge of w and x_1. The ex ante action x_1 is itself chosen to maximize expected profit conditional on the forecast f, anticipating the future optimal choice of x_2. Throughout the analysis, we assume that profit is concave in farm actions, and that cost and revenue functions are continuously differentiable and satisfy the necessary regularity conditions for the dominated convergence theorem to apply. We also assume the firm is price-taking, and that prices are not influenced by weather or forecast realizations.[5] The firm's initial problem can be written as solving for the ex ante static value function $\overline{v}(f)$, where we assume that the farmer is risk neutral:[6]

$$\begin{aligned}\overline{v}(f) &= \max_{x_1} \mathbb{E}\big[\max_{x_2} \pi(x_1, x_2, w)|f\big] \\ &= \mathbb{E}\big[\pi(x_1^\star, x_2^\star, w)|f\big],\end{aligned}$$

with $x^\star = \{x_1^\star, x_2^\star\}$ denoting the set of expected profit-maximizing actions, where the first is ex ante optimal given the forecast, and the second is ex post optimal given the weather realization.

In the rest of this section, we use this framework to justify our focus on forecasts and ex ante actions when estimating cost responses to heat shocks. We further use the model to highlight how our approach avoids a likely bias in the usual estimation of cost responses and direct effects of weather. We specify this bias as stemming from expectational errors in ex ante adaptation, and recover a proxy for this bias that we will directly estimate in our data.

2.1 Implications of Optimization and the Envelope Theorem

Profit maximization implies the two following first order conditions at the optimal levels of action:

$$\begin{aligned}&\mathbb{E}\left[p\frac{\partial q}{\partial x_1}|_{x^\star} \;\Big|\; f\right] = \mathbb{E}\left[\frac{\partial c}{\partial x_1}|_{x^\star} \;\Big|\; f\right] \\ &p\frac{\partial q}{\partial x_2}|_{x^\star} = \frac{\partial c}{\partial x_2}|_{x^\star}\end{aligned}$$

[5] This assumption is justified by our empirical focus on weather and forecast variation after controlling for year and region-year fixed effects that partial out most general equilibrium effects from our analysis.

[6] We assume a risk neutral farmer to simplify the exposition of the role of expectational errors and adjustment frictions in the estimation of marginal responses to realized weather with ex post panel data.

Notice that the first relation implies expected optimality of the ex ante actions (conditional on the forecast), while the second implies the optimality of ex post actions once the weather is realized. An immediate application of the envelope theorem relates to the identification of the direct effect of small weather variations on profit, in our case on expected profit. Using the two equalities from above, we can write that at the optimum:

$$\begin{aligned}\mathbb{E}\left[\frac{d\pi}{dw}|_{x^\star}\ |f\right] &= \mathbb{E}\left[\frac{\partial \pi}{\partial w}|_{x^\star} + p\frac{\partial q}{\partial x_1}|_{x^\star} - \frac{\partial c}{\partial x_1}|_{x^\star} + p\frac{\partial q}{\partial x_2}|_{x^\star} - \frac{\partial c}{\partial x_2}|_{x^\star}\ \Big|\ f\right] \\ &= \mathbb{E}\left[\frac{\partial \pi}{\partial w}|_{x^\star}\ \Big|\ f\right] \\ &= \int \frac{\partial \pi(x_1^\star, x_2^\star, w)}{\partial w}\phi(w|f)dw.\end{aligned}$$

This implies that exogenous variation in weather should identify the direct effect of weather on optimized expected farm profit, exclusive of the effects of agent re-optimization. The optimality of farm actions also implies that the marginal cost of ex post responses to realized weather can be recovered from the measurement of the marginal benefit from those same ex post responses, and similarly that the expected marginal costs of ex ante actions can be recovered from their benefit counterpart:

$$\begin{aligned}& p\frac{\partial q}{\partial x_2}|_{x^\star}\frac{\partial x_2}{\partial w}|_{x^\star} = \frac{\partial c}{\partial x_2}|_{x^\star}\frac{\partial x_2}{\partial w}|_{x^\star} \\ & \mathbb{E}\left[p\frac{\partial q}{\partial x_1}|_{x^\star}\frac{\partial x_1}{\partial w}|_{x^\star}\ \Big|\ f\right] = \mathbb{E}\left[\frac{\partial c}{\partial x_1}|_{x^\star}\frac{\partial x_1}{\partial w}|_{x^\star}\ \Big|\ f\right]\end{aligned}$$

These equalities are useful in cases where cost data is scarce, but revenue data readily available. We call these applications revealed-preference-based estimation approaches. Such applications are often used in the climate economics literature to measure the consequences of heat shocks, which can then be extrapolated into consequences of climate change (Hsiang, 2016), and to measure costs of adaptation to heat then extrapolated into costs of adaptation to climate change (Carleton et al., 2022). However, these arguments are generally applied in a meaningfully different data context where realized yields, rather than expected profits, are observed.

2.2 Ex Post Data and the Value of Information Bias

Estimating the effect of weather on realized yields or profits can introduce a bias driven by expectational errors, which relates to the value of information. We identify two conditions for this bias to arise: first realized variables are used as a proxy for expected ones in the analysis, and second the adaptation actions taken under imperfect information prior to the realization of weather have adjustment costs. In other words, this second condition requires that ex post actions cannot perfectly substitute for ex ante actions. A simple way to represent this second condition in our model is to assume that x_1 contains a different set of actions than x_2. An example would be that x_1 corresponds to crop choice which is decided early in the growing season, and cannot be modified once the season has started. This condition could also be met if x_1 and x_2 correspond to the same set of actions, but bear different costs throughout the growing season.

Whether this bias is substantial or not further depends on additional contextual conditions: whether expectational errors are symmetrically distributed in the data (which we will refer to as a support condition), and the symmetry of the profit function in these expectational errors (a symmetry condition). Together, support and symmetry will determine the sign and magnitude of this bias.

Bias Approximation: To outline this bias, it is useful to introduce $x_1^\star(f)$, the ex ante optimal ex ante action, and $x_1^\star(w)$, the ex post optimal ex ante action, the one taken if one had known perfectly about incoming weather. Unless forecasts are perfect, or if there are no adjustment costs so all actions can be considered to be x_2 actions, it will generally be the case that:

$$\pi(x_1^\star(f), x_2^\star(w), w) \neq \pi(x_1^\star(w), x_2^\star(w), w)$$

In cases where weather is perfectly forecasted, or if we have no ex ante adaptation, we can still invoke the envelope theorem, and the total derivative of realized profit with respect to weather simplifies exactly to the partial derivative of realized profit to weather:

$$\frac{d}{dw}\pi(x_1^\star(w), x_2^\star(w), w) = \frac{\partial \pi(x_1^\star(w), x_2^\star(w), w)}{\partial w} + \underbrace{\frac{\partial \pi}{\partial x_1}}_{=0}\frac{\partial x_1^\star(w)}{\partial w} + \underbrace{\frac{\partial \pi}{\partial x_2}}_{=0}\frac{\partial x_2^\star(w)}{\partial w} = \frac{\partial \pi(x_1^\star(w), x_2^\star(w), w)}{\partial w}$$

Because $x_2^\star(w)$ is ex post optimal, the envelope theorem implies $\frac{\partial \pi}{\partial x_2}\frac{\partial x_2^\star(w)}{\partial w} = 0$. Furthermore, if weather is perfectly forecasted, we also have that $\frac{\partial \pi}{\partial x_1}\frac{\partial x_1^\star(w)}{\partial w} = 0$. This implies that in cases without ex ante adaptation or when weather is perfectly forecasted, the total derivative of realized farm profit with respect to weather is equal to the partial derivative of profit with weather. In other words, we can identify the direct effect of weather on farm profit using exogenous variation in weather.

In the presence of imperfectly forecasted weather and ex ante adaptation, the total derivative of realized profit with respect to weather is no longer equal to the direct effect of weather on profit. Using a first-order Taylor approximation of the total derivative of realized farm profit relative to weather, we show in appendix A that we can write:

$$\frac{d\pi(x_1^\star(f), x_2^\star(w), w)}{dw} \approx \frac{\partial \pi(x_1^\star(w), x_2^\star(w), w)}{\partial w} \underbrace{-\frac{\partial^2 \pi}{\partial x_1^2}\left(\frac{\partial x_1^\star(w)}{\partial w}\right)^2 \varepsilon}_{\equiv\ \Omega\ \text{(Bias)}} \tag{1}$$

This total derivative differs from the direct effect by a second term, denoted Ω. Ω captures the marginal change in the loss from ex post mis-adaptation as realized weather moves away from a given forecast. In that sense, it is related to the value of ex ante information about weather. Point-wise, given the concavity of the profit function, Ω carries the sign of the forecast error ε. In expectation, under a symmetric error distribution, its sign is determined by the asymmetry of the profit function, i.e. whether the profit consequences of mis-adaptation are larger when heat is under-predicted than when it is over-predicted.

Implications for the Bounding of Costs: We can also relate this bias to the bounding of marginal cost responses from marginal benefits. Here again, the issue arises from the substitution of realized variables for their expected counterpart. We show in appendix A that we can write the wedge between realized marginal costs and benefits of ex ante actions as follows, with λ a parameter capturing the relative responsiveness of the farm to forecasts relative to its responsiveness to weather:

$$\begin{aligned}\left(p\frac{\partial q(x_1^\star(f), x_2^\star(w), w)}{\partial x_1} - \frac{\partial c(x_1^\star(f), x_2^\star(w), w)}{\partial x_1}\right)\frac{\partial x_1^\star(f)}{\partial f} &= \frac{\partial \pi(x_1^\star(f), x_2^\star(w), w)}{\partial x_1}\frac{\partial x_1^\star(f)}{\partial f} \\ &= \frac{\partial^2 \pi}{\partial x_1^2}\frac{\partial x_1^\star(w)}{\partial w}\frac{\partial x_1^\star(f)}{\partial f}\varepsilon \\ &= -\lambda\Omega, \text{ with } \lambda = \frac{\frac{\partial x_1^\star(f)}{\partial f}}{\frac{\partial x_1^\star(w)}{\partial w}}.\end{aligned} \tag{2}$$

Point-wise, the difference between the realized marginal costs and benefits of anticipatory actions differs by a term which is first order in the forecast error. One way to interpret this result is simply that forecast errors generate ex post mis-optimization, which moves us away from the point where marginal costs and benefits equal each other once the weather is realized.

Empirical Relevance of Information-related Bias: The bias driven by expectational errors can be larger or smaller depending on: 1) whether the forecast error distribution is symmetric around the realization and 2) the shape of the profit function in relation to the forecast errors. The first will determine the support on which this bias is estimated.[7] The second will determine whether we can expect expectational errors to offset each other over the support of the error distribution. When the errors' support is symmetric, we can expect that the bias will be negligible for profit functions that are symmetric in the forecast errors, i.e. if farmers experience comparable losses for similarly large positive and negative forecast errors. On the contrary, if the profit function is asymmetric in the errors, the bias is unlikely to disappear.

A static asymmetry in the profit function can arise through two channels, which are both relevant for many applications in climate economics. The first is a curvature configuration: if the benefit (revenue) side is convex in the forecast error while the cost side is approximately linear, the expectation of realized benefits of the anticipatory actions can exceed the expectation of realized costs over a symmetric error distribution. The second, which we find is the relevant one in our setting, is purely temporal: the benefits of anticipatory actions are realized immediately, as farmers avoid the losses from exposure to unprepared-for heat, while their costs are dynamic switching costs realized only in subsequent seasons. Statically, this makes under-preparation costly and over-preparation nearly free, even though the costs and benefits of the anticipatory action equate dynamically. Section 7 provides direct evidence for this temporal channel, which is why we reject the static application of the revealed-preference-based method but cannot reject its dynamic counterpart.

[7] In our case, our conditional forecast error distribution, after including department and year fixed effects, is symmetric around the realized weather. Our estimates will then be average effects over a symmetric distribution of errors.

Examples of symmetric and asymmetric profit functions: We provide two simple examples of symmetric and asymmetric profit functions. These are used to highlight how the bias varies with the symmetry of the profit functions.

Consider first the symmetric profit function $\pi(x_1, x_2, w) = -\Big(w - x_1\Big)^2$. In this context, farmers are required to precisely target weather, over-adapting and under-adapting are similarly costly. In this case, the optimal ex ante action given that the farmer received a forecast f is $x_1^\star = f$, and $\frac{\partial \pi}{\partial x_1}|_{x_1^\star} = -2\varepsilon$, as well as $\frac{\partial^2 \pi}{\partial x_1^2}|_{x_1^\star} = -2$. Here, we do have that in expectation, for a symmetric distribution of the forecast errors:

$$\mathbb{E}\Bigg[\frac{d\pi(x_1^\star(f), x_2^\star(w), w)}{dw}\Bigg] \approx \mathbb{E}\Bigg[\frac{\partial \pi(x_1^\star(w), x_2^\star(w), w)}{\partial w}\Bigg] + \underbrace{\mathbb{E}\Bigg[+2\Big(\frac{\partial x_1^\star(w)}{\partial w}\Big)^2 \varepsilon\Bigg]}_{=0}$$

Consider in turn an asymmetric profit function. Let the curvature of the profit function change with the sign of the forecast error such that: $\pi(x_1, x_2, w) = -\gamma_-\Big(w - x_1\Big)^2 \mathbb{1}\Big\{\varepsilon < 0\Big\} - \gamma_+\Big(w - x_1\Big)^2 \mathbb{1}\Big\{\varepsilon \geq 0\Big\}$. We assume that under-preparing for heat (when $\varepsilon < 0$) is more costly that over-preparing ($\gamma_- > \gamma_+$), with a limit case of $\gamma_+ \to 0$ representing statically costless over-adaptation. We now have $\frac{\partial^2 \pi}{\partial x_1^2}|_{\varepsilon<0} = -2\gamma_-$, and $\frac{\partial^2 \pi}{\partial x_1^2}|_{\varepsilon \geq 0} = -2\gamma_+$. This implies in turn the following:

$$\mathbb{E}\Bigg[\frac{d\pi(x_1^\star(f), x_2^\star(w), w)}{dw}\Bigg] \approx \mathbb{E}\Bigg[\frac{\partial \pi(x_1^\star(w), x_2^\star(w), w)}{\partial w}\Bigg] + 2(\gamma_- - \gamma_+)\mathbb{E}\Bigg[\varepsilon \mathbb{1}\Big\{\varepsilon < 0\Big\}\Bigg]$$

With a symmetric error distribution and a profit function more sensitive to under-preparation ($\gamma_- > \gamma_+$), the expected bias is negative. Because the direct effect of heat on profit is itself negative, a negative bias makes the realized-weather estimate more negative than the true direct effect: standard estimates overstate the direct damage of extreme heat, because the realized-weather coefficient also absorbs the within-season losses farmers incur when they mis-adapt to it.

2.3 Our Approach: Ex Ante Responses to Forecasts

Ex Ante Costs: The first goal of this paper is to recover the marginal costs of ex ante responses to heat, and to recover them from variation in forecasts while controlling for realized weather. This approach is an alternative to the revealed-preference-based approach outlined in section 2.1. We can write the total derivative of realized costs with respect to forecasts, conditional on realized weather, as:[8]

$$\begin{aligned}\frac{dc(x_1^\star(f), x_2^\star(w), w)}{df}|_w &= \frac{\partial c(x_1^\star(f), x_2^\star(w), w)}{\partial x_1}\frac{\partial x_1^\star(f)}{\partial f} \\ &\quad + \frac{\partial c(x_1^\star(f), x_2^\star(w), w)}{\partial x_2}\frac{\partial x_2^\star(w)}{\partial f} + \frac{\partial c(x_1^\star(f), x_2^\star(w), w)}{\partial w}\frac{\partial w}{\partial f} \\ &= \frac{\partial c(x_1^\star(f), x_2^\star(w), w)}{\partial x_1}\frac{\partial x_1^\star(f)}{\partial f}\end{aligned}$$

Here the second term of the first line disappears because, conditional on weather, ex post actions do not vary with forecasts and do not impact costs. The same logic applies to the third term. Because we directly observe costs, we can focus on the estimation of their marginal response in the data. Furthermore, we can assume the costs of ex ante responses to be fairly deterministic, in many instances immediately realized when actions are taken. As such, we can also take the response of realized costs to forecasts as a good proxy for the response of expected costs to forecasts, which is the true policy object influencing farm decisions. With this approach, we can then obtain an unbiased estimate of ex ante cost responses, trading a narrower focus on within-season ex ante actions for cleaner identification. This narrower focus is highlighted by the disappearance of the second term in the equation above, e.g., our estimation will miss the costs of ex post responses and ex ante responses that are not moved by seasonal forecasts.

Profit Effects: In the paper, we also focus on the marginal response of profits and revenues to forecasts, conditional on realized weather. We argue that these estimates are a good proxy for the expectational bias discussed in equation (1) and equation (2). Our design recovers the marginal response of realized profit to the forecast, holding realized weather fixed:

$$\begin{aligned}\frac{d\pi(x_1^\star(f), x_2^\star(w), w)}{df}|_w &= \frac{\partial \pi(x_1^\star(f), x_2^\star(w), w)}{\partial x_1}\frac{\partial x_1^\star(f)}{\partial f} \\ &\approx \frac{\partial^2 \pi(x_1^\star(w), x_2^\star(w), w)}{\partial x_1^2}\Big(x_1^\star(f) - x_1^\star(w)\Big)\frac{\partial x_1^\star(f)}{\partial f} \\ &\approx \frac{\partial^2 \pi}{\partial x_1^2}\frac{\partial x_1^\star(f)}{\partial f}\frac{\partial x_1^\star(w)}{\partial w}\varepsilon\end{aligned}$$

In the second line, we use a first-order Taylor expansion of the marginal effect of forecasts on profits evaluated at the ex ante optimal action $x_1^\star(f)$, around the ex post optimal action $x_1^\star(w)$. We can then relate our estimate to the bias in equation (1), and to the expression for equation (2):

$$\frac{d\pi(x_1^\star(f), x_2^\star(w), w)}{df}|_w \approx -\lambda\Big(-\frac{\partial^2 \pi}{\partial x_1^2}\Big(\frac{\partial x_1^\star(w)}{\partial w}\Big)^2\varepsilon\Big) = -\lambda\Omega$$

In cases where $\lambda = 1$, the estimated marginal effect recovers the negative of the theoretical bias term of equation (1). This equality will stand if the policy function governing ex ante actions

[8] This total derivative with respect to forecasts, holding weather fixed, means that we leverage variation in the forecast errors ε.

maps expectations to inputs with the same functional form as perfect foresight maps realizations to these same actions ($\lambda = 1$). In general, we will recover the negative of the bias times the sensitivity parameter λ.[9] We can interpret this statistic as a proxy for the ex ante value of information. In the paper, we will find that this statistic is positive, and provide evidence that the profit function is asymmetric in the forecast errors. As such, because this statistic is the negative of the bias, our results imply that standard revealed-preference-based methods overstate the direct damages of heat.

3 Empirical Context and Data

3.1 Exposure to Weather in French Agriculture

We use 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, following Schlenker et al. (2005); Braun and Schlenker (2023). The adoption rate of commodity futures trading by farmers is also very low in Europe (Michels et al., 2019). This implies that farmers face the full consequences of unexpected weather shocks, and further motivates the use of forecasts by farmers in order to reduce such damages.

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 A24 shows estimated 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, temperatures below 0°C are not very common, and on average extreme heat temperature remains around 30°C. Figure A7 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 moderate and hot temperature days in France over our time period.

To characterize weather sensitivity in our sample, 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 figure A12 looks at the aggregate effect at the farm level. These results show that crop yields have

[9] In the presence of risk aversion, we could expect the x_1 policy function in cases where weather is perfectly vs. imperfectly forecasted to differ, and λ can serve as a reduced form representation of this sensitivity to risk.

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.

Our main empirical 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 with lead times of four months ahead (Canal, 2015). For our analysis, we use hindcasts produced by the European Center for Medium-Range Weather Forecasts (ECMWF), which are available from 1994 onwards. 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 and which we find drive the bulk of our results. Today, different companies offer 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 provides three-month-ahead weather bulletins which are made public on its website.

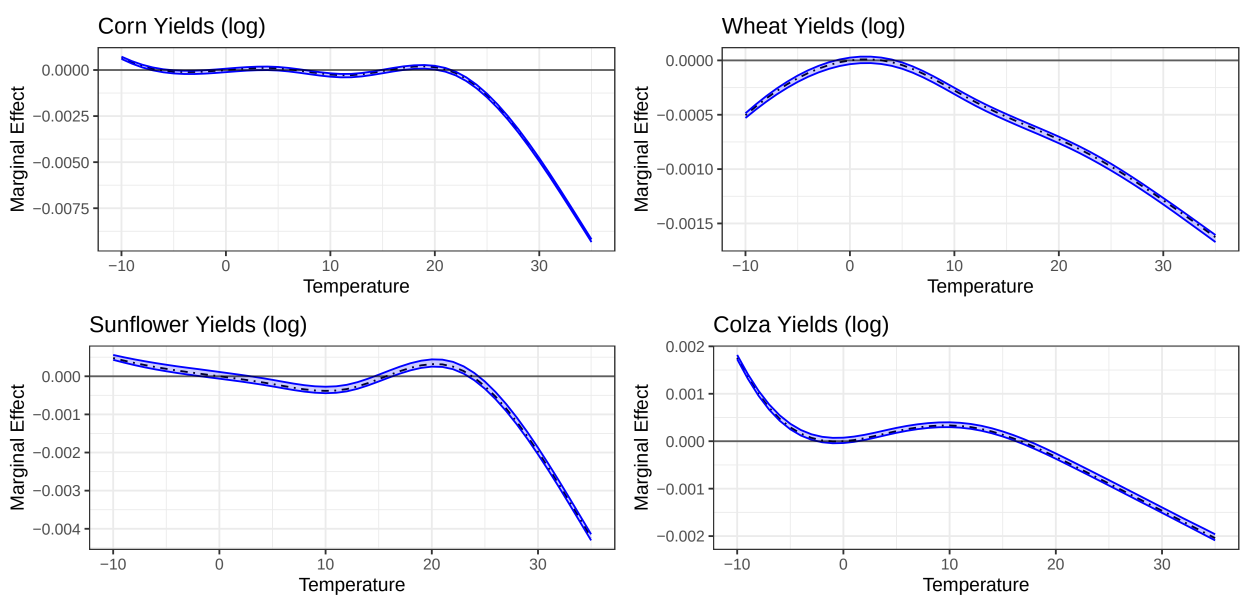

Figure 1: Temperature Effects on Crop-Specific Yields

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.

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.[10]

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 C.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.[11] 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.[12] While farm 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.

[10] 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.

[11] 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.

[12] 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).[13] 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.[14] We localize 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 A17 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.

Finally, department-level average agricultural wages are taken from the continuous labor survey "Enquête Emploi."

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.[15] 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.

[13] 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.

[14] 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.

[15] Plots surveyed are selected among the farms that benefit from the European Union's Common Agricultural Policy. The survey focuses on land plots defined as the set of contiguous land for which the same crop is cultivated, with homogeneous agricultural practices (fertilizer and pesticide use for example). For each crop, the survey selects the minimum number of regions covering at least 95% of that crop's production, and within each region the minimum set of departments accounting for at least 90% of the region's production. Within departments, the survey selects farms with at least .1 hectare cultivated, and less than 200ha. A unique plot is selected within each farm. For the waves that we study, around 20,000 plots are sampled each time.

3.3 Climate Data

3.3.1 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 (see, e.g., Schlenker and Roberts, 2009), we measure temperature exposure using moderate temperature or growing degree days (GDD) and hot temperature degree days (HDD), sometimes referred to in the literature as killing degree days or extreme degree days. 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 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.[16]

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 A16 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 A10 and figure A11 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.[17]

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

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

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

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.

In this paper, 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. In appendix B.1.5, we use a specification in every point similar to the main regression equation used throughout this paper, and presented in equation (3). We plot in these figures the marginal impact of a higher seasonal HDD forecast on short run month-specific forecasts about expected cold weather. Conditional on all our baseline controls and fixed effects, these show that our seasonal HDD forecasts correlate strongly with short-term temperature forecasts received by farmers throughout the year: specifically with forecasts emitted on the first day of every month—for the next 30 days—predicting how cold the weather is expected to be.[18] As such our seasonal forecast for HDDs correlates well with short run forecasts received by farmers throughout the year, and will be able to track farmers' responses taken not only in the spring before the onset of the hot summer, but also in the fall in preparation for upcoming heat. These results confirm that farmers can predict a hot summer from observations of abnormally hot winters, and preemptively change their production decisions. In appendix B.1.5 we also document some heterogeneity in this response. For most of our sample, the bottom 90% of the distribution in terms of average forecasted HDD realizations, a higher seasonal HDD forecasts is significantly associated with hotter predicted weather in October through December, as well as in May and June. For the top 10% of the sample, where forecasted HDDs are much more frequent, a higher seasonal HDD forecast will only correlate with

[18]French farmers receive repackaged forecasts from Météo France throughout the year, which predict the probability that the upcoming months will be more or less hot than average. These are the main type of longer-run forecasts we expect them to use when making crop choice decisions. Here we proxy the likelihood that weather will be colder with a measure which computes the forecasted number of degree hours spent below resp. 1.64, 1.96 and 2.58 standard deviations below the average department-specific hour of the day of the month temperature. See https://meteofrance.fr/actualite/publications/les-tendances-climatiques-trois-mois

hotter predicted weather in the summer. While being in keeping with the seasonal aggregations used throughout the agricultural climate change literature, our seasonal heat forecasts will hence be able to track the cost and benefits of actions taken by farms throughout the year, in expectation of extreme heat.

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 A5 and figure A6, 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 A3, 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.[19] The error distribution for HDDs exhibits a slight skew, consistent with the calibration plots showing that forecasts tend to under-predict extreme heat. Importantly, once we include fixed effects, in this case department and year which are the levels at which seasonal forecasts vary, both HDD and GDD forecast distributions are centered around the realizations (figure A2). In the rest of our analysis, we can then take these forecasts to be unbiased conditional on similar or tighter fixed effect structures.

Splitting our sample between high and low average HDD quantiles reveals that forecast errors are slightly larger in departments where HDD realizations are more frequent. The bottom half of our sample has an average error of 1.21, meaning that on average forecasts under-predict HDDs by 1.21 units (standard deviation of 2.14). In the top half, that error is of 2.72 (standard deviation of 3.73). However, this difference is not statistically significant, and the gap in forecast errors between these groups is smaller than the gap in average HDD realizations (1.41 vs 3.43), which means that forecasts do track to some extent this spatial heterogeneity in realized 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 anticipatory actions, as well as test for the empirical validity of the revealed-preference-based approach, we use an estimating equation that regresses profits, total costs, or total revenue on both realizations and forecasts of weather. In particular, we

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

use the following estimating equation.

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

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.[20]

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 optimized decisions like profit, 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. When focusing on profit, as we have argued in the previous section, the estimation of direct effects off of weather realization is likely contaminated by a bias dependent on the value of information.

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 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 $\zeta_{r(j)}^1 t$ and $\zeta_{r(j)}^2 t^2$. These account for potential sub-national trends that would correlate with weather and our outcomes of interest.[21] Finally, farm fixed effects, γ_j, and year fixed effects, η_t, mean that effects

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

[21] 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

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 The Structure of Ex Ante Actions

In this section, we document two key elements that jointly generate an expectational error bias in static applications of the envelope theorem: first, that some actions with meaningful adjustment costs are taken ex ante based on imperfect expectations about weather, and second, that realized profits are asymmetric with respect to the sign of these expectational errors.[22] We start by documenting that French farmers respond to HDD forecasts by altering their crop mix, reducing their reliance on heat-sensitive crops and expanding the footprint of more resilient varieties. This behavior provides evidence of ex ante adaptation to heat. Because these actions can only be taken before seasonal weather is fully realized, they entail meaningful adjustment costs: once sown, crop mixes are fixed for the season. We then demonstrate that realized profit is asymmetric in HDD forecast errors.

5.1 Crop Switching as Evidence of Ex Ante Actions

Crop switching is widely recognized as a crucial mechanism for mitigating agricultural damages from heat and climate change (Rising and Devineni, 2020). This strategy operates along two distinct margins: sensitivity and exposure. First, regarding sensitivity, farmers anticipating hotter weather can substitute toward more heat-resilient crops to minimize potential losses. As highlighted by the specifications in figure 1, pairs such as corn and wheat, or sunflower and colza, exhibit differing heat tolerances; substituting corn for wheat, or sunflower for colza, effectively lowers a farm's baseline sensitivity to heat shocks. Second, crop switching alters the timing of weather exposure. As shown in figure A10 and figure A11, winter wheat in France is typically sown in November and harvested before August, whereas corn is sown in May and harvested in October. Shifting land into wheat when a hot summer is forecasted allows farmers to entirely bypass late-summer heat shocks, thereby reducing the effective accumulation of HDDs. Finally, as established in subsection 3.3,

(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.

[22] As outlined in section 2, these two elements can be combined with the symmetric distribution of conditional forecast errors documented in figure A2, and the use of realized profit as a proxy for expected profit. Together, they imply that usual approaches to recover static costs of adaptation and direct effects of weather will be biased by the inclusion of the effect of expectational errors.

farmers receive short-term forecasts throughout the year, providing the necessary information to draw inferences about summer heat anomalies and preemptively adjust their exposure through these crop allocation decisions.

To test for this channel, we estimate our baseline model from equation (3) using the land area dedicated to specific crops within a farm as the outcome variable. Figure 2 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 heat sensitivity observed in the output results we discuss later (table A13).[23] The increase in land dedicated to peas and beans matches cropping patterns observed by Aragón et al. (2021).

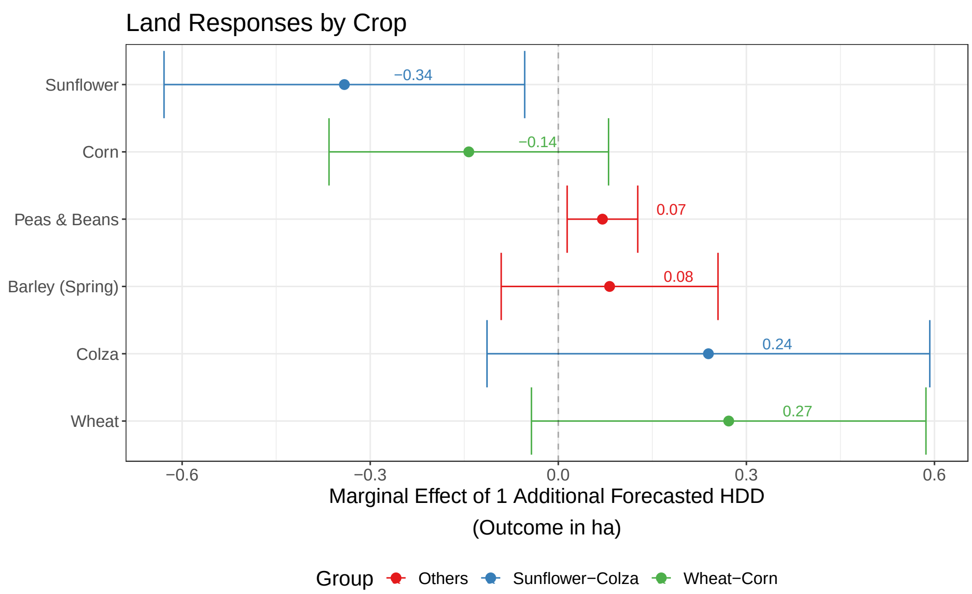

Figure 2: Decomposition of Cropland Responses

Notes: Estimates display the response to one month ahead forecasted heating degree days of crop-specific land allocations at the farm level. The regressions 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. Standard errors are clustered two-way department-by-year, and confidence bands display the 95% confidence interval.

Additional Channels of Ex Ante Responses In the appendix, we investigate an additional margin of ex ante responses: whether farmers adjust the timing of their within-season production actions based on weather forecasts. For this analysis, we utilize plot-level data from the PKGC survey detailing the dates of ploughing, sowing, and harvesting for our primary crops. Given that

[23]Results are shown for stacked sunflower-colza and corn-wheat regressions in table A8. These allow to test whether farms are moving their crop mix in a specific way, within oil crops from sunflower to colza and within cereals from corn to wheat. Wald tests reject the null of an equal response for colza and sunflower at the 5% level (p=0.048), but fail to reject for corn and winter wheat (p=0.29). Note that crop switching in response to heat shocks are not bound to remain within grain category, substitution away from corn is likely to be towards another crop sown in spring, which is not the case for most French wheat.

this dataset is a repeated cross-section with coarser sampling than our primary panel, we replace farm fixed effects with department fixed effects. Furthermore, we cluster the standard errors at the department level rather than the department-by-year level to accommodate the limited number of time periods available ($T = 3$). Aside from these adjustments, the specification remains identical to the baseline model in equation (3). The results, presented in figure A18, demonstrate that farmers actively modify their production schedules in response to heat forecasts, with significant heterogeneity across crop types. While the timing for wheat remains largely unresponsive, farmers make substantial timing adjustments for more heat-sensitive crops. Specifically, dates for ploughing and sowing are shifted either forward or backward to shield crops from peak heat shocks or to optimize growth in a warmer-than-average season.[24]

The presence of ex ante responses implies that forecast errors can generate ex post mis-adaptation. We document next the consequences of this mis-adaptation respectively for small and large negative and positive forecast errors.

5.2 Non-Linear Response to Forecast Error

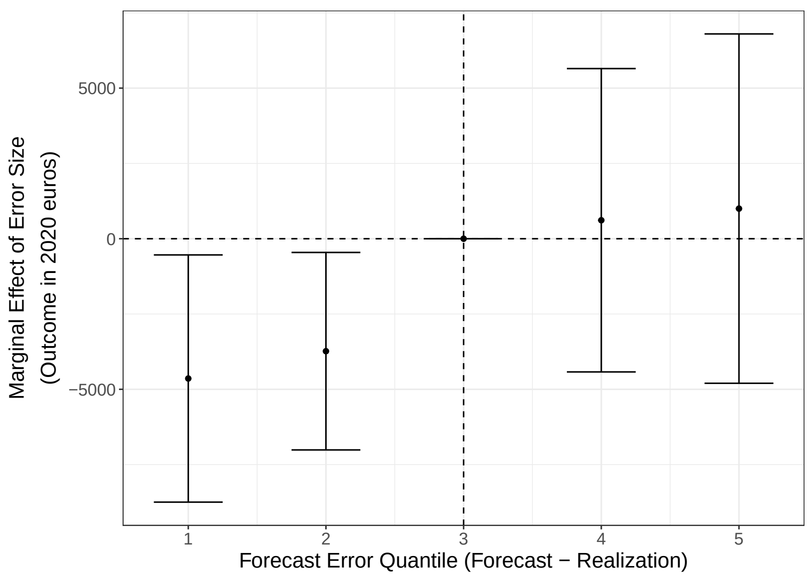

Figure 3: Profit and Sign of the Forecast Error Effect

Notes: Estimates display the non-linear response of farm profit to forecast errors. Error quantiles are labeled from 1 to 5, increasing in forecast error from the largest negative error where the realization is largest compared to the forecast, up to the largest positive error, where the forecast is largest compared to the realization. The omitted error category contains errors between -0.1 and 0.1. The regression is otherwise 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. Standard errors are clustered two-way department-by-year, and confidence bands display the 95% confidence interval.

[24] Table A14 through table A16 provide the full set of corresponding results.

Building on the evidence that farmers undertake ex ante adaptation subject to significant adjustment costs, we now show that their profit functions exhibit asymmetry relative to HDD forecast errors. The combination of these two features provides support for the information-based bias formalized in equation (3). To empirically test this, we decompose the forecast effects by both the direction and magnitude of the expectation error. We employ a specification identical to equation (3), modifying it only to separate the forecasted HDD variable into conditional quantiles. For positive forecast errors (forecasts exceeding realizations) of 0.1 or greater, we split the observations into two bins at the median of this positive-error distribution. An analogous median split is performed for negative forecast errors (-0.1 or lower). The intermediate range between -0.1 and 0.1 constitutes the central bin and serves as the omitted reference category. These errors are labeled in figure 3 from 1 to 5 by increasing forecast error value.[25] This figure plots the marginal effect on profit of an additional forecasted HDD within each bin.[26]

The results show an asymmetry: negative forecast errors significantly reduce farm profit, whereas positive errors yield small, statistically insignificant effects. Specifically, the largest negative forecast errors (bin 1) reduce average profits by €4.6k (5.3% of average profits), while smaller negative errors reduce profits by €3.7k (4.3%). The differences between the two coefficients should, however, be interpreted with caution, as larger errors are less frequent and yield larger standard errors in the estimate. In contrast, over-predicting heat generates no meaningful profit penalty; the estimates for both small and large positive errors are economically modest, at €0.6k (0.7%) and €1.0k (1.1%), respectively, and statistically indistinguishable from zero.

The negative profit effects stemming from under-forecasts of heat suggest that French farmers do use seasonal forecasts to form their expectations about incoming weather. Indeed, relative to precise heat forecasts, imprecise negative forecasts lead to profit losses. While farmers might discount large forecast errors, they still align their production decisions on them to the extent that these errors generate costly under-preparedness to heat.

Overall, the results are intuitive: negative forecast errors likely lead farmers to under-forecast future heat shocks, rendering them under-prepared for the impact of heat. On the contrary, over-forecasts could induce farmers to take protective actions which limit the negative consequences of realized heat on production. Actions in excess of the realized heat then matter little for the realized profit, at least within the time span of a growing season. Given the symmetric conditional distribution of HDD forecast errors show in figure A2, these results imply that in our sample, 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 to shield their profit from potential negative shocks. We next aggregate forecast error results into a single statistic capturing the marginal profit, revenue, and cost effects of additional forecasted HDDs.

[25] HDD forecasts tend to under-predict HDD realizations, and the median negative forecast error is -1.5, and the median positive one is 0.5. This skewness explains why coefficients for negative errors are slightly better identified.

[26] The table with all regression estimates is provided in table A28, with the cost and revenue dynamics further illustrated in figure A22.

6 Static Consequences of Ex Ante Actions

In this section, we estimate the static profit, cost, and revenue responses to GDD and HDD forecasts, highlighting that profits respond positively to forecasted HDDs (due to profit losses from under-forecasts, as shown above). We further discuss the interpretation of these positive profit effects as an estimate of the value of information that likely biases static applications of the envelope theorem used to recover both the direct effects of weather on farm profit and the costs of adaptation indirectly from revenue estimates. Finally, we explore several potential explanations for this positive static profit response.

6.1 Revenue, Cost, and Profit Effects

Table 1: Farm-Level Profit, Revenue, and Cost Responses 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: $^{*}p < .10$, $^{**}p < .05$, $^{***}p < .01$.

We first show the effect of forecasted moderate and extreme temperature on farm profits, revenues and costs, controlling for realizations of temperature. Table 1 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 cost responses. The profit variable corresponds to the farm's gross operating income, and includes value added, subsidies, expenses for insurance, and insurance indemnities.[27]

GDDs 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, local, idiosyncratic, adverse weather (either lower GDD or higher HDD) corresponds to a negative TFP shock, and 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. On the other hand, changes in the timing of harvest could allow for a positive revenue response without substantial changes in costs.

Table 1 shows that profit and revenues respond positively to forecasted GDDs and HDDs. In contrast, costs respond positively to forecasted GDDs but exhibit an insignificant, negative response to forecasted HDDs. And, most importantly, 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 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 actions. In our sample, these effects are also not statistically significant. Given that the central goal of the paper is to identify and quantify cost responses to forecasts, 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 near zero, and that costs and revenue have coefficients of similar magnitude. This result suggests that for moderate temperatures, farmers are taking costly actions to arrive at an increase in revenue, resulting in little net effect on profit. The effects on both revenues and costs are substantial: the standard deviation of forecasted GDDs is about 240, so a typical change in GDDs will lead to a roughly 6% change in sales and costs. This is in line with the known sensitivity of the agricultural sector to weather.

The effect of forecasted HDDs shows a different pattern. Profits and sales increase while costs do not. An increase in the HDD forecast by one degree-day one month in advance leads to an increase in profit by €2k, and an increase in revenue of €1k. Given typical revenue per farm of about €150k and a standard deviation of forecasted HDD of just over 1 (see table A2), this

[27] 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 C.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 A10.

coefficient indicates that a typical change in forecasted HDD causes revenues 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 substantial or significant effect on production costs.

These results are in line with the decomposition by the sign and magnitude of HDD forecast errors shown in figure 3 and figure A22. Because negative forecast errors are harmful to farm revenues, but positive ones are relatively harmless, an increase in forecasted HDDs raises farm revenues by reducing the likelihood of harmful mis-adaptation. Then, because costs show no response to HDD forecasts, the overall pattern we obtain is one where profits are increasing in forecasted HDDs.

These results capture the value of information, quantifying the economic gains from forecasts that enable farmers to mitigate production harms. Following our discussion in section 2, a positive value of information also implies a negative bias in the static revealed-preference-based approach to estimating the direct effects of weather shocks. Finally, the cost and revenue responses in table 1 confirm that the static profit function is asymmetric with respect to forecast errors: revenues exhibit asymmetry, while costs remain flat and unresponsive to these errors.

6.2 Information Timing

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

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

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 $FP_{d(j)t}^{\ell}$.

This regression serves two main purposes. First, it identifies the precise timing of information arrival, and is thus useful for understanding what type of action our analysis is recovering if different farm responses face different magnitudes of adjustment costs or need to occur at different times. When running regressions with a set of forecasts, variation in one forecast, conditional on all others, will capture the information received at that lead time. Second, and relatedly, while our focus remains on within-season adaptation, this design allows us to assess responses from a larger set of actions that might require longer lead time.

Equation (4) essentially decomposes the effect we show using a single forecast horizon in section 6.1. 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 identification, 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).[28] 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.

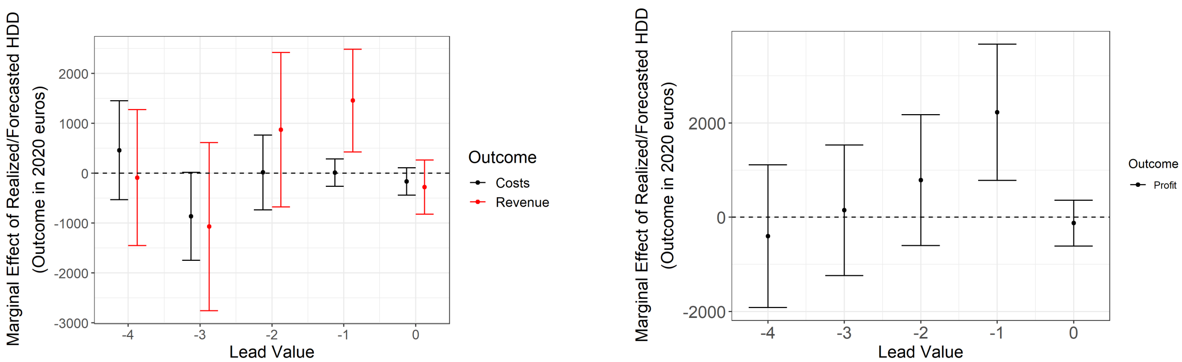

(a) Cost and Revenue

(b) Profits

Figure 4: Timing of Forecasts on Farm Profit, Cost and Revenue

Notes: Estimates display the response of farm costs and revenue (panel a) and profits (panel b) to forecasts of hot degree days received at different lead times from one to four months ahead (lead value -1 to -4), as well as the effect of realized hot degree days (lead 0). The regressions are based on equation (4), 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. Standard errors are clustered two-way department-by-year, and confidence bands display the 95% confidence interval.

Figure 4 plots the coefficients associated with hot degree days from estimating two different versions of equation (4). In particular, the figure shows β_ℓ^f for $\ell \in [1, 4]$ and shows β^w for the lead value of 0 for regressions with profit (panel a), costs and revenues (panel b) 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.

We see that the effect on revenues associated with a one-month-ahead forecast is statistically significant and of a similar magnitude to the effect found in table 1. The cost effect of a one-month-ahead forecast of HDDs remains centered around zero. Revenue responses are not statistically significant at the 5% level for horizons longer than one month, but the point estimates indicate that information close to the shock allows for revenue-enhancing responses while information farther from the shock leads to, if anything, a decrease in revenues. Costs are near zero for all horizons aside from three months ahead. The fact that cost effects do not consistently get stronger as the forecast horizon increases suggests that convex adjustment costs are not a major factor driving the results that we find (Downey et al., 2023). The three-month-ahead forecast effect could indicate that there are adjustments that are uniquely available at a quarterly frequency. Each month, the French

[28]In a model with non-convex adjustment costs, a specific forecast lead might generate a larger response at a given horizon because the optimal timing of a particular responses matches that lead value the most. Changes in forecast accuracy by lead could also cause different responses to different leads. 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.

meteorological agency also delivers a publicly available document with forecasts up to three month ahead, so the three months ahead forecast is the first forecast easily observable by farmers, which could explain why farmers do not respond to four-month-ahead forecasts.[29] The same regression looking at profits shows a clear pattern: news arriving one month ahead has a significant effect on profit while all other leads exhibit smaller and insignificant effects.

6.3 What is Driving Cost, Revenue, and Profit Responses?

In this section, we explore the robustness of our results to mis-measurement and offsetting movements, and identify the positive revenue gains from heat forecasts as directly coming from output gains. In the following section, we will explore dynamic explanations that can give rise to this asymmetric static profit structure.

Offsetting Prices and Quantities: We examine the revenue effect in more detail asking whether the response comes from agent reoptimization or equilibrium effects. To do so, we estimate versions of the equation (3) where the outcome variables are physical output (simply the sum of all produced quantities), the output price index, or storage (sum of all stored quantities).[30]

The results are shown in figure 5 (see table A7 for the table version). One can see that output rises in response to a change in HDD forecasts, consistent with the results shown above in table 1. Output prices, in contrast, do not change on average (though the effect is imprecisely estimated). Flat prices and increased quantities confirm that farmers are individually responding to forecasts, and that gains in farm profit and revenues are not a result of general equilibrium forces. Storage finally shows a relatively large, though insignificant, response, indicating that farm actions might have dynamic consequences—a point we return to below—and suggests that our revenue results in table 1 are an underestimate of the consequences of farm responses.

In table A13, 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 (0.7%), 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.[31]

[29] For an example of these publicly available and processed forecasts, see https://meteofrance.fr/actualite/publications/les-tendances-climatiques-trois-mois.

[30] We show in table A13 the crop-specific responses of output quantity to temperature realizations and forecasts. 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 C.1 gives a formal definition of the different variables.

[31] 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

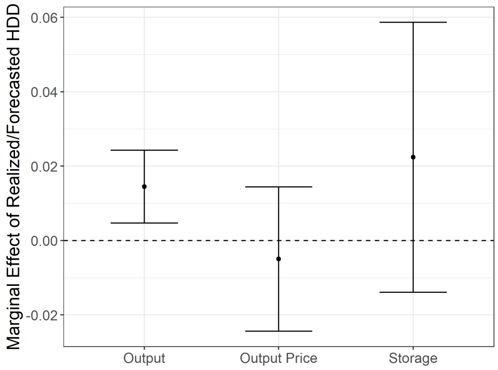

(a) Revenue Decomposition

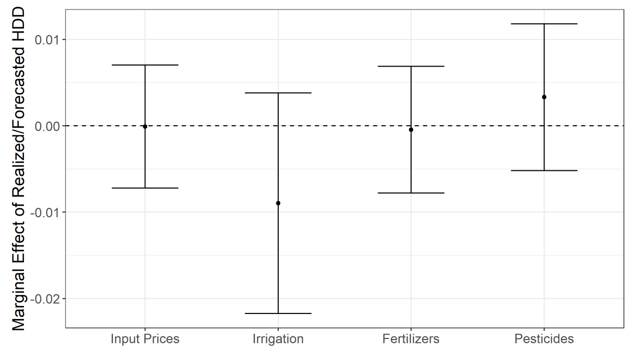

(b) Cost Decomposition

Figure 5: Prices and Quantities for (log) Outputs and Inputs

Notes: Estimates display the response to one month ahead forecasted hot degree days of, in panel a, farm output, output price, storage levels and, in panel b, irrigation, fertilizers and pesticides expenses. We also show the response of input prices as observed in an agricultural input price survey at the store-year-input level. The regressions 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. Standard errors are clustered two-way department-by-year, and confidence bands display the 95% confidence interval.

On the right hand side of figure 5, we look at the cost side and rule out potential agent re-optimization that would produce no net effect on costs: differentiating between equilibrium effects where input price movement would offset varying input use, and individual behavior where changes in the input mix would leave total costs flat. First, we show a separate regression focusing on input prices, where prices are observed at the store-year-product level, and the regression includes store and product fixed effects in addition to year fixed effects. The regression shows a relatively precise zero effect of forecasted HDDs on input prices, ruling out offsetting price and quantity movements on the cost side. We give price responses per input category in table A12, confirming that none of the input prices respond to forecasts.

We then 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.[32] None of these inputs respond to forecasts, ruling out offsetting input-specific responses.[33]

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. To address this, we decompose land and labor, the two inputs for which we observe sub-categories, and analyze the

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 from their endowed weather which might not necessarily be the effective weather they end up facing.

[32]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 C.1 for detailed definitions of the different variables.

[33]We show in table A11 results for a larger set of inputs, including land and labor. None of land, labor, fertilizers, phytosanitary products, seeds, or irrigation respond to forecasted HDDs. Labor does respond to forecasted GDDs.

response of these sub-categories directly measured in levels rather than expenses (resp. hectares and hours). Decomposing land responses in figure A20, we see that the zero land effect 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 response.

When decomposing labor responses in figure A21, most categories remain unresponsive, while we note that the two most flexible types of labor respond negatively to expected hot weather. 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.

Further Robustness: We further examine the robustness of our results to alternative specifications, including adding lags of realized weather and removing region-specific time trends. We also vary the cutoffs used to define GDDs versus HDDs and introduce a measure of freezing degree days. Additionally, we test realized weather disaggregated to the municipality level, paired with our department-level forecasts. For this exercise, realized weather is computed using the full continuous range of hourly realizations rather than only four daily observations. Furthermore, we investigate heterogeneity in farm responses to forecasts. In all cases, the estimates remain consistent with our baseline findings. These robustness checks are reported in appendix E.

7 Dynamic Consequences of Ex Ante Actions

We have previously shown that farms experience a positive value of information from increased HDD forecasts, as these reduce the likelihood of costly mis-adaptation. We have further shown that this pattern is related to an underlying asymmetry of the static profit function in forecast errors. In this section, we provide an explanation for this asymmetry. We show that profit gains in the initial period are eroded over time by significant profit losses, which are driven by delayed increases in production costs. The net effect on the present discounted value of profit over four periods is statistically indistinguishable from zero. As such, the static profit function is asymmetric with respect to forecast errors because losses from mis-adaptation are immediately realized, while the costs of adaptation materialize only over time. We further explore the mechanisms giving rise to these dynamic cost increases.

7.1 Empirical Framework

To test for the dynamic effects of heat forecasts, 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 $t+3$. Our empirical design, based on local projections (Jordà, 2005; Jordà and Taylor, 2025), implements this strategy. The estimating equation is the same as equation (3), except that we now look at outcome y for farm j in year $t+h$, where $h \in [-2, 3]$, and include additional controls for past weather. Our estimating equation is:

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

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.

The addition of past weather controls is motivated by the new threats of endogeneity caused by the dynamic framework, specifically for the [t-2,t-1] results where w_{t-h} is both directly related to y_{t-h}, and can relate dynamically to $FHDD_t$. To isolate the causal effect of the forecast at time t on outcomes at time $t+h$ requires that the forecast is not correlated with other factors that determine outcomes in the surrounding years. 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.

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 A18 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.

7.2 Profits

We now present the main results of our dynamic analysis, focusing on the effect of farm responses to forecasted heat on the stream of profits over subsequent growing seasons. Figure 6 plots the dynamic response of profits to a forecasted HDD shock at time t (the table of results is available in table A19). The coefficients reveal a clear pattern of intertemporal trade-offs. We find a positive and significant contemporaneous profit effect of $+1.9k$ euros. This initial gain, however, is followed by statistically significant profit losses in the subsequent two periods, with effects of $-1.8k$ euros in year t+1 and $-1.4k$ euros in year $t+2$. The effect in year $t+3$ is small and not statistically significant.

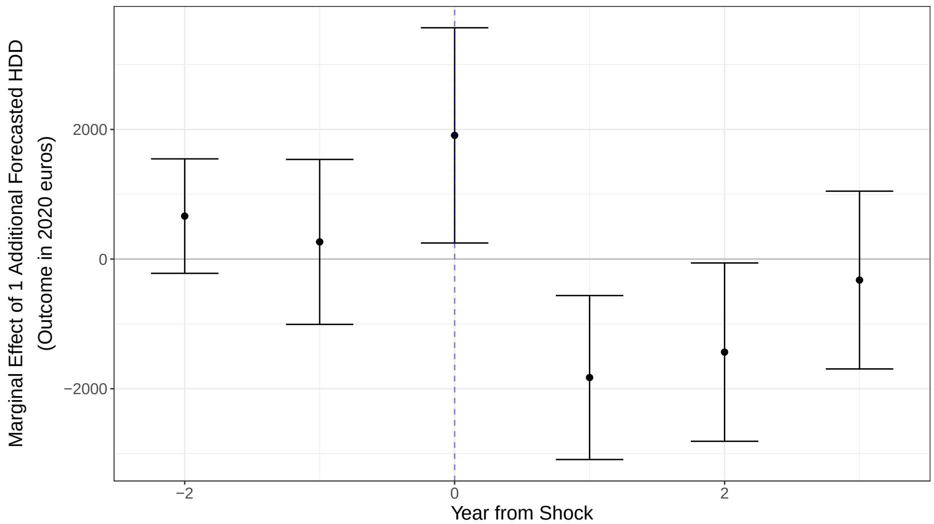

Figure 6: Dynamic Effects of Ex Ante Responses on Profit

Notes: Estimates display the dynamic response of farm profit to forecasted hot degree days at time t. The regressions are based on equation (5), 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. Standard errors are clustered two-way department-by-year, and confidence bands display the 95% confidence interval.

These estimates allow 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 ex ante responses is $-1.1k$ euros, 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. While we estimated a significant bias in the static framework, we cannot reject the null hypothesis that the dynamic revealed preference methods hold as, over time, the value of information vanishes.

7.3 Evidence for Costs in Subsequent Growing Seasons

We next examine the 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 A20.

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 (denoted by 0 on the x-axis). However, this is followed by a statistically significant cost increase of €$1k$ in year $t+1$, and an NPV of €1.1k. This pattern of delayed cost increases suggests that 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.

Consistent with the profit results, costs do arise but only in future periods rather than at the time of the shock. In our context, the net welfare consequences of a forecasted heat shock come not from direct, un-adapted damages, but almost entirely from the future costs of the actions taken to avoid those damages.

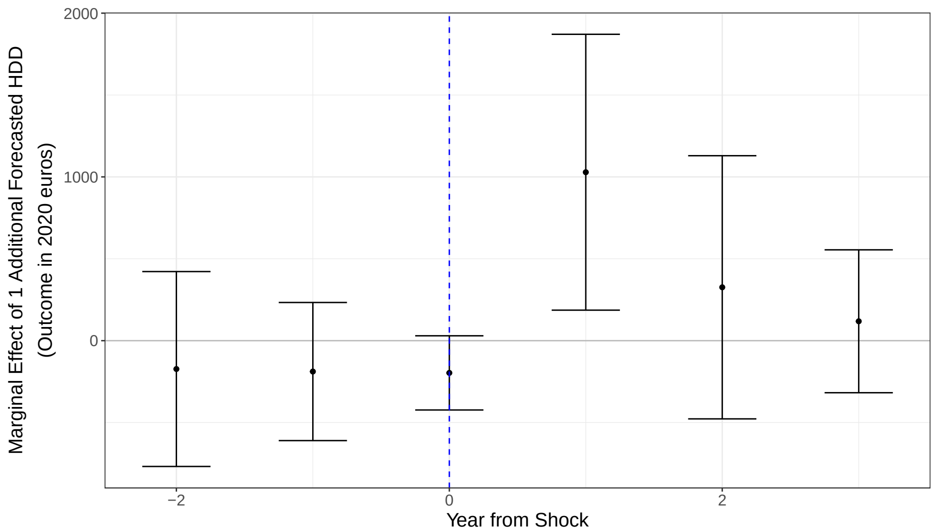

Figure 7: Dynamic Effects of Ex Ante Responses on Production Costs

Notes: Estimates display the dynamic response of farm costs to forecasted hot degree days at time t. The regressions are based on equation (5), 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. Standard errors are clustered two-way department-by-year, and confidence bands display the 95% confidence interval.

This interpretation—that crop switching incurs delayed costs—is well-supported by the literature on dynamic land use. Results from Livingston et al. (2008) and Scott (2013) show that deviating from an optimal crop rotation can have 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, corroborate 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. 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 $t+1$. These findings provide direct 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.

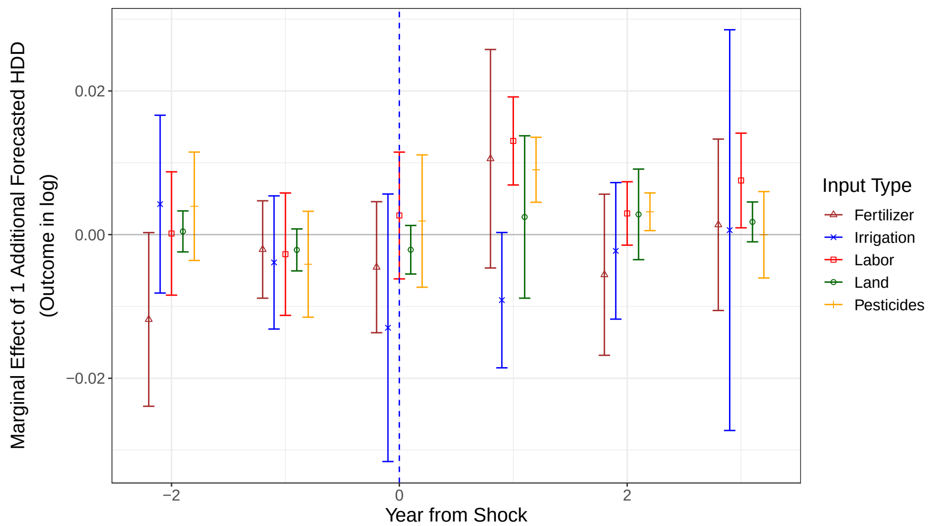

Figure 8: Dynamic Effects of Ex Ante Responses on (Log) Input Use

Notes: Estimates display the dynamic response of different (log) farm input bills to forecasted hot degree days at time t. The regressions are based on equation (5), 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. Standard errors are clustered two-way department-by-year, and confidence bands display the 95% confidence interval.

8 Implications for Climate Change

So far we have analyzed the response of farms to forecasts about one-off, seasonal variation in weather. We now focus on the implications these results have for the cost of adapting to repeated shocks from climate change.

Our first exercise assumes that actions in response to forecasts will scale with the frequency of those weather conditions. This exercise calculates the future costs of adaptation under the assumption that the way farmers respond to variation in heat, and the cost of that response, remains the same as it is today.

Of course, it is unlikely that adaptation costs, or the set of available adaptation actions will remain constant as climate changes. Our second exercise works to relax that assumption by using spatial variation in local climates as a proxy for temporal change in climate, following the Ricardian cross-sectional tradition (Mendelsohn et al., 1994; Auffhammer, 2018; Carleton et al., 2022).

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

In section 7, we estimated the average, farm-level cost of responding to a one-off, forecasted, one unit increase in HDD to be €1.1k. Given a baseline average of 2.4 HDDs per year in our sample period, if we assume that realized HDDs are well forecasted, this implies that current responses to extreme heat cost 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).[34] The distribution of these biases is shown in figure A15. 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.

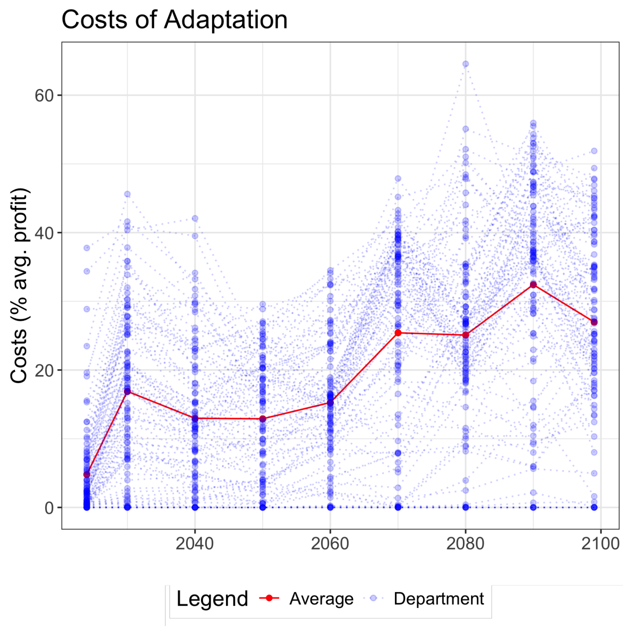

Figure 9: Costs of Adaptation Under Climate Change

Notes: The figure displays on the projected evolution of the costs of adaptation, using our estimate for the costs of responding to HDDs in figure 7. The projections are based on yearly HDD realizations in France according to the ECMWF CMIP6 model under the SSP3-7.0 scenario. The red line traces the average value, while the blue dots indicate department-specific realizations.

The average number of HDDs in France 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 says that under our assumption of a stable response to heat, climate-induced warming would cause a substantial financial burden on the French agricultural sector.

8.2 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

[34]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.

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—an example being 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.[35] 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 in 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 A21. Including these interactions has two consequences. The first one is that the coefficients associated with forecasted HDDs in levels become much larger in value for both the profit and the revenue regressions. Second, the coefficients associated with the interacted terms (average temperature) are negative and also statistically significant. These suggest that there are potentially decreasing returns to adaptation as the average temperature increases, 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.

9 Conclusion

Climate change is expected to have significant and wide-ranging impacts on human and non-human systems, from migration to human health to productivity across multiple sectors (see Hsiang, 2025; Lemoine et al., 2025, for recent reviews). Costs of adaptation are central to estimating these impacts: the effects of climate change combine the damages from direct exposure with the costs agents bear when adapting, and the relative size of these costs determines how much adaptation occurs and thus how exposed agents ultimately are. Yet these costs are rarely observed directly, and the standard approach infers them indirectly, bounding unobserved costs with observed benefits. We instead study the costs of ex ante responses to heat shocks in French agriculture, a setting whose centrality to climate policy and France's weight as an agricultural exporter make it of direct interest. We do so using precise, farm-level accounting data over 1994–2018 merged with realized and forecasted weather. Conditioning on realized weather, we use forecasts to identify the costs and benefits of farmers' anticipatory responses, recovering cost estimates directly rather than bounding them by benefits, and letting us examine the conditions under which the revealed-preference-based approaches common in climate economics deliver unbiased estimates of adaptation costs and direct weather effects.

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

Our central methodological finding is that those approaches can be biased when realized profit (resp. revenue) is used as a proxy for expected profit (resp. revenue), and that the bias has a clear economic interpretation: it is the value of information. When agents take costly actions in advance under imperfect forecasts, and when the profit function is asymmetric in forecast errors, realized-weather estimates of direct damages also absorb the within-season losses farmers incur when they mis-adapt to imperfectly forecasted heat, and therefore overstate true damages. We provide direct evidence for both ingredients in French agriculture. Farmers adapt ex ante by switching their crop mix toward more heat-resilient crops, an action with meaningful adjustment costs that must be taken before the season's weather is known. And the static profit function is asymmetric in forecast errors: under-preparing for heat is costly within the season, while over-preparing is not. Together these imply a positive value of information, which we estimate as the positive effect of heat forecasts on contemporaneous profit, a quantity that, under the usual static application of the envelope theorem, should be zero. A positive value of information corresponds to a negative bias in the realized-weather estimate, since the direct effect of heat is itself negative, this makes measured damage larger in magnitude, so standard methods overstate the harm from extreme heat.

The profit function asymmetry arises because the timing of costs and benefits differs. The profit gains from anticipating heat appear in the year of the shock, but the costs of the underlying crop-switching emerge only in subsequent seasons, through higher fertilizer, pesticide, and labor use consistent with disrupted crop rotations, so much so that over a four-year horizon the net present value is statistically indistinguishable from zero. Because losses from mis-adaptation are immediate while adaptation's costs are deferred, the value-of-information bias is a static phenomenon that vanishes dynamically. This timing also carries the headline cost estimate: negligible in the year of the shock but substantial over time, the cost of adapting to extreme heat amounts to roughly 3% of average annual profit today and, holding the response function fixed under CMIP6 projections, rises to on the order of 30% of current-day profit by century's end, magnitudes invisible to a static analysis, underscoring both the value of attending to dynamics and the practical relevance of a bias that leads within-season methods to overstate the damages from extreme heat.

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.

——— (2022): "Climate Adaptive Response Estimation: Short and long run impacts of climate change on residential electricity and natural gas consumption," *Journal of Environmental Economics and Management.*

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., A. Jina, E. M. Kelley, G. Lane, and H. Sahai (2025): "The Value of Forecasts: Experimental Evidence from India," *Working Paper.*

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

Butler, E. E. and P. Huybers (2013): "Adaptation of US maize to temperature variations," *Nature Climate Change.*

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.

Carter, C., X. Cui, D. Ghanem, and P. Mérel (2018): "Identifying the Economic Impacts of Climate Change on Agriculture," *Annual Review of Resource Economics*, 10:361-380.

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

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

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

Costinot, A., D. Donaldson, and C. Smith (2016): "Evolving Comparative Advantage and the Impact of Climate Change in Agricultural Markets: Evidence from 1.7 Million Fields around the World," *Journal of Political Economy*, 124, 205–248, publisher: University of Chicago Press-Chicago, IL.

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.

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

EPA (2023): "Supplementary Material for the Regulatory Impact Analysis for the Final Rulemaking, "Standards of Performance for New, Reconstructed, and Modified Sources and Emissions Guidelines for Existing Sources: Oil and Natural Gas Sector Climate Review"'; EPA Report on the Social Cost of Greenhouse Gases: Estimates Incorporating Recent Scientific Advances," Tech. rep., Environmental Protection Agency, docket ID No. EPA-HQ-OAR-2021-0317.

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*.

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

Hsiang, S. (2025): "The Global Economic Impact of Climate Change: An Empirical Perspective," Tech. Rep. w34357, National Bureau of Economic Research.

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. (2025): "Impacts of climate change on global agriculture accounting for adaptation," *Nature*, 642, 644–652.

Jordà, O. (2005): "Estimation and Inference of Impulse Responses by Local Projections," *American Economic Review*, 95, 161–182.

Jordà, O. and A. M. Taylor (2025): "Local Projections," *Journal of Economic Literature*, 63, 59–110.

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.*

Kleven, H. J. (2021): "Sufficient Statistics Revisited," *Annual Review of Economics*, 13, 515–538, publisher: Annual Reviews.

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

Lemoine, D., C. Hausman, and J. G. Shrader (2025): "A Guide to Climate Damages," Tech. Rep. w34348, National Bureau of Economic Research.

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.

Mérel, P., E. Paroissien, and M. Gammans (2024): "Sufficient statistics for climate change counterfactuals," *Journal of Environmental Economics and Management*, 124, 102940.

Michels, M., J. Mollman, and O. Musshoff (2019): "Understanding the Intention to Use Commodity Fugures Contracts," *Agricultural Finance Review.*

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

Obolenski, M. (2025): "Protect or Prepare? Crop Insurance and Adaptation in a Changing Climate," *Working Paper.*

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.*

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

Schelling, T. C. (1992): "Some Economics of Global Warming," *American Economic Review*, 82, 1–17.

Schlenker, W., 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 Review of Economics.*

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.*

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 Model Derivations

Derivation of equation (1): We characterize the wedge between the direct effect of weather on profit, and the object recovered from realized profit data when the ex ante action is chosen under an imperfect forecast.

We write the total derivative of the realized profit with respect to weather as:

$$\frac{d}{dw}\pi(x_1^\star(f), x_2^\star(w), w) = \frac{\partial \pi(x_1^\star(f), x_2^\star(w), w)}{\partial w} + \frac{\partial \pi}{\partial x_1}\underbrace{\frac{\partial x_1^\star(f)}{\partial w}}_{=0} + \underbrace{\frac{\partial \pi}{\partial x_2}}_{=0}\frac{\partial x_2^\star(w)}{\partial w} = \frac{\partial \pi(x_1^\star(f), x_2^\star(w), w)}{\partial w}$$

We obtain this expression because $x_2^\star(w)$ is ex post optimal, the envelope theorem implies $\frac{\partial \pi}{\partial x_2}\frac{\partial x_2}{\partial w}|_{x^\star} = 0$. Furthermore, because the ex ante action $x_1^\star(f)$ is chosen prior to the realization of weather, it cannot mechanically respond to the ex post weather shock, meaning $\frac{\partial \pi}{\partial x_1}\frac{\partial x_1}{\partial w}|_{x^\star} = 0$.

We take a first-order Taylor expansion of the marginal effect evaluated at the ex ante optimal action $x_1^\star(f)$, around the ex post optimal action $x_1^\star(w)$:

$$\frac{\partial \pi(x_1^\star(f), x_2^\star(w), w)}{\partial w} \approx \frac{\partial \pi(x_1^\star(w), x_2^\star(w), w)}{\partial w} + \frac{\partial^2 \pi}{\partial w \partial x_1}\Big(x_1^\star(f) - x_1^\star(w)\Big)$$

By definition of the ex post optimal action, the first-order condition is $\frac{\partial \pi}{\partial x_1}\Big(x_1^\star(w), x_2^\star(w), w\Big) = 0$. Differentiating this identity with respect to w yields $\frac{\partial^2 \pi}{\partial x_1^2}\frac{\partial x_1^\star(w)}{\partial w} + \frac{\partial^2 \pi}{\partial x_1 \partial w} = 0$, which implies $\frac{\partial^2 \pi}{\partial x_1 \partial w} = -\frac{\partial^2 \pi}{\partial x_1^2}\frac{\partial x_1}{\partial w}$.

Finally, we treat the optimal ex ante action as a policy function mapping information to inputs. We use a Taylor expansion of this policy function around the true weather w to express the difference in actions as a function of the forecast error:

$$\Big(x_1^\star(f) - x_1^\star(w)\Big) \approx \frac{\partial x_1^\star(w)}{\partial w}\Big(f - w\Big) = \frac{\partial x_1^\star(w)}{\partial w}\varepsilon.$$

This yields:

$$\frac{d}{dw}\pi(x_1^\star(f), x_2^\star(w), w) = \frac{\partial \pi(x_1^\star(f), x_2^\star(w), w)}{\partial w} \approx \frac{\partial \pi(x_1^\star(w), x_2^\star(w), w)}{\partial w} - \frac{\partial^2 \pi}{\partial x_1^2}\Big(\frac{\partial x_1^\star(w)}{\partial w}\Big)^2 \varepsilon$$

Derivation of equation (2): Next, we characterize the wedge between the realized marginal cost and benefits of ex ante actions. Using the expression for the profit function, we can write:

$$\Big(p\frac{\partial q(x_1^\star(f), x_2^\star(w), w)}{\partial x_1} - \frac{\partial c(x_1^\star(f), x_2^\star(w), w)}{\partial x_1}\Big)\frac{\partial x_1^\star(f)}{\partial f} = \frac{\partial \pi(x_1^\star(f), x_2^\star(w), w)}{\partial x_1}\frac{\partial x_1^\star(f)}{\partial f}$$

We can then use a Taylor expansion of $\frac{\partial\pi(x_1^\star(f),x_2^\star(w),w)}{\partial x_1}$ around $x_1^\star(w)$, which yields:

$$\frac{\partial\pi(x_1^\star(f),x_2^\star(w),w)}{\partial x_1} \approx \frac{\partial\pi(x_1^\star(w),x_2^\star(w),w)}{\partial x_1} + \frac{\partial^2\pi}{\partial x_1^2}\Big(x_1^\star(f) - x_1^\star(w)\Big)$$

By optimality of ex ante actions under perfect foresight, the first term disappears. Using the same previous derivation, we can then write that $x_1^\star(f) - x_1^\star(w) \approx \frac{x_1^\star(w)}{\partial w}\varepsilon$. We finally get:

$$\begin{aligned}\left(p\frac{\partial q(x_1^\star(f),x_2^\star(w),w)}{\partial x_1} - \frac{\partial c(x_1^\star(f),x_2^\star(w),w)}{\partial x_1}\right)\frac{\partial x_1^\star(f)}{\partial f} &= \frac{\partial\pi(x_1^\star(f),x_2^\star(w),w)}{\partial x_1}\frac{\partial x_1^\star(f)}{\partial f} \\ &= \left(\frac{\partial^2\pi}{\partial x_1^2}\frac{x_1^\star(w)}{\partial w}\varepsilon\right)\frac{\partial x_1^\star(f)}{\partial f}\end{aligned}$$

B Additional Figures and Tables

B.1 Figures

B.1.1 Precision of Forecasts

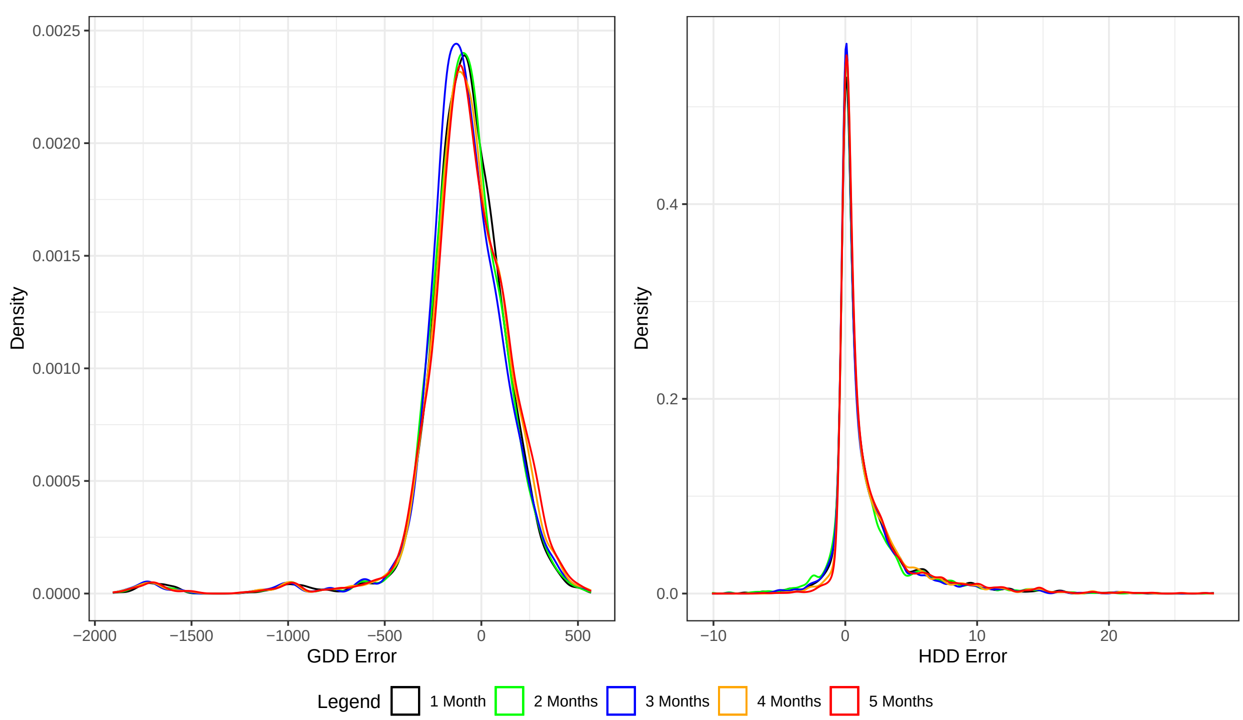

Figure A1: Distribution of Temperature Forecast Error

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.

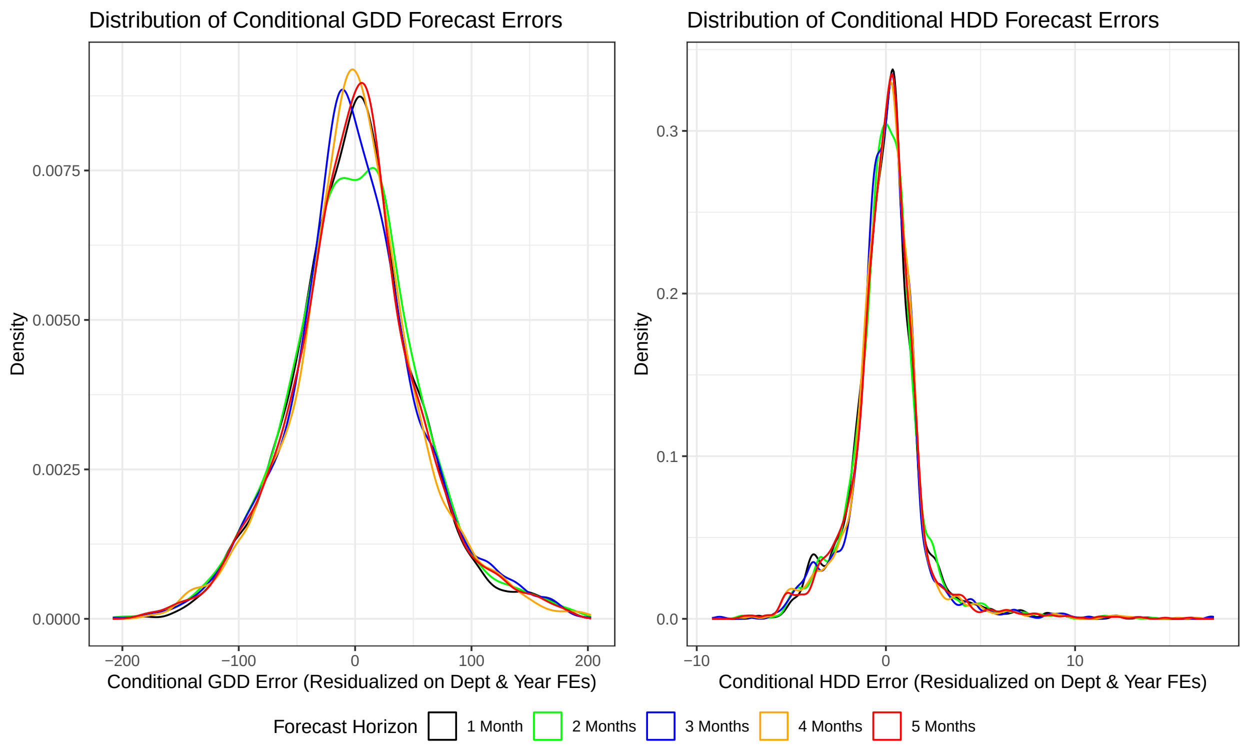

Figure A2: Conditional Distribution of Temperature Forecast Error

Notes: These estimated kernel densities show the empirical distribution of the forecast errors (realized minus forecasted) conditional on department and year fixed effects, 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. The inclusion of fixed effects matches the regression design used throughout the paper, and shows that conditional on them, forecasts can be taken as unbiased, symmetrically distributed around the realizations.

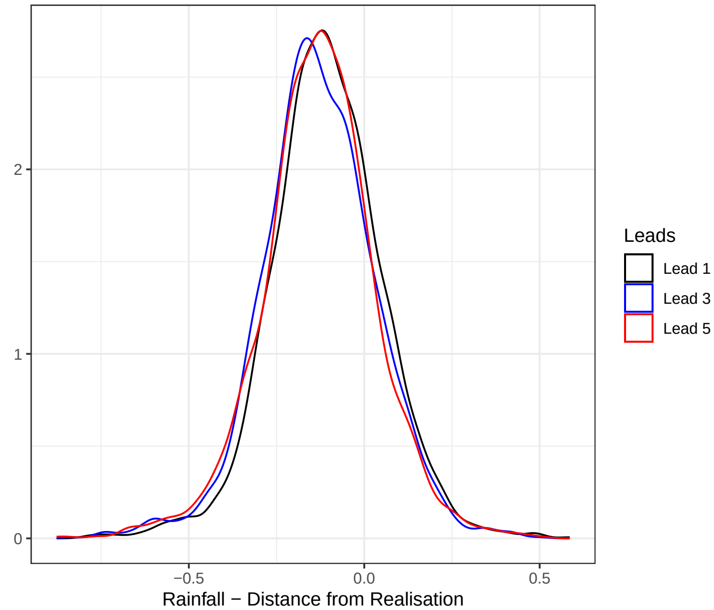

Figure A3: Distribution of Rainfall Forecast Error

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.

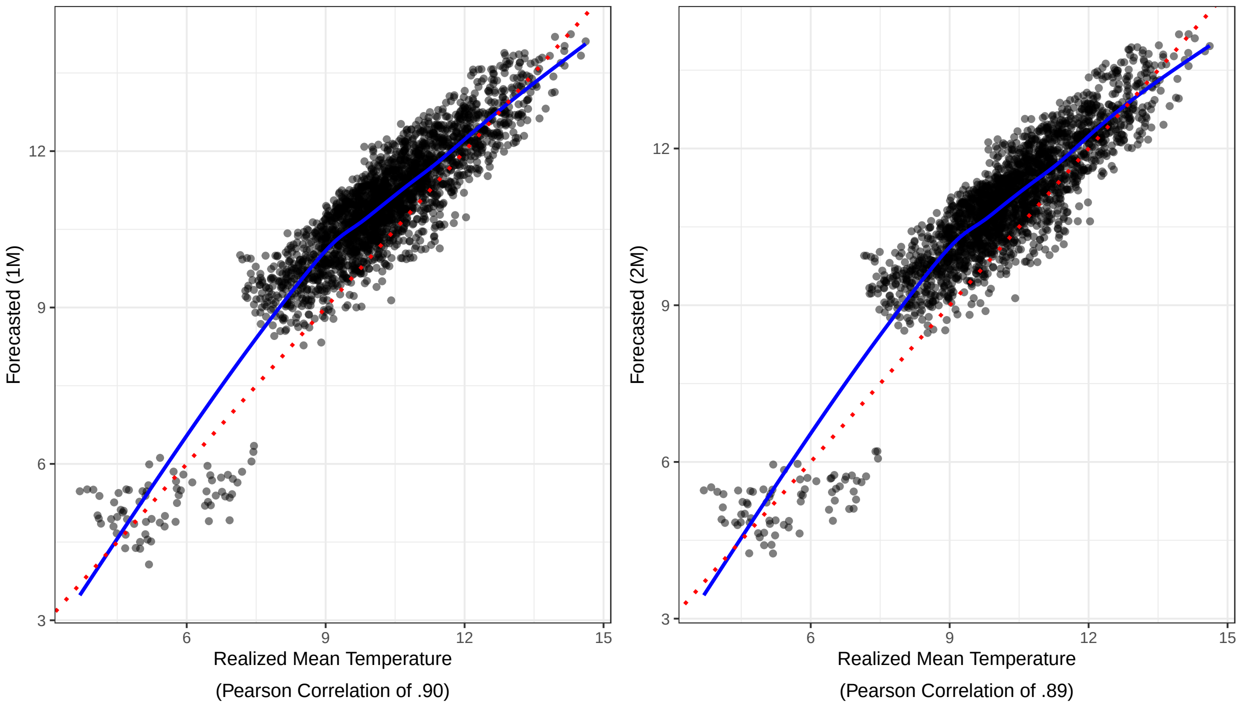

Figure A4: Correlation Realized and Forecasted Mean Temperature

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.

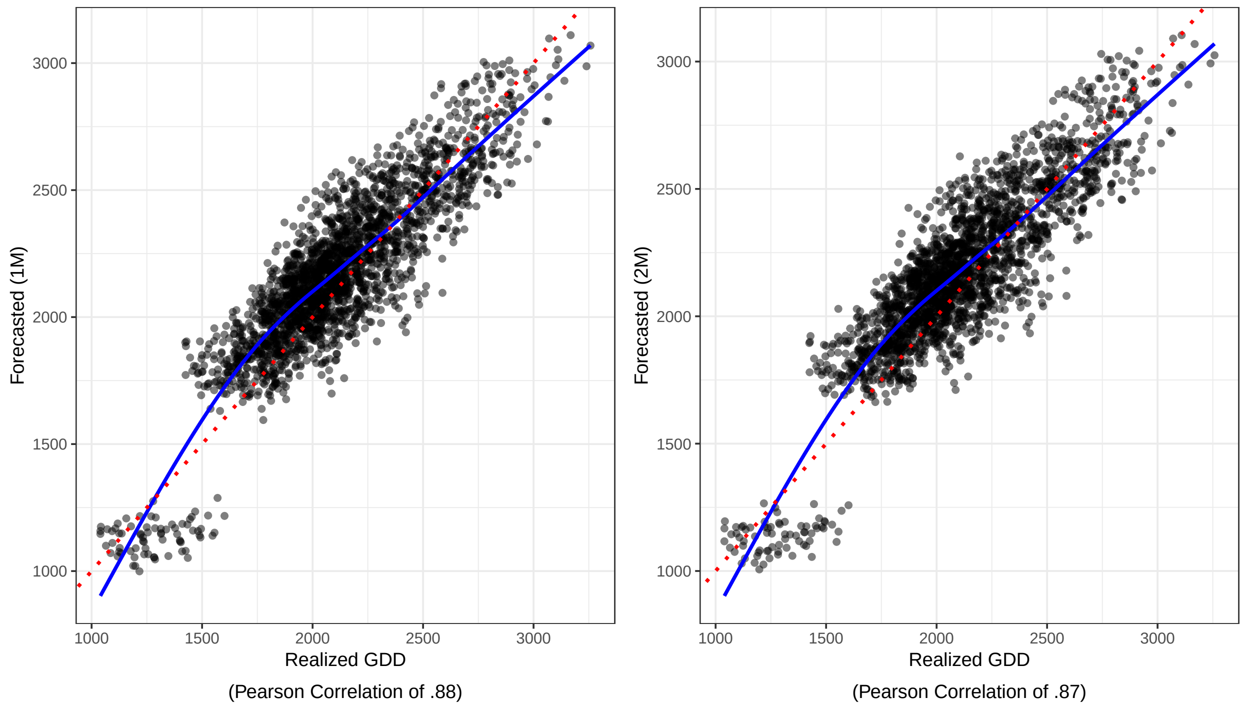

Figure A5: Correlation Realized and Forecasted GDDs

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.

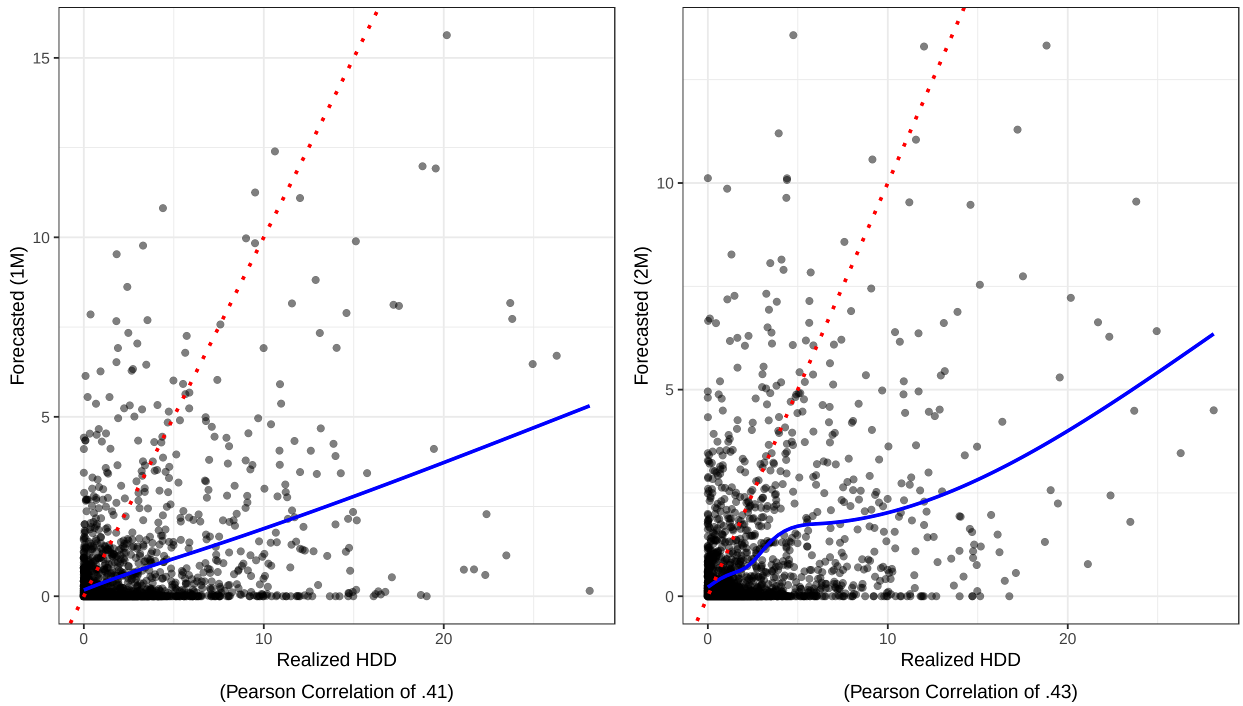

Figure A6: Correlation Realized and Forecasted HDDs

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.

B.1.2 Climate in France

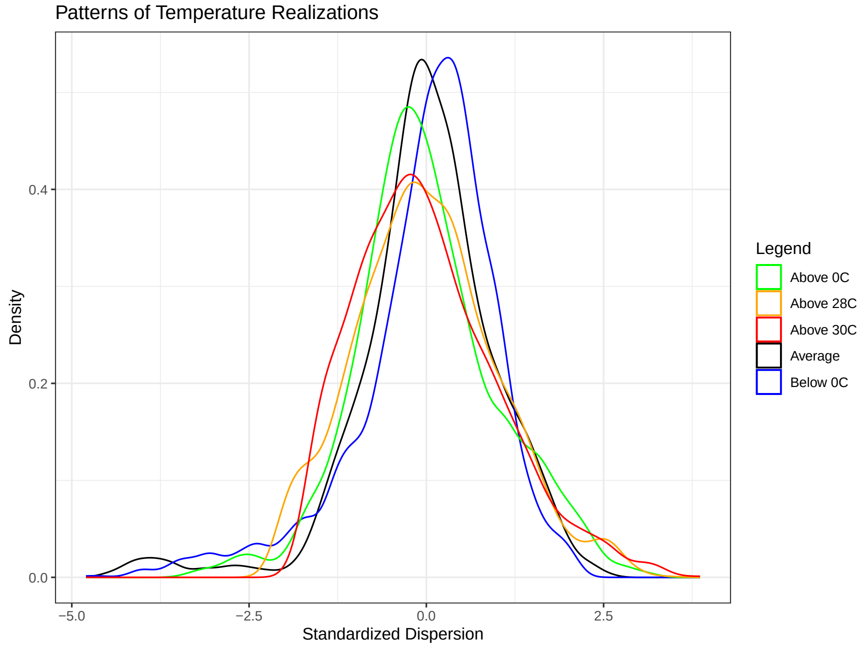

Figure A7: Distribution of the Standardized Conditional and Unconditional Mean Temperature

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

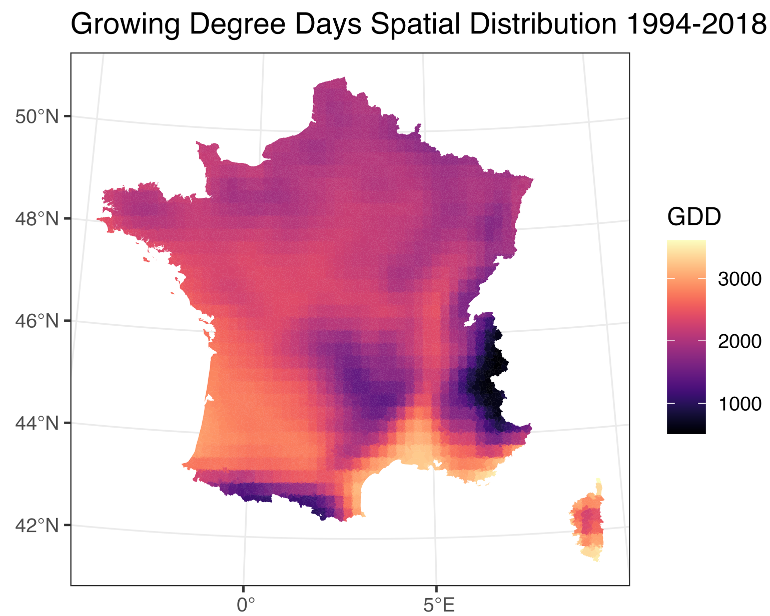

Figure A8: Distribution of Growing Degree Days

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).

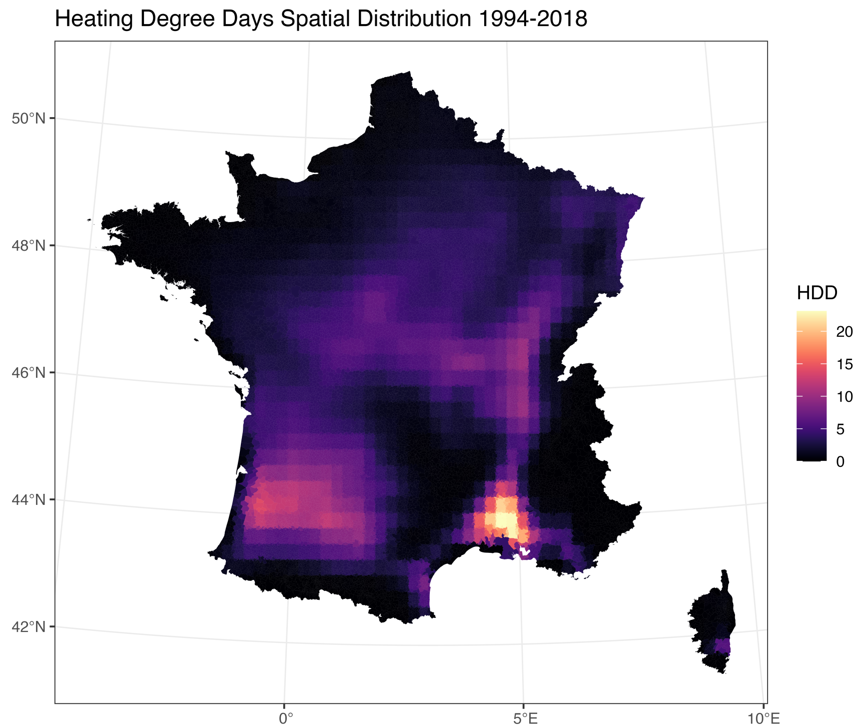

Figure A9: Distribution of Heating Degree Days

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.

B.1.3 Growing Seasons

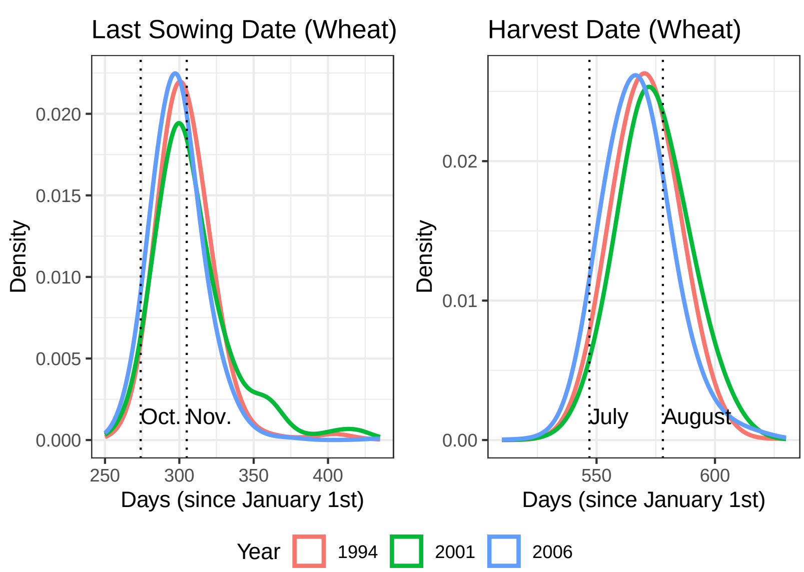

Figure A10: Wheat Growing Season in France

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.

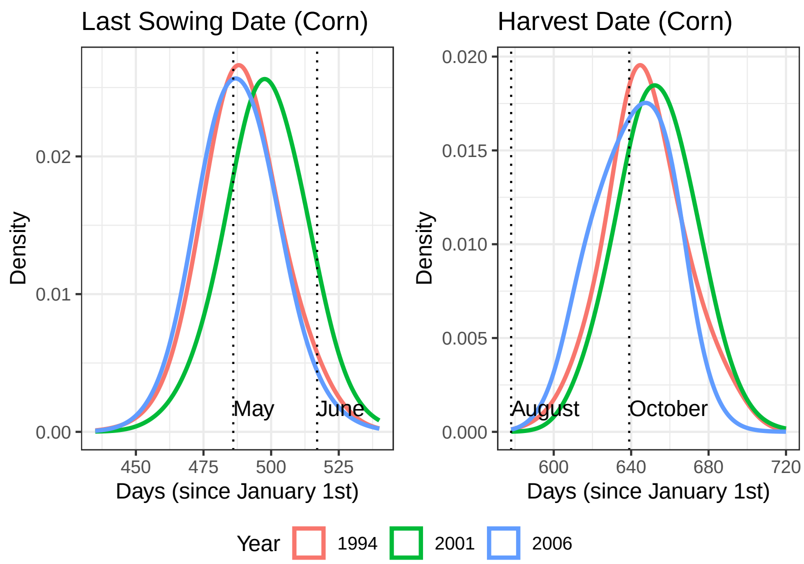

Figure A11: Corn Growing Season in France

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.

B.1.4 Yield and Output Effects of Weather Variation

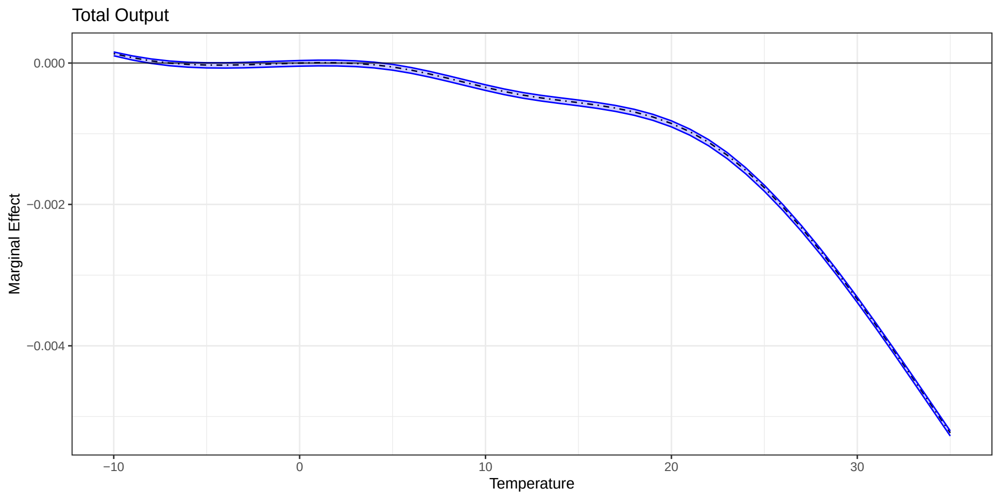

Figure A12: Temperature Effects on Farm Total Output

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

B.1.5 Forecasted Heat & Other Signals

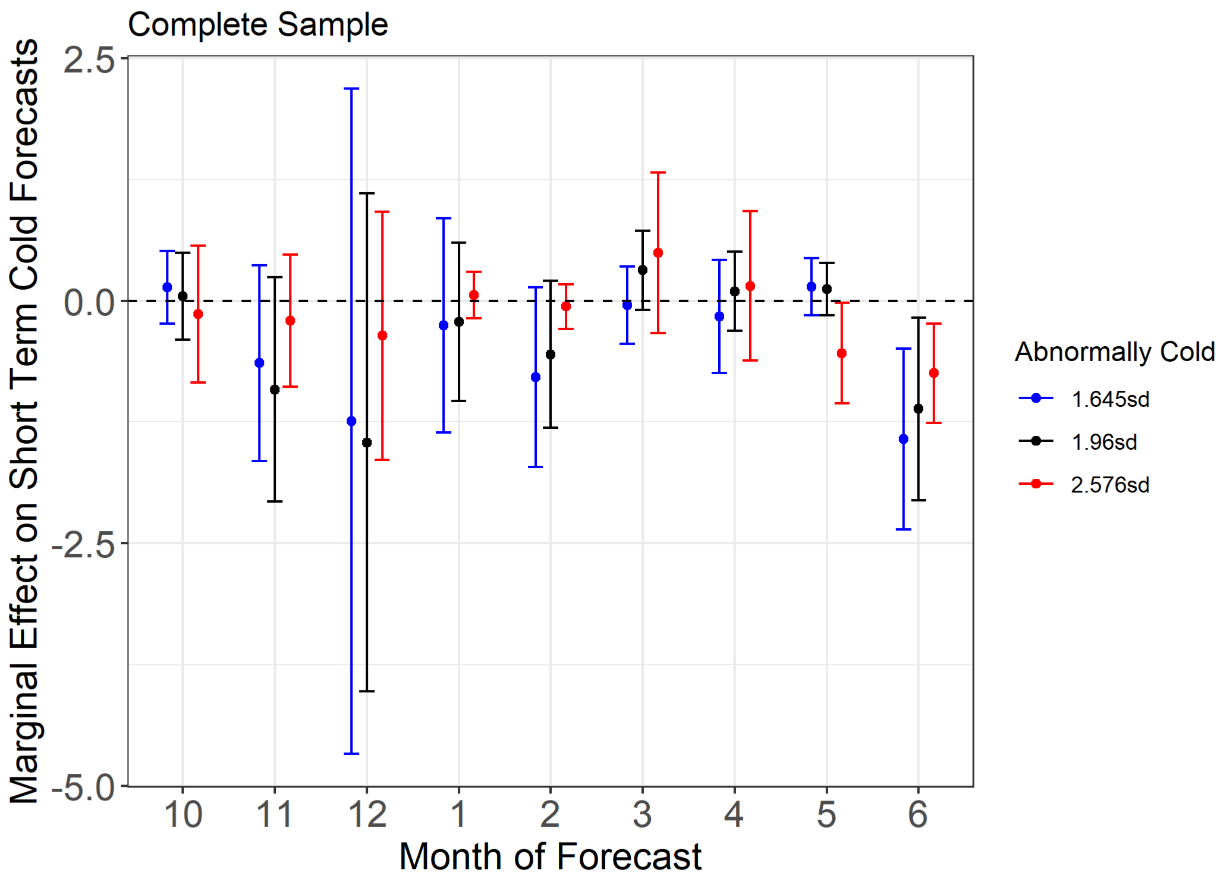

Figure A13: Relation Between Growing Season Forecasted HDD & Short Term Cold Forecasts

Notes: This figure shows effect of a growing-season forecast of HDDs on month-specific forecasts of days which are colder than usual. These short term forecasts correspond to the forecast received by farmers on the first day of the month, for the rest of the month. Outcome variables are cold days, corresponding to the integral of cold weather realizations below a threshold corresponding resp. to 1.96 standard deviations from the historical department-month-day-hour average, 1.645 deviations and 2.576 deviations. In this sense, the forecasts correspond to predicted temperature significantly colder than the average. Forecasts that are 2.576 deviations away from the average hence correspond to the coldest realizations.

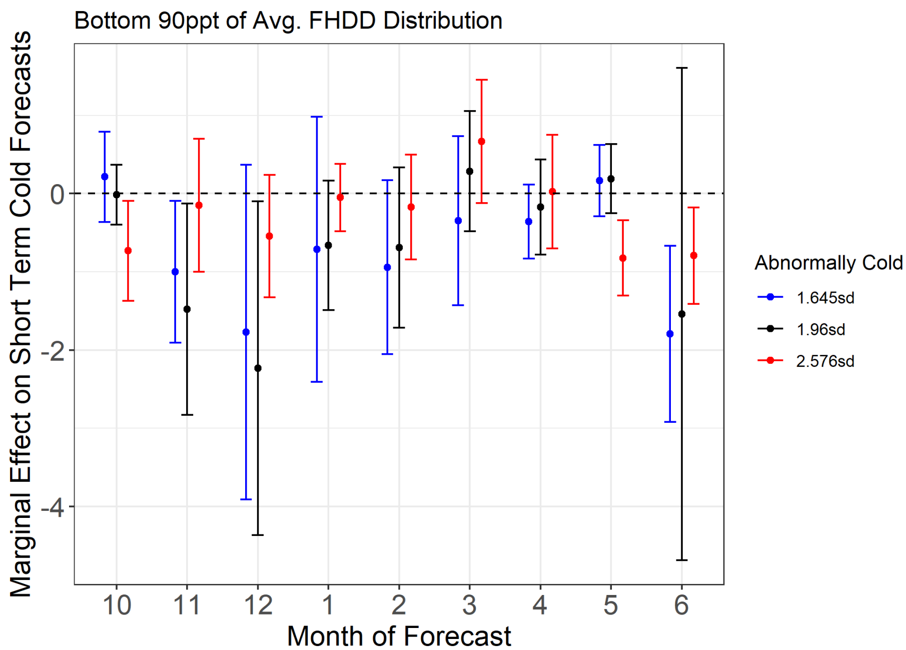

(a) Results for Departments with Small FHDD Values

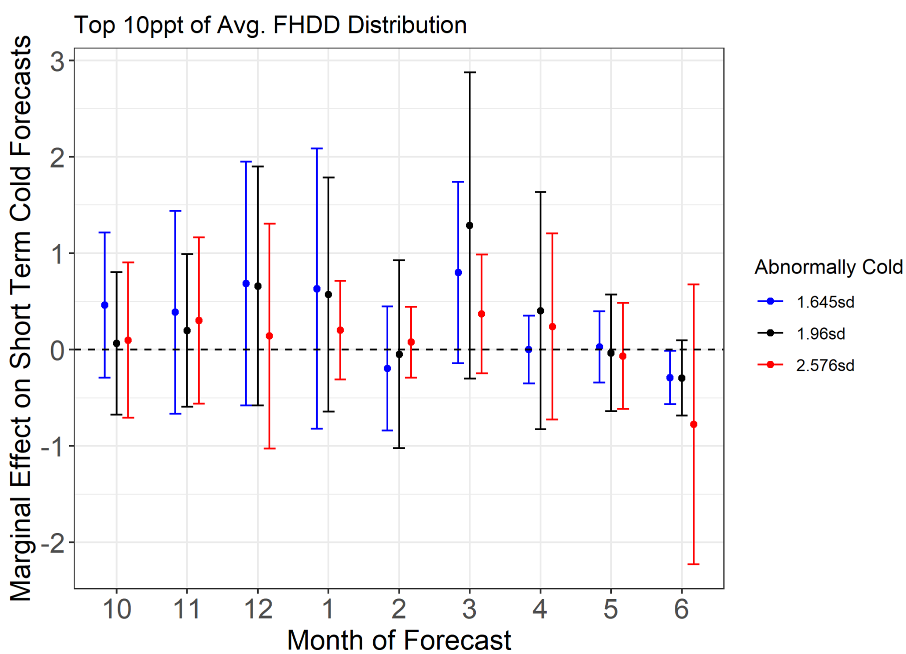

(b) Results for Departments with Large FHDD Values

Figure A14: Relation Between Growing Season Forecasted HDD & Short Term Cold Forecasts (Heterogeneity)

Notes: This figure shows effect of a growing-season forecast of HDDs on month-specific forecasts of days which are colder than usual. These short term forecasts correspond to the forecast received by farmers on the first day of the month, for the rest of the month. Outcome variables are cold days, corresponding to the integral of cold weather realizations below a threshold corresponding resp. to 1.96 standard deviations from the historical department-month-day-hour average, 1.645 deviations and 2.576 deviations. In this sense, the forecasts correspond to predicted temperature significantly colder than the average. Forecasts that are 2.576 deviations away from the average hence correspond to the coldest realizations. The top figure corresponds to the results on the part of our sample where FHDDs are less frequent and smaller, while the second corresponds to the tail of the sample where they are large and more frequent. The first sample corresponds to the bottom 90 percent of the distribution of average department-level FHDD values, and the second to the top 10 percent of that distribution.

B.1.6 Bias of Climate Projections

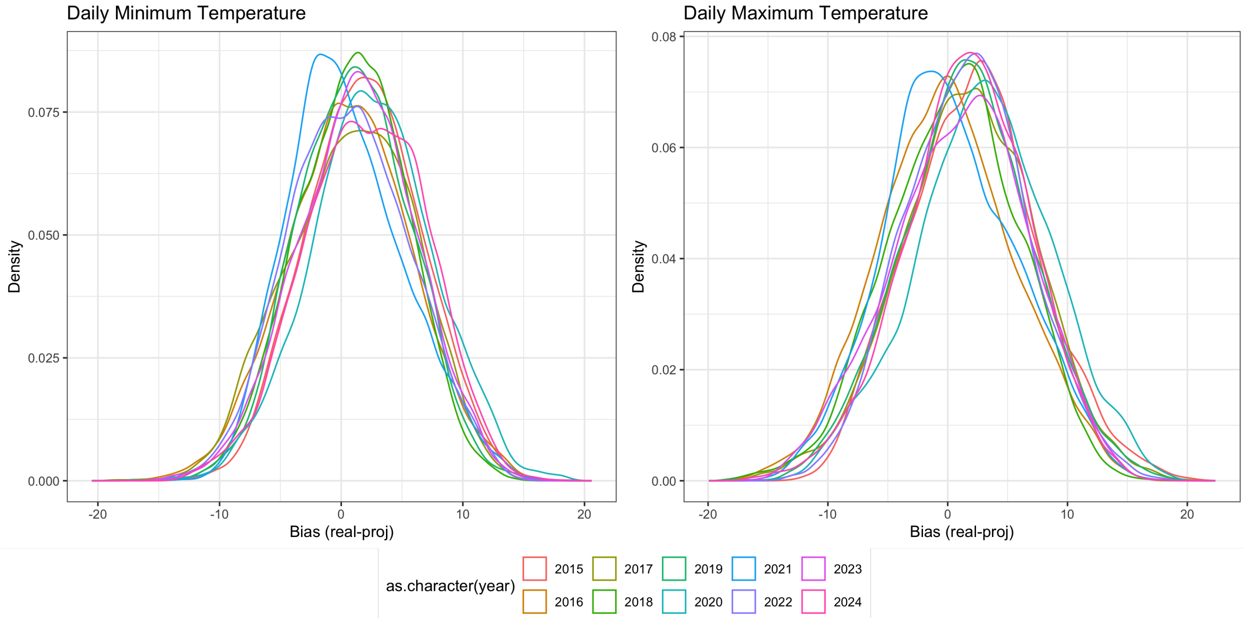

Figure A15: Distribution of Biases between CMIP6 and ERA5

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.

B.1.7 Multi-Product Farms

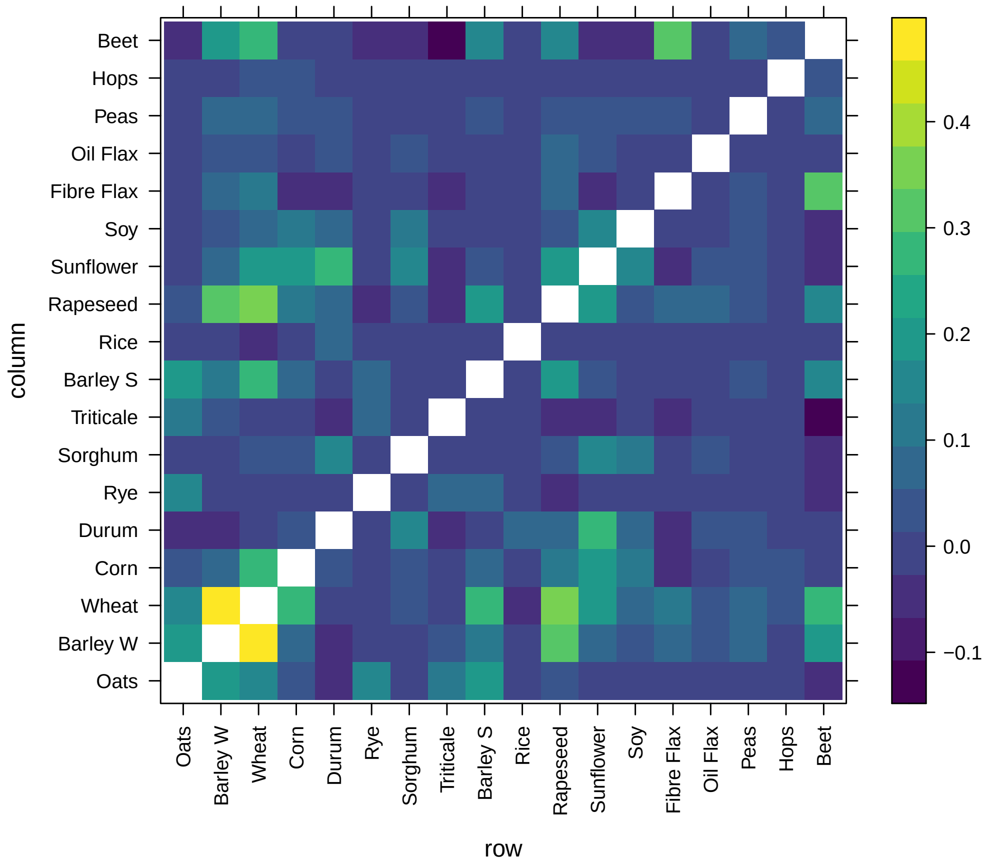

Figure A16: Conditional Probabilities to Grow Crop Pairs

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

B.1.8 Input Price Indices

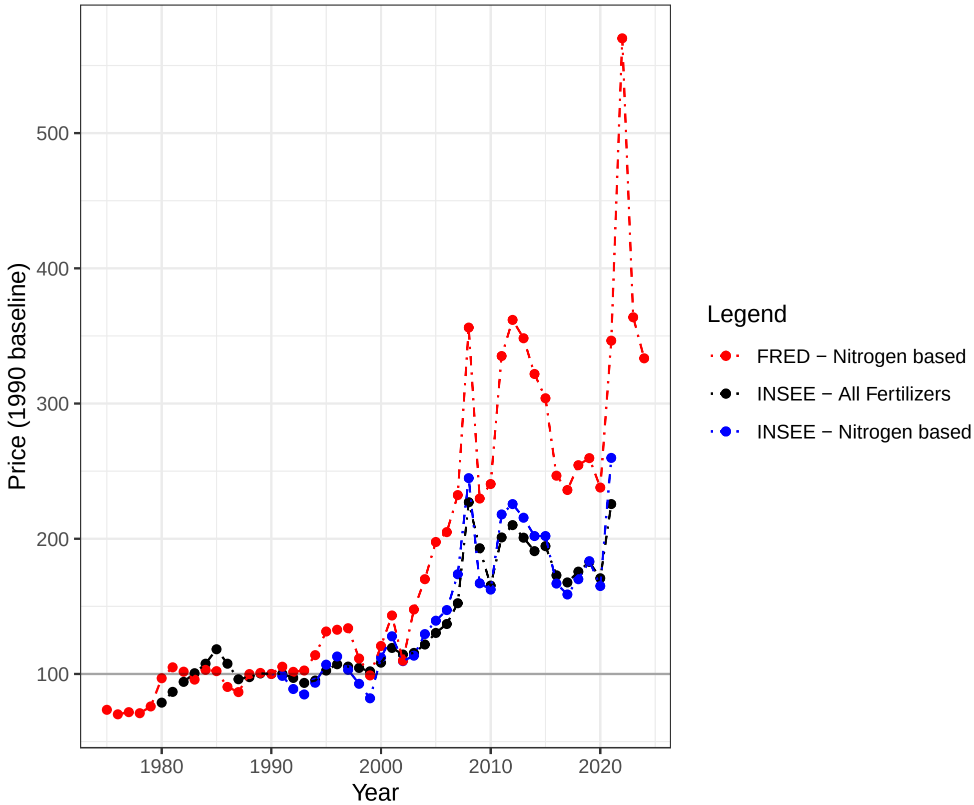

Figure A17: Comparison of Input Price Indices Series

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.

B.1.9 Additional Results

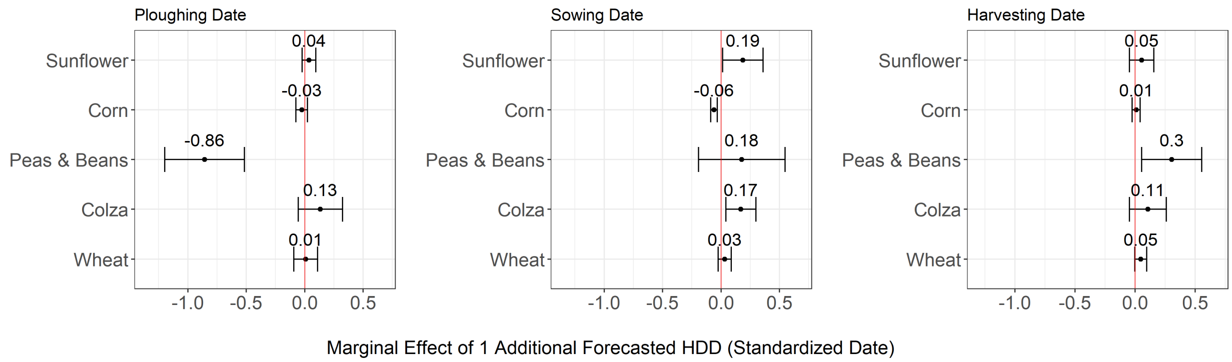

Figure A18: Crop-Specific Timing Responses

Notes: Estimates display the response to one month ahead forecasted heating degree days of crop-specific timing decisions at the plot level, using the PKGC repeated cross-sectional data. The regressions are based on equation (3), but replace farm fixed effects with department level ones. 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 PKGC. Standard errors are clustered at the department level to account for the very small number of time periods, and confidence bands display the 95% confidence interval.

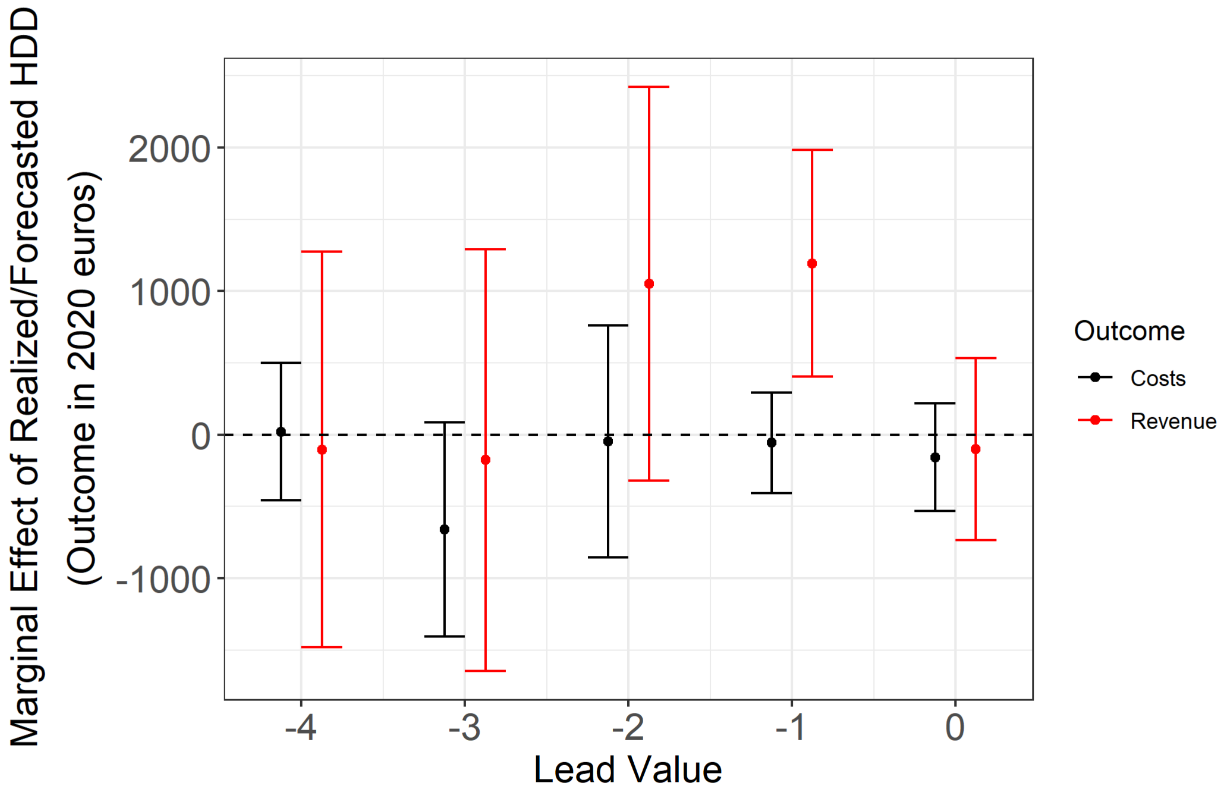

Figure A19: Effect of Forecasted HDDs

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

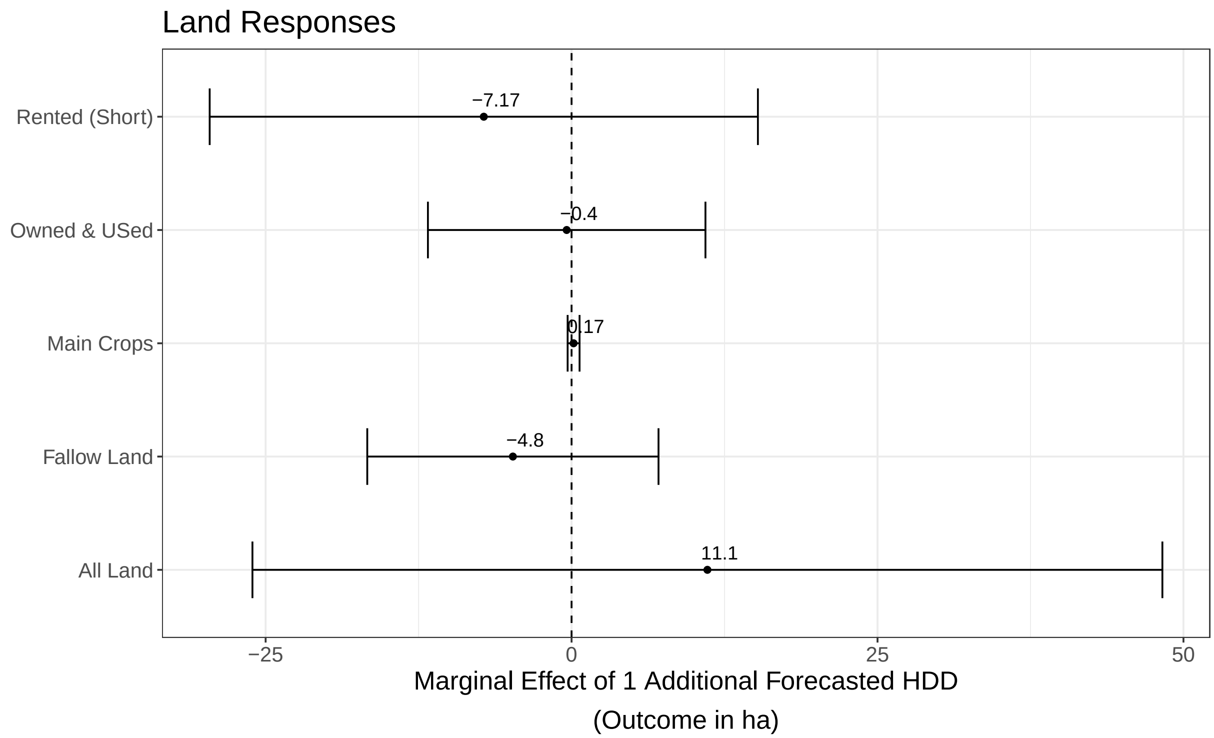

Figure A20: Decomposition of the Land responses: Ownership and Use

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.

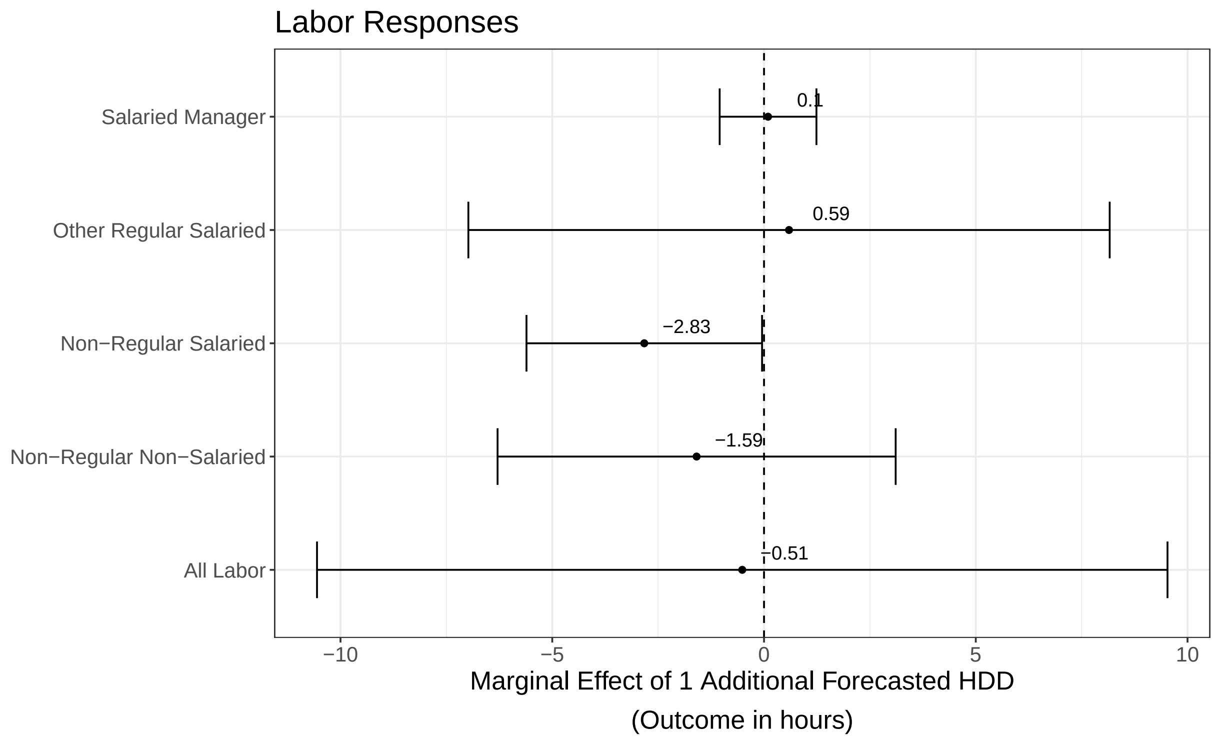

Figure A21: Decomposition of Labor Responses

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.

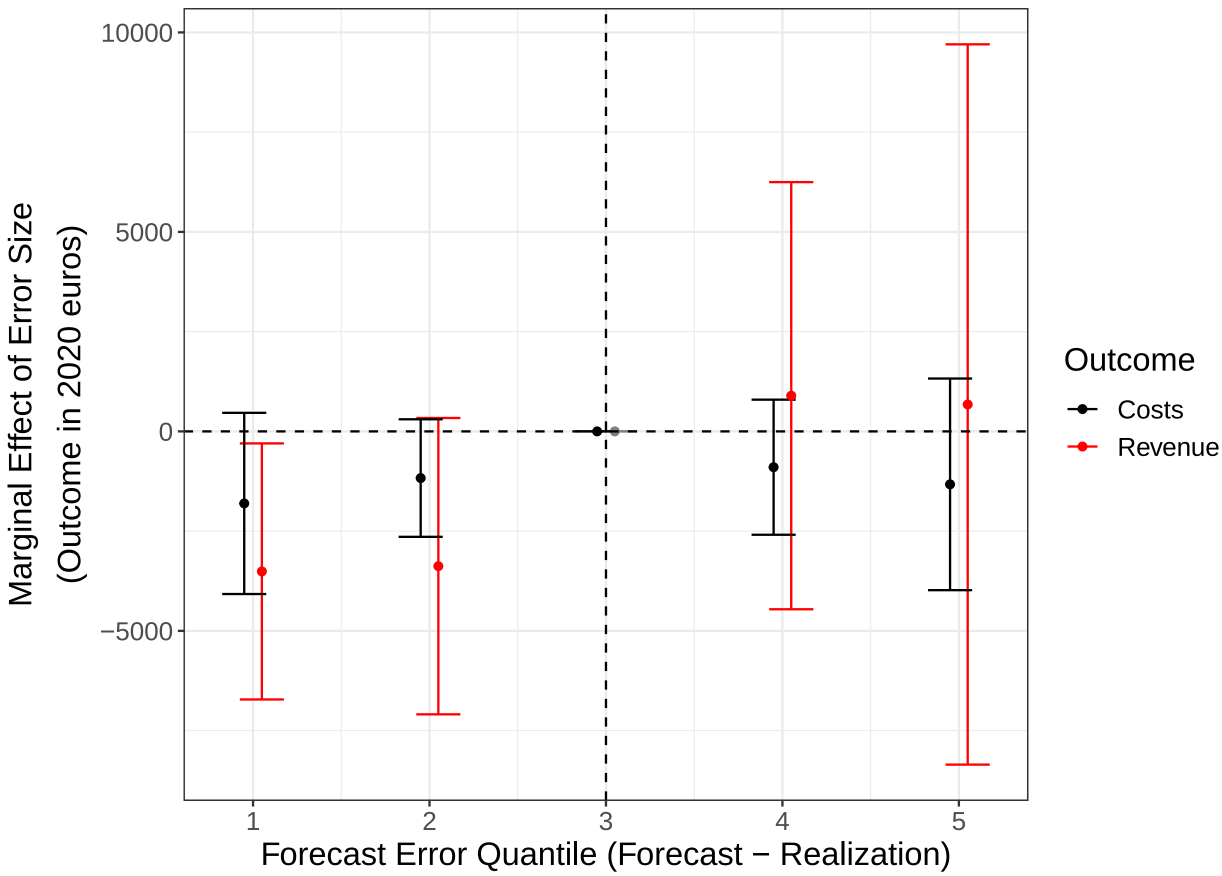

Figure A22: Cost and Revenue & Sign of the Forecast Error

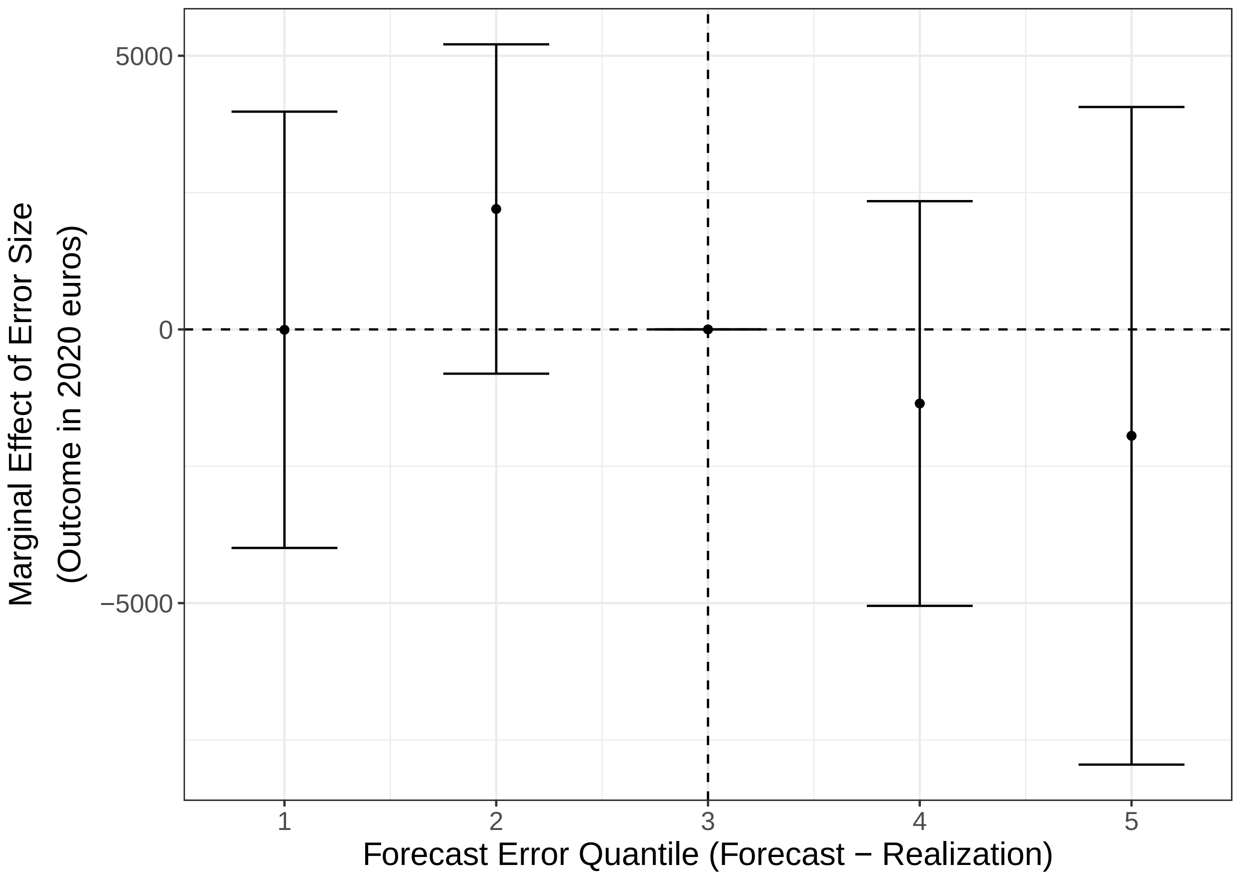

Figure A23: Profit & Sign of the GDD Forecast Error

B.2 Tables

Table A2: Descriptive Statistics - Farm-Level Dataset

Statistic	N	Mean	St. Dev.	Pctl(25)	Pctl(75)
Rainfall	18,917	0.75	0.17	0.63	0.85
Rainfall (F1)	18,917	0.85	0.13	0.76	0.92
GDD	18,917	2,140	264	1,957	2,292
GDD (F1)	18,917	2,182	238	2,022	2,305
HDD	18,917	2.36	3.2	0.24	3.13
HDD (F1)	18,917	0.44	1.12	0	0.4
Sales (€)	18,917	155,386	113,740	78,739	199,717
Total Costs (€)	18,917	123,249	85,427	66,903	156,468
Intermediate Inputs (€)	18,917	103,749	67,149	59,074	132,589
Value Added (€)	18,917	58,357	62,858	18,306	80,436
Profit (€)	18,917	86,695	70,268	40,066	114,771
Price Index (€/t)	18,917	263	296.4	140.5	230.6
Storage (sum, 0.1 t)	18,917	88.7	1,877	-280	500
Storage (index, 0.1 t)	16,075	65.4	983.7	-67.9	176
Output (sum, 0.1 t)	18,917	7,793	5,033	4,252	10,170
Output (index, 0.1 t)	18,917	2,401	2,926	769	2,847
Production (corn, 0.1 t)	10,346	3,464	4,046	877	4,682
Price (corn) (€/t)	9,958	150	39.1	122.9	172.8
Sales (corn) (€)	10,346	53,508	69,739	11,163	68,978
Quantity Sold (corn, 0.1 t)	10,346	3,437	4,133	817	4,617
Production (wheat, 0.1 t)	17,336	3,840	3,005	1,665	5,248
Price (wheat) (€/t)	17,181	163.1	37.95	136.4	184
Sales (wheat) (€)	17,336	61,538	53,986	24,345	82,922
Quantity Sold (wheat, 0.1 t)	17,336	3,799	3,162	1,560	5,166
Irrigation (€)	18,917	683.1	2,665	0	0
Labor (€)	18,917	2,684	1,330	1,600	3,200
Phytosanitary (€)	18,917	23,072	16,499	11,388	30,872
Fertilizer (€)	18,917	31,777	20,727	17,456	41,061
Land (ha)	18,917	145.04	86.19	8.379	18.7
Seeds (€)	18,917	12,231	9,686	5,847	15,882

Table A3: Descriptive Statistics - Store-Level Dataset

Statistic	N	Mean	St. Dev.	Pctl(25)	Pctl(75)
Pesticide Prices	2,126	615.5	718.3	258	697
Fertilizer Prices	3,098	3,254	2,959	2,450	3,783
Seed Prices	1,831	343.7	1,777	95.7	171.8

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.	Pctl(25)	Pctl(75)
Land Prices	3,405	5,754	3,085	3,717	6,834

Table A6: Descriptive Statistics - Weather Outcomes within Samples

Statistic	N	Mean	St. Dev.	Pctl(25)	Pctl(75)
Panel A	Entire Sample				
Rainfall	18,926	0.76	0.17	0.63	0.85
Rainfall (F1)	18,926	0.85	0.13	0.76	0.92
GDD	18,926	2,140	264	1,957	2,292
GDD (F1)	18,926	2,182	237.7	2,022	2,306
HDD	18,926	2.361	3.2	0.24	3.13
HDD (F1)	18,926	0.44	1.12	0	0.4
Panel B	Worst Region-Years				
Rainfall	1,425	0.68	0.14	0.56	0.75
Rainfall (F1)	1,425	0.88	0.11	0.79	0.95
GDD	1,425	2,418	255.7	2,196	2,587
GDD (F1)	1,425	2,356	276.1	2,136	2,548
HDD	1,425	10.4	3.51	7.35	12
HDD (F1)	1,425	1.41	2.07	0.02	2.1

Table A7: 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 A8: 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 (department & year) 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 A9: 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 A10: 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 A11: 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 A12: 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 A13: 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 A14: Timing Response to Forecasts: Ploughing

Dependent Variable:	Standardized Ploughing Date				
Model:	(1)	(2)	(3)	(4)	(5)
Variables					
GDD	-0.0012	0.0010	-0.0053	-0.0021	-0.0149**
	(0.0018)	(0.0010)	(0.0040)	(0.0022)	(0.0057)
GDD (F)	0.0010	-0.0004	-0.0155**	-0.0048	-0.0136*
	(0.0030)	(0.0016)	(0.0062)	(0.0053)	(0.0071)
HDD	-0.0046	-0.0025	0.0344	-0.0714**	0.2189*
	(0.0175)	(0.0135)	(0.0249)	(0.0335)	(0.1149)
HDD (F)	0.0085	-0.0253	0.1344	0.0372	-0.8561***
	(0.0520)	(0.0254)	(0.0968)	(0.0295)	(0.1743)
Crop	Wheat	Corn	Colza	Sunflower	Peas & Beans
Mean Date	290.7	420.3	234.4	368.0	365.4
Fixed-effects					
Department	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes
Fit statistics					
Observations	7,313	8,562	3,069	3,016	2,015
R^2	0.37	0.46	0.22	0.28	0.34

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

Notes: Outcomes are standardized at the crop level, to make comparisons of effects across crops more straightforward. The mean ploughing dates are also shown in the table. These dates run from 0 (the first day of the calendar year) to 730, the last day of the following year. As such ploughing for wheat happens in the fall, and for corn in the winter. One-way department 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: Timing Response to Forecasts: Sowing

Dependent Variable:	Standardized Sowing Date				
Model:	(1)	(2)	(3)	(4)	(5)
Variables					
GDD	-0.0017	0.0020	-0.0051	0.0066	0.0104***
	(0.0015)	(0.0014)	(0.0036)	(0.0044)	(0.0035)
GDD (F)	0.0047**	-0.0048**	-0.0126**	-0.0095	0.0348***
	(0.0023)	(0.0019)	(0.0047)	(0.0093)	(0.0124)
HDD	-0.0109	0.0194	0.0131	-0.2667***	-0.1907
	(0.0172)	(0.0181)	(0.0246)	(0.0813)	(0.1434)
HDD (F)	0.0317	-0.0599***	0.1703**	0.1861**	0.1777
	(0.0281)	(0.0142)	(0.0653)	(0.0880)	(0.1885)
Crop	Wheat	Corn	Colza	Sunflower	Peas & Beans
Mean Date	301.5	490.5	247.8	484.1	449.8
Fixed-effects					
Department	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes
Fit statistics					
Observations	9,987	9,307	4,436	3,430	2,143
R^2	0.25	0.27	0.16	0.23	0.31

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

Notes: Outcomes are standardized at the crop level, to make comparisons of effects across crops more straightforward. The mean sowing dates are also shown in the table. These dates run from 0 (the first day of the calendar year) to 730, the last day of the following year. One-way department 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: Timing Response to Forecasts: Harvesting

Dependent Variable:	Standardized Harvesting Date				
Model:	(1)	(2)	(3)	(4)	(5)
Variables					
GDD	-0.0005	-0.0004	-0.0009	-0.0036	0.0018
	(0.0014)	(0.0013)	(0.0031)	(0.0024)	(0.0076)
GDD (F)	0.0044*	-0.0046*	-0.0017	0.0079	0.0195
	(0.0025)	(0.0024)	(0.0047)	(0.0069)	(0.0118)
HDD	-0.0069	-0.0179	-0.0029	-0.0542	-0.0568
	(0.0143)	(0.0200)	(0.0326)	(0.0492)	(0.0690)
HDD (F)	0.0470*	0.0083	0.1063	0.0532	0.3040**
	(0.0255)	(0.0168)	(0.0784)	(0.0519)	(0.1277)
Crop	Wheat	Corn	Colza	Sunflower	Peas & Beans
Mean Date	569.2	648.7	561.0	627.8	574.4
Fixed-effects					
Department	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes
Fit statistics					
Observations	9,974	9,186	4,419	3,409	2,138
R^2	0.54	0.22	0.38	0.23	0.49

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

Notes: Outcomes are standardized at the crop level, to make comparisons of effects across crops more straightforward. The mean harvesting dates are also shown in the table. These dates run from 0 (the first day of the calendar year) to 730, the last day of the following year. One-way department 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: Cost and Revenue Reactions to Alternative Weather

Dependent Variables:	Revenue	Costs
Model:	(1)	(2)
Variables		
GDD	2.73	-1.26
	(10.6)	(8)
GDD (F)	21.3	35.9**
	(37.1)	(14.8)
HDD	-534	-106
	(178)	(90.8)
HDD (F)	628**	-13.6
	(261)	(128)
FDD	22.7	-2.68
	(30)	(12.4)
FDD (F)	-23.4	-96
	(107)	(84.3)
Mean	155,638	123,304
Unique Farms	2,625	2,625
Fixed-effects		
Farm	Yes	Yes
Year	Yes	Yes
Fit statistics		
Observations	18,428	18,428
R^2	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 A18: 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 A19: 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 A20: 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 A21: 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.

C Data Details

C.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.** 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 is 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.

C.2 Weather Data

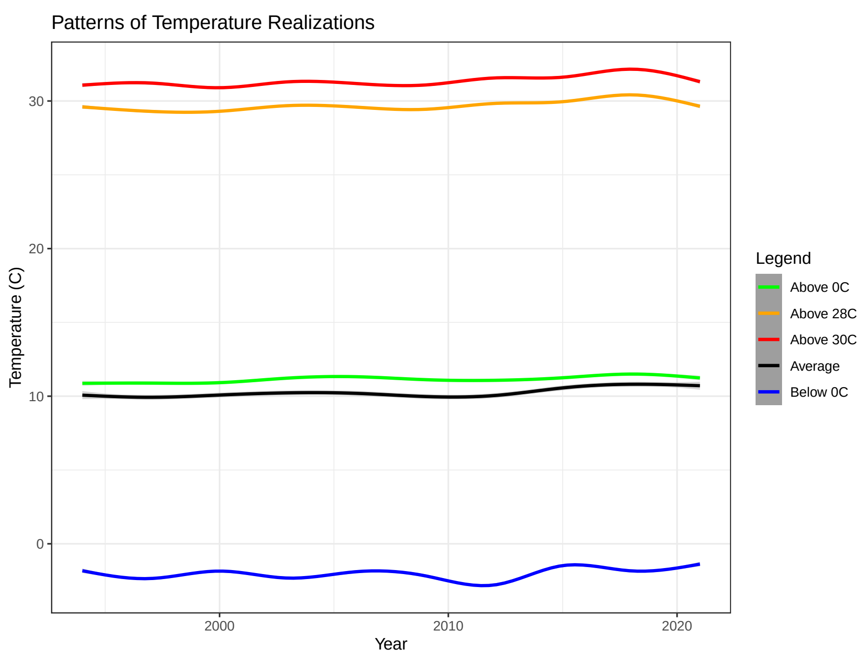

Figure A24: Unconditional and Conditional Average Temperature Realizations

D Validation of Results

Table A22: Corn-Specific Outcomes

Dependent Variables:	Output	Price	Revenue	Quantity Sold
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 A23: 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.

E Robustness

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

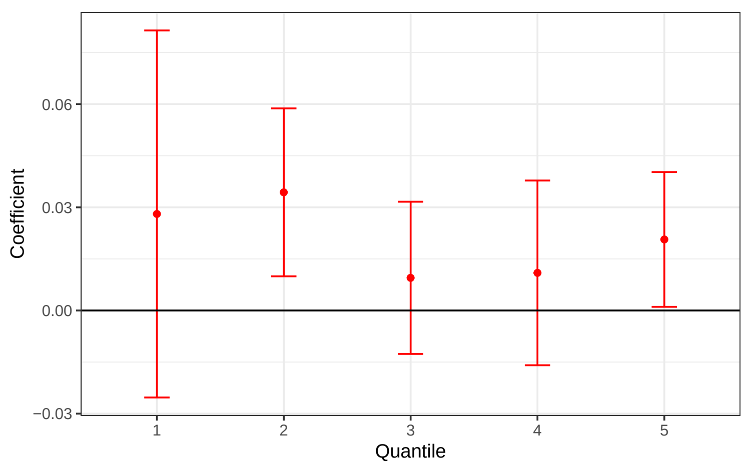

Figure A25: Varying Response to Fixed Lead of HDD

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.

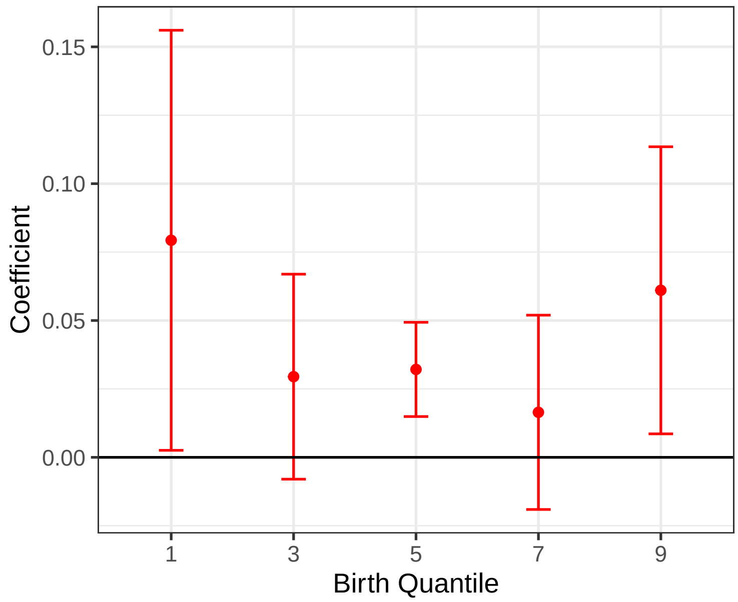

Figure A26: Profit Response to Forecasts across Manager's Age

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.

E.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 A24: Farm-Level Profit - Lead 1 with Lags

Dependent Variables:	Value Added	Profit
Model:	(1)	(2)
Variables		
GDD	7.83	-2.9
	(14.8)	(14.4)
GDD (lag)	-3.25	-7.94
	(9.94)	(9.86)
GDD (F)	24.8	9.61
	(47.1)	(45.9)
HDD	-22.8	134
	(333)	(331)
HDD (lag)	353*	429*
	(201)	(219)
HDD (F)	2,202**	1,781**
	(849)	(762)
Mean	58,357	86,695
Unique Farms	2,217	2,217
Fixed-effects		
Farm	Yes	Yes
Year	Yes	Yes
Fit statistics		
Observations	16,314	16,314
R^2	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 Table A10, 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.

E.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 A25: Farm-Level Profit - Alternative Weather

Dependent Variables:	Value Added	Profit
Model:	(1)	(2)
Variables		
GDD	6.97	-4.1
	(11.2)	(11.3)
GDD (F)	-12.4	-20.3
	(44.7)	(44.6)
HDD	61.3	129
	(177)	(171)
HDD (F)	943***	797***
	(296)	(269)
FDD	68.2*	64.4
	(39.1)	(37.7)
FDD (F)	189	142
	(158)	(156)
Mean	58,332	87,330
Unique Farms	2,625	2,625
Fixed-effects		
Farm	Yes	Yes
Year	Yes	Yes
Fit statistics		
Observations	18,428	18,428
R^2	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 A17 we show the costs and revenue responses to these alternative weather calculations.

E.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 1.

Table A26: Farm-Level Profit - No Time Trend

Dependent Variables:	Revenue	Costs
Model:	(1)	(2)
Variables		
GDD	3.47	-1.27
	(10.7)	(7.36)
GDD (F)	47.8	37.7**
	(34.4)	(14.4)
HDD	-212	-184
	(303)	(149)
HDD (F)	1,342***	39.3
	(455)	(270)
Mean	155,386	123,249
Unique Farms	2,603	2,603
Fixed-effects		
Farm	Yes	Yes
Year	Yes	Yes
Fit statistics		
Observations	18,917	18,917
R^2	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.

E.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 A27: Cost and Revenue Reactions to Disaggregated Realized Weather

Dependent Variables:	Revenue	Costs
Model:	(1)	(2)
Variables		
GDD	1.51	-2.43
	(16.7)	(5.36)
GDD (F)	42.9	52.7**
	(45.7)	(22.5)
HDD	43.1	39.9
	(328)	(114)
HDD (F)	1,289***	52.9
	(256)	(37.9)
Mean	164,532	131,969
Unique Farms	1,890	1,890
Fixed-effects		
Farm	Yes	Yes
Year	Yes	Yes
Fit statistics		
Observations	12,754	12,754
R^2	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 measured at the village level, while forecasts remain measured at the department one.

E.6 Decomposing along the Forecast Error Sign

Table A28: Cost and Revenue Responses - Heterogeneity along the sign of Forecasts Errors: Forecasts - Realization (1 month forecasts)

Dependent Variables:	Revenue	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.