Heat Stress, Damage Control, and Structural Maladaptation

Thiam Moustapha[†]

[Click Here for Most Recent Version]

June 2026

Abstract

I develop a mechanism in which farmers can respond rationally to heat stress and still fail to offset yield losses. The mechanism is biological timing: the phenological phase associated with the pest-related signal guiding pesticide adjustment can differ from the phase in which heat most strongly affects yield. When these signals are misaligned, defensive expenditure may be imperfectly targeted to the yield-critical threat, a pattern I term *structural maladaptation*. Across four major crops in Quebec over 1991–2022, yield falls significantly with heat exposure, while estimated unit-cost responses are positive for all four crops and statistically significant for oats, barley, and soybeans. The strongest evidence comes from barley, the only crop for which the retained pesticide-response and yield-critical phases differ. For barley, the interaction between yield-critical heat and pesticide expenditure is negative and statistically significant: as sowing-phase heat rises, the conditional pesticide–yield association becomes less favorable. For oats and corn, the corresponding interactions are statistically indistinguishable from zero, while soybeans exhibit a positive interaction; all three are phase-aligned crops. A falsification exercise further shows that the barley interaction is not statistically distinguishable from zero when indexed to growing-phase heat. These results provide evidence that biological timing can constrain agricultural adaptation and motivate greater attention to phase-specific information and agronomic decision support rather than input expansion alone.

JEL Classification: Q12, Q54, Q16, D22, C23

Keywords: climate adaptation, damage control, defensive intensification, killing degree days, phenological phase mismatch, structural maladaptation, agricultural productivity, Quebec

[†]PhD candidate, University of Quebec at Montreal. Email: thiam.moustapha@uqam.ca

1 Introduction

Climate change threatens agricultural systems globally, but its economic consequences depend not only on the magnitude of thermal shocks but also on how producers adjust. A large empirical literature estimates reduced-form temperature effects on agricultural outcomes and documents substantial behavioral responses, including changes in pesticide expenditure, crop allocation, and planting decisions (Bareille et al., 2024; Burke and Emerick, 2016; Jagnani et al., 2021; Kawasaki, 2023). These reduced-form approaches capture the net outcome of direct physiological damage and endogenous adjustment, but generally do not identify why observed adaptation fails to eliminate residual losses. This paper asks what features of the production environment can generate such *incomplete adaptation*.

I develop a mechanism based on biological timing. A crop's phenological calendar can create a wedge between the climate signal associated with pest pressure and pesticide adjustment and the phase in which heat is most damaging to yield. When these windows differ, an input response that is rational given the observable signal may nevertheless be imperfectly aligned with the yield-critical threat. I refer to this timing-based failure as *structural maladaptation*.

This paper studies four major crops in Quebec (oats, barley, corn, and soybeans) over the period 1991–2022, using two complementary datasets. A thirty-two-year regional yield panel provides estimates of the yield response to heat across nine agricultural regions. Farm-level input expenditure surveys covering three waves allow estimation of input-expenditure responses and the unit cost of production. Climate variables are Killing Degree Days (KDD), the accumulated degree-days above a crop-specific threshold, constructed separately for three phenological phases: preparation, sowing, and growing season. The phase boundaries are prescribed by the provincial crop insurance authority on agronomic and institutional grounds, making them predetermined with respect to realized within-season heat and farmer-specific input choices.

The empirical strategy draws on two sources of variation. In the regional panel, identifi-

cation comes from within-region year-to-year variation in phase-specific KDD after absorbing region and year fixed effects. In the CECPA sample, input responses are identified from within-region variation across the three survey waves, conditional on region and survey-wave fixed effects. The cross-crop design provides a within-sample comparison between one crop for which the retained pesticide-response and yield-critical phases differ (barley) and three crops for which they coincide (oats, corn, and soybeans). This phase-alignment contrast provides the central empirical comparison for the maladaptation mechanism.

The analysis is grounded in an enriched damage-control production technology (Lichtenberg and Zilberman, 1986; Sexton et al., 2007) in which variable inputs are partitioned into two functionally distinct classes. Pesticides act as damage-control inputs: they suppress pest pressure without shifting the yield ceiling. Fertilizers and seeds act as frontier-expanding inputs: they raise potential yield but do not directly suppress pests. The marginal return to each class responds to a different component of the climate signal, generating opposing intensive-margin predictions that together account for the full heterogeneity of observed input responses. Five testable predictions emerge from the model. First, pesticide expenditure rises when the damage-control channel dominates, but may contract when the yield-value channel dominates. Second, frontier investment contracts when the survival-share channel dominates and expands when recovery demand dominates. Third, incomplete adaptation: the combined input response does not fully offset the direct physiological damage to potential yield. Fourth, *cost escalation*: residual yield loss raises unit cost through the output channel alone or jointly with expenditure adjustment. Fifth, structural maladaptation: for crops in which the pest-pressure phase differs from the yield-critical phase, defensive expenditure responds to an imperfect proxy for the heat shock most relevant for yield. When the two phenological windows are imperfectly correlated, pesticide adjustment is attenuated relative to the response that would occur under perfect phase alignment.

The empirical analysis yields five main sets of findings. First, heat exposure significantly reduces yields for all four crops, with semi-elasticities ranging from -0.015 to -0.026 per

killing degree-day. The damage, however, is not uniform: the phenologically critical window and temperature threshold differ sharply across crops. For cold-season cereals (oats, barley), the sowing and establishment phase is the yield-critical window, where above-threshold heat is associated with significantly lower yield (Porter and Gawith, 1999; Rezaei et al., 2015). For warm-season crops (corn, soybeans), the damage threshold is 30 °C and the critical window is grain fill and reproductive development (Boote et al., 2018; Lobell et al., 2013). These phase-specific patterns, assigned on agronomic grounds, are consistent with the crop-physiology literature and suggest that annual temperature aggregates confound economically distinct mechanisms. The yield penalties are statistically significant even after conditioning on region fixed effects, year fixed effects, and phase-specific precipitation controls, consistent with observed adjustment failing to fully offset heat-related damage.

Second, area allocation responds to observable signals but does not eliminate the adaptation challenge. Barley exhibits a significant retrospective area contraction: sown area in year t is lower following greater growing-season heat in year $t-1$, consistent with retrospective adjustment. Corn area is positively associated with current early-season warmth, with preparation-phase heat providing the cleanest pre-planting signal. No statistically detectable area response is found for oats or soybeans.

Third, input-expenditure responses are sharply crop- and input-specific. Oats reduce fertilizer and seed expenditure without a detectable pesticide response. Barley combines higher pesticide expenditure with lower fertilizer and seed expenditure. Corn increases expenditure on all three input categories, whereas soybeans reduce pesticide expenditure while increasing fertilizer and seed expenditure. These patterns correspond to distinct adaptation regimes in the model and show that heat-induced adjustment cannot be characterized as uniformly intensive or uniformly retrenched.

Fourth, the estimated unit-cost coefficient is positive for all four crops and statistically significant for oats, barley, and soybeans; the corn estimate is positive but imprecise. For oats and barley, total expenditure rises while yield falls, so the two accounting channels

point in the same direction. For soybeans, the total-expenditure response is statistically indistinguishable from zero while yield remains negatively associated with heat, consistent with an important output-side contribution to unit-cost growth.

Fifth, the significant negative pesticide–heat interaction is specific to barley, the only phase-mismatched crop. For oats and corn the interaction estimates are negative but statistically indistinguishable from zero, while the soybean estimate is positive and significant. The theory imposes no sign restriction for these phase-aligned crops. More importantly, replacing the barley sowing-phase interaction with growing-phase heat reduces the coefficient to approximately zero, while a joint specification yields a substantially more negative interaction for sowing than for growing heat. This phase specificity is consistent with the phase-mismatch diagnostic.

This paper connects to five literatures. The first is the climate-yield literature. A large literature has quantified the yield response to temperature using panel fixed-effects estimators applied to long historical records (Deschênes and Greenstone, 2007; Fisher et al., 2012; Lobell et al., 2013; Schlenker and Roberts, 2009). These studies establish the nonlinear temperature-yield relationship and the dominance of above-threshold degree-days as the relevant thermal variable. This paper builds on this foundation in three respects. First, I disaggregate the thermal dose into phase-specific components and show that the relevant threshold and the relevant phase differ systematically across crops, a dimension that annual aggregates cannot resolve. Second, I document the full cold-crop/warm-crop asymmetry using a common identification design applied separately to each crop. Third, this paper treats the yield equation not as the endpoint of the analysis but as one component of a joint damage-control system: the yield results discipline the input-expenditure and unit-cost estimates that follow. The yield literature provides the question; the production technology literature provides the mechanism through which the answer operates.

The second literature is the damage-control production literature. The damage-control framework introduced by Lichtenberg and Zilberman (1986) and extended by Sexton et al.

(2007) reconceptualizes pesticides as damage reducers rather than as complements to the production frontier; Carpentier and Letort (2012) show that ignoring this distinction biases cross-input elasticities. This literature has developed largely in isolation from the climate economics tradition: damage-control papers rarely model the source of the damage as an exogenous climate variable, and climate-yield papers rarely embed production technology in a damage-control framework. This paper bridges this gap. Embedding the damage-control technology in a model of phase-specific climate risk generates joint predictions about yields, input expenditure, and unit costs that neither literature delivers alone, in particular the conditions under which input response and yield response move in opposite directions, and the unit-cost consequences that follow. Testing these predictions requires observing how farm expenditure patterns covary with phase-specific heat exposure.

The third literature concerns farmer adaptation. Recent work has established that farmers respond to climate shocks through input adjustment. Jagnani et al. (2021) document that Kenyan smallholders increase pesticide use under early-season temperature shocks; Bareille et al. (2024) find weather-induced changes in pesticide purchases among French producers; and Kawasaki (2023), using Japanese crop data, further show that weather effects on pests and diseases depend strongly on the timing of exposure. Bareille and Chakir (2024) structurally separate the agronomic from the adaptation component of the yield response. This literature establishes that defensive adjustment occurs but leaves open whether such adjustment fully offsets heat-related losses and whether its effectiveness depends on the timing of the underlying climate signal. This paper contributes to this strand in two respects. I document persistent negative yield responses alongside observed input-expenditure adjustment and provide evidence that the conditional pesticide–yield relationship deteriorates under the phase mismatch observed for barley. The distinction matters for policy design: where adaptation is constrained by input availability, expanded access may help, whereas timing-based misalignment raises the potential value of phase-specific agronomic information. The present estimates do not rank these interventions by productivity or welfare effects.

The fourth literature is the crop phenology literature. The agronomic literature has long recognized that thermal sensitivity is stage-dependent. Porter and Gawith (1999) document that cardinal temperatures for development and sterility differ substantially across phenological stages in wheat; Rezaei et al. (2015) show across cereals that heat stress during reproductive development reduces grain number while heat during grain filling reduces grain weight. Hatfield et al. (2011) and Hatfield and Prueger (2015) synthesize this evidence across crops and stress the importance of temporal resolution for reliable damage prediction. My contribution to this strand is econometric and economic: I translate stage-dependent physiological vulnerability into phase-specific KDD variables suitable for panel estimation, and use the cross-phase structure to provide evidence for an economic mechanism: the informational wedge between the pest-related signal guiding within-season adjustment and the yield-critical heat exposure that may not yet be fully observed when those input decisions are made. Documenting the economic relevance of this wedge places the paper in direct dialogue with the maladaptation literature.

The fifth strand is the maladaptation and adaptation cost literature. Barnett and O'Neill (2010) formalize maladaptation as an adaptation response that reduces welfare or increases vulnerability relative to the baseline. In agricultural economics, Burke and Emerick (2016) find that longer-run adaptation mitigates less than half and more likely none of the large short-run productivity losses from extreme heat; Blanc and Schlenker (2017) warn that panel estimates conflate short-run damage with long-run adaptation depending on the identifying time horizon. This paper introduces a mechanism-based account of maladaptation that is distinct from all three. My use of the term is narrower and production-side: structural maladaptation denotes a timing-induced attenuation of protective adjustment, rather than a welfare ranking of observed behavior relative to a counterfactual adaptation policy. The mechanism developed here does not require adjustment frictions, resource constraints, or a particular time horizon of identification. It can arise from a structural feature of the crop's biological calendar that may not be resolved simply by increasing the level of input use or by

improving access to the same within-season signal. When the observable signal is misaligned with the yield-critical phase, the resulting loss of adaptive effectiveness can contribute to persistent output losses and associated unit-cost pressure.

The paper is organized as follows. Section 2 develops the damage-control framework and derives the five testable predictions formally. Section 3 describes the two datasets and the construction of phase-specific KDD. Section 4 presents the empirical strategy for each estimating equation. Section 5 reports the results in six steps mirroring the empirical design: yield damage, extensive-margin area adjustments, input-expenditure heterogeneity, budget-share reallocation, cost escalation, and structural maladaptation. Section 6 presents robustness checks. Section 7 translates the estimated coefficients into near-term yield and unit-cost projections under SSP2–4.5 (2025–2030 horizon). Section 8 concludes with policy implications and directions for future research.

2 Theoretical Framework

The damage-control framework rests on five structural elements. First, a damage-control production technology in which variable inputs are partitioned into damage-control inputs (pesticides) and frontier-expanding inputs (fertilizers, seeds), each responding to a distinct climate signal (Lichtenberg and Zilberman, 1986; Sexton et al., 2007). Second, an intensive margin for pesticides whose direction depends on whether the damage-control or yield-value channel dominates (Bareille et al., 2024; Jagnani et al., 2021; Kawasaki, 2023). Third, an intensive margin for frontier inputs whose direction depends on whether the survival-share or recovery-demand channel dominates. Fourth, incomplete adaptation: input adjustment fails to fully offset the direct damage to potential yield. A unit-cost consequence follows when the resulting output contraction exceeds any proportional reduction in expenditure. Fifth, phase mismatch: when the phenological window that drives the pest-related signal and input demand differs from the window that governs yield formation, defensive investment is

calibrated to an imperfect signal of the yield-critical threat, rendering defensive adjustment informationally constrained and potentially ineffective against the yield-critical shock.

2.1 Production Environment and the Damage Control Technology

Consider a risk-neutral farmer producing a set of crops $\mathcal{C} = \{c_1, \ldots, c_C\}$ over a growing season indexed by year t. For each crop c, the farmer cultivates area $S_{c,t} \geq 0$ (in hectares) and chooses a vector of variable inputs per hectare $\mathbf{x}_{c,t} \in \mathbb{R}_+^K$ comprising pesticides, fertilizers, and seeds. These inputs are partitioned into two functionally distinct classes:

$$\mathbf{x}_{c,t} = \big(\mathbf{f}_{c,t},\ x_{\text{pest},c,t}\big), \tag{2.1}$$

where $\mathbf{f}_{c,t} = (x_{fert,c,t},\ x_{seed,c,t})$ collects frontier-expanding inputs (fertilizers and seeds) that raise potential yield through improved nutrient supply and genetic yield potential, and $x_{\text{pest},c,t}$ is the damage-control input (pesticides) that suppresses biological damage without shifting the production frontier. This partition is not merely taxonomic: it determines the channel through which each input operates, the information signal to which its demand responds, and the unit-cost consequences of its intensification under heat stress.

Realized yield per hectare follows the damage-control production function introduced by Lichtenberg and Zilberman (1986):

$$y_{c,t} = A_{c,t}\Big(\theta_{c,t}^{\text{yield}},\ \mathbf{f}_{c,t}\Big)\ \cdot\ \big[1 - D\big(N_{c,t}(\theta_{c,t}^{\text{pest}},\ x_{\text{pest},c,t},\ \mathbf{z}_{c,t})\big)\big], \tag{2.2}$$

where the three components are defined as follows:

(i) $A_{c,t}(\theta_{c,t}^{\text{yield}},\ \mathbf{f}_{c,t}) > 0$ is potential yield, the output the crop would achieve in the absence of any biological damage. It depends on yield-relevant heat exposure $\theta_{c,t}^{\text{yield}}$ through direct physiological channels (enzyme denaturation, disruption of photosynthesis, and pollen sterility; Hatfield et al., 2011; Lobell et al., 2013; Schlenker and Roberts, 2009)

and on frontier inputs $\mathbf{f}_{c,t}$ through nutrient supply and genetic yield potential. Fertilizers and seeds raise the ceiling of what the crop can produce; they do not protect against biological damage once the ceiling is set.

(ii) $N_{c,t}(\theta^{\text{pest}}_{c,t}, x_{\text{pest},c,t}, \mathbf{z}_{c,t}) \geq 0$ is the level of effective pest pressure (pathogens, insects, and weeds) that reaches the crop after partial suppression by pesticides. Only the damage-control input $x_{\text{pest},c,t}$ enters this function; frontier inputs affect potential yield $A_{c,t}$ but do not directly suppress pests.

(iii) $D(\cdot) : [0, \bar{N}] \to [0, 1)$ is a biological damage function: the fraction of potential output destroyed by pests, defined on the biologically relevant domain $[0, \bar{N}]$ where $\bar{N}$ is the maximum feasible pest pressure. It satisfies $D(0) = 0$, $D' > 0$, $D'' > 0$, and $D(\bar{N}) < 1$.

The production function (2.2) embeds two distinct production logics within a single framework. Frontier-expanding inputs ($\mathbf{f}_{c,t}$) raise the yield ceiling $A_{c,t}$: their marginal product is the additional output that better nutrients or superior genetics can achieve in the absence of biological damage. The damage-control input ($x_{\text{pest},c,t}$) does not shift $A_{c,t}$; it suppresses pest pressure $N_{c,t}$, thereby recovering a larger fraction of the already-determined potential output. This distinction has a sharp implication: the marginal returns to the two input classes respond to different climate signals. The return to pesticides is governed by ambient pest pressure (itself driven by $\theta^{\text{pest}}_{c,t}$), whereas the return to fertilizers and seeds is governed by the yield ceiling $A_{c,t}$, which depends on $\theta^{\text{yield}}_{c,t}$. When these two heat exposures belong to different phenological phases, the information set that disciplines pesticide demand diverges from the signal that determines whether frontier investment is worthwhile (Carpentier and Letort, 2012; Lichtenberg and Zilberman, 1986; Sexton et al., 2007).

I impose the following structural assumptions throughout.[1]

(A1) **Direct climate damage:** yield-critical heat reduces potential yield.

[1]The main curvature and sign restrictions used below are: $A_{\theta^{\text{yield}}} < 0$; $N_{x_{\text{pest}}} < 0$, $N_{x_{\text{pest}}x_{\text{pest}}} \geq 0$; $N_{x_{\text{pest}}\theta^{\text{pest}}} \geq 0$; $N_{\theta^{\text{pest}}} > 0$; $D'' > 0$; $A_{f_k} > 0$, $A_{f_k f_k} \leq 0$; $A_{f_k \theta^{\text{yield}}}$ unrestricted.

(A2) **Concave damage control:** pesticides reduce pest pressure at decreasing returns.

(A3) **Heat–pest channel:** pest-signal heat raises effective pest pressure and weakly reduces the marginal suppressive effectiveness of each unit of pesticide applied.

(A4) **Convex damage:** marginal damage per unit of pest pressure is increasing.

(A5) **Frontier input productivity:** each frontier input raises potential yield at decreasing returns; the cross-partial $\partial^2 A / \partial f_k \, \partial \theta^{\text{yield}}$ is unrestricted in sign, reflecting whether stressed plants benefit from additional nutrients or lose the capacity to absorb them, a distinction central to the observed heterogeneity across crops.

2.2 Climate, Pest Pressure, and Phenological Specificity

Let $\boldsymbol{\theta}_{c,t} = (\theta_{c,\tau,t})_{\tau \in \mathcal{T}}$ denote the vector of phase-specific heat-stress exposures, where $\tau \in \mathcal{T} = \{\text{prep, sow, grow}\}$ indexes the three phenological windows of the crop calendar (defined precisely in Section 3). Two functionally distinct components of $\boldsymbol{\theta}_{c,t}$ are relevant:

- $\theta^{\text{pest}}_{c,t} \equiv \theta_{c,\tau^{\text{pest}},t}$: heat stress during the phase τ^{pest} that most directly stimulates pest proliferation and disease pressure. The farmer observes $\theta^{\text{pest}}_{c,t}$ within the growing season and uses it to calibrate input applications.

- $\theta^{\text{yield}}_{c,t} \equiv \theta_{c,\tau^*,t}$: heat stress during the phenologically critical phase τ^* for yield formation, which determines the realized value of $A_{c,t}(\cdot)$ and hence the actual agronomic damage. This phase may differ from τ^{pest}, and its effect may not be directly observable at the time of input application.

Effective pest pressure is governed by the function $H : \mathbb{R}_+ \times \mathbb{R}_+ \times \mathbb{R}^J_+ \to \mathbb{R}_+$:

$$N_{c,t} = H\big(\theta^{\text{pest}}_{c,t},\ x_{\text{pest},c,t},\ \mathbf{z}_{c,t}\big)\,, \tag{2.3}$$

where $\mathbf{z}_{c,t}$ denotes exogenous agronomic controls (precipitation, water deficit, frost days). Frontier inputs $\mathbf{f}_{c,t}$ do not enter H: fertilizers and seeds alter the yield ceiling but do not directly suppress pathogen, insect, or weed populations. Under (A2) and (A3), H satisfies $H_\theta > 0$ (heat raises pest pressure), $H_p < 0$ (pesticides suppress it), and $H_{p\theta} \geq 0$ (heat weakens pesticide effectiveness).[2] The cross-partial $H_{p\theta} \geq 0$ captures a structural tension: as heat raises pest pressure while potentially weakening the marginal effectiveness of pesticide use, the return to damage control becomes increasingly sensitive to the balance between the damage-control and yield-value channels, a tension central to the incomplete adaptation result established in Proposition 1.

2.3 The Farmer's Optimization Problem

The farmer's decision problem is structured by a two-stage timing protocol:

1. **Pre-season stage** (before the growing season begins): the farmer chooses planted area $S_{c,t}$ for each crop c, based on expectations over climate, output prices, and lagged observables. Land is allocated under the endowment constraint $\sum_{c \in \mathcal{C}} S_{c,t} \leq \bar{S}_t$, where $\bar{S}_t$ is the total farmable area.

2. **Within-season stage** (during the growing season): after observing realized heat stress $\theta^{\text{pest}}_{c,t}$, the farmer chooses variable inputs $\mathbf{x}_{c,t}$ per hectare for each planted crop.

The problem is solved by backward induction, characterizing within-season input demand first and then pre-season area allocation.

[2] Notation: $H_\theta \equiv \partial H / \partial \theta^{\text{pest}}$, $H_p \equiv \partial H / \partial x_{\text{pest}}$, $H_{p\theta} \equiv \partial^2 H / \partial x_{\text{pest}} \, \partial \theta^{\text{pest}}$.

Within-season input choice. Conditional on planted area $S_{c,t}$ and observed heat signal $\theta^{\text{pest}}_{c,t}$, the farmer maximizes expected per-hectare profit:

$$\begin{aligned}(\mathbf{f}^*_{c,t},\, x^*_{\text{pest},c,t}) &= \arg\max_{\mathbf{f}\geq\mathbf{0},\, x_{\text{pest}}\geq 0} \;\; \Pi_{c,t}(\mathbf{f}, x_{\text{pest}}) \\ &\equiv\; P_{c,t}\,\mathbb{E}\Big[y_{c,t}\Big(\theta^{\text{yield}}_{c,t},\, \mathbf{f},\, x_{\text{pest}}\Big)\;\Big|\;\theta^{\text{pest}}_{c,t}\Big] - \sum_{k\in\{fert,seed\}} w_{k,t}\, f_k - w_{\text{pest},t}\, x_{\text{pest}},\end{aligned} \tag{2.4}$$

where $P_{c,t}$ is the output price, $w_{k,t}$ is the unit cost of frontier input k, and $w_{\text{pest},t}$ is the unit cost of pesticides. The expectation is taken with respect to $\theta^{\text{yield}}_{c,t}$, which may not be perfectly observable at the time of input application. Since $\theta^{\text{yield}}_{c,t}$ may belong to a different phenological window than $\theta^{\text{pest}}_{c,t}$, the farmer conditions on a signal that contains only partial information about the yield-relevant damage.

The two classes of inputs are governed by distinct first-order conditions, reflecting their different channels of action.

Damage-control input (pesticides). The optimal pesticide level equates the expected marginal value of damage reduction to its price:

$$P_{c,t}\;\cdot\;\mathbb{E}\left[A_{c,t}\Big(\theta^{\text{yield}}_{c,t},\, \mathbf{f}^*_{c,t}\Big)\cdot\Big(-D'(N_{c,t})\Big)\cdot\frac{\partial H}{\partial x_{\text{pest}}}\;\middle|\;\theta^{\text{pest}}_{c,t}\right] = w_{\text{pest},t}. \tag{2.5}$$

Since $\partial H/\partial x_{\text{pest}} < 0$ by (A2), the left-hand side is strictly positive. The marginal value of pesticides depends on $A_{c,t}$: when heat destroys potential yield, the value of protecting what remains declines in proportion, a mechanism that proves central to the soybean disengagement result.

Frontier inputs (fertilizers and seeds). The optimal level of each frontier input equates the expected marginal contribution to potential yield, weighted by the surviving fraction of

output, to its price:

$$P_{c,t} \cdot \mathbb{E}\left[\frac{\partial A_{c,t}}{\partial f_k}\left(\theta_{c,t}^{\text{yield}}, \mathbf{f}_{c,t}^*\right) \cdot \left[1 - D\left(N_{c,t}^*\right)\right] \,\middle|\, \theta_{c,t}^{\text{pest}}\right] = w_{k,t}. \tag{2.6}$$

The return to frontier inputs is scaled by $[1 - D(N_{c,t}^*)]$, the fraction of potential yield that survives pest damage. When pest pressure is high, because $\theta_{c,t}^{\text{pest}}$ is large, this survival fraction shrinks, reducing the marginal value product of fertilizers and seeds even though frontier inputs do not enter the pest function directly. Assumption (A5) ensures concavity of $\Pi_{c,t}$ in $\mathbf{f}$. Under the concavity conditions stated in Appendix A.3, the within-season optimum is unique.

2.4 Defensive Input Demand and the Lemma of Defensive Intensification

Equations (2.5)–(2.6) characterize the farmer's optimal pesticide and frontier-input choices. The question of interest is how these choices respond to a rise in the pest-signal heat θ^{pest}: the signs of $\partial x_{\text{pest}}^*/\partial\theta^{\text{pest}}$ and $\partial f_k^*/\partial\theta^{\text{pest}}$. To characterize these comparative statics, it is useful to rewrite the first-order conditions in compact form. Define

$$M(\mathbf{f}, s) \;\equiv\; \mathbb{E}\big[A_{c,t}(\theta^{\text{yield}}, \mathbf{f}) \,\big|\, \theta^{\text{pest}} = s\big]$$

as the conditional expected potential yield given the observable pest signal, where $s \in \mathbb{R}_+$ denotes a realization of θ^{pest}, and

$$\Phi(x_{\text{pest}}, s) \;\equiv\; 1 - D(H(s, x_{\text{pest}}, \mathbf{z}))$$

as the surviving output share after pest damage. The first-order conditions (2.5)–(2.6) then factor as $P \cdot M \cdot \Phi_x = w_{\text{pest}}$ and $P \cdot M_{f_k} \cdot \Phi = w_k$, where $\Phi_x \equiv D'(N)|H_p| > 0$. Because M captures the yield-value channel and Φ captures the survival-share channel, changes in pest-

signal heat shift both objects simultaneously, creating the opposing responses characterized below.

Lemma 2.1 (Defensive intensification)**.** *Suppose assumptions (A1)–(A4) hold. Define*

$$q_p \equiv \underbrace{M_s(\mathbf{f}, \theta^{pest}) \cdot \Phi_x(x_{pest}, \theta^{pest})}_{\textit{yield-value channel}} + \underbrace{M(\mathbf{f}, \theta^{pest}) \cdot \Phi_{xs}(x_{pest}, \theta^{pest})}_{\textit{damage-control channel}}, \tag{2.7}$$

where $M_s \equiv \partial M/\partial \theta^{pest}$, $\Phi_x \equiv D'(N)|H_p| > 0$, *and* $\Phi_{xs} \equiv D''(N)H_\theta|H_p| - D'(N)H_{p\theta}$. *Then the marginal-return force on pesticide demand satisfies:*

$$\operatorname{sign}\left(\frac{\partial x^{\mathrm{BR}}_{pest}}{\partial \theta^{pest}_{c,t}}\right) = \operatorname{sign}(q_p), \tag{2.8}$$

where x^{BR} *denotes the best-response pesticide choice with frontier inputs held fixed. Defensive intensification holds when* $q_p > 0$*; retrenchment when* $q_p < 0$. *The sign of* q_p *depends on the relative strength of the damage-control and yield-value channels. Phase overlap affects the yield-value component through* M_s*, but does not by itself determine the direction of pesticide adjustment.*[3]

Proof. Differentiate the factored first-order condition $P \cdot M \cdot \Phi_x = w_{\mathrm{pest}}$ with respect to θ^{pest} and apply the implicit function theorem. The second-order condition signs the denominator negative; the numerator reduces to q_p as in (2.7). Full derivation in Appendix A.2. □

Remark. The damage-control condition $D''H_\theta|H_p| > D'H_{p\theta}$ holds automatically when $H_{p\theta} = 0$ ((A3) with equality) and more generally whenever pest pressure is strongly responsive to heat or damage is sufficiently convex. When the phases are largely distinct, M_s may be close to zero, making the damage-control channel more likely to dominate. Greater phase overlap strengthens the yield-value channel, but the sign of q_p ultimately depends on the relative magnitude of the two forces. Section 2.5 maps the resulting regime to the four Quebec crops.

[3] A sufficient condition is $D''(N)H_\theta|H_p| > D'(N)H_{p\theta}$ and $M_s \geq 0$. When $M_s < 0$ (correlated phases), intensification instead requires $M\Phi_{xs} > |M_s|\Phi_x$. See Appendix A.2.

Lemma 2.2 (Frontier disengagement)**.** *Suppose assumptions (A2), (A3), (A4), and (A5) hold. Define*

$$q_k \equiv \underbrace{M_{f_k s}(\mathbf{f}, \theta^{pest}) \cdot \Phi(x_{pest}, \theta^{pest})}_{\textit{recovery-demand channel}} + \underbrace{M_{f_k}(\mathbf{f}, \theta^{pest}) \cdot \Phi_s(x_{pest}, \theta^{pest})}_{\textit{survival-share channel}}, \tag{2.9}$$

where $M_{f_k s} \equiv \partial^2 M / \partial f_k \, \partial \theta^{pest}$ *and* $\Phi_s \equiv -D'(N) H_\theta < 0$*. Then the marginal-return force on frontier input demand satisfies:*

$$\operatorname{sign}\left(\frac{\partial f_k^{\mathrm{BR}}}{\partial \theta_{c,t}^{pest}} \right) = \operatorname{sign}(q_k), \tag{2.10}$$

where f_k^{BR} *denotes the best-response frontier input with the damage-control input held fixed. Frontier disengagement holds when* $q_k < 0$*; recovery-demand when* $q_k > 0$*.*[4]

Proof. Differentiate the factored first-order condition $P \cdot M_{f_k} \cdot \Phi = w_k$ with respect to θ^{pest} and apply the implicit function theorem. The second-order condition signs the denominator negative; the numerator reduces to q_k as in (2.9). Full derivation in Appendix A.2. □

The economic intuition of Lemma 2.2 is direct: when pest-signal heat rises, the surviving output share Φ shrinks ($\Phi_s < 0$), reducing the marginal value product of frontier inputs and inducing disengagement. This is the survival-share channel ($M_{f_k} \Phi_s < 0$). The recovery-demand channel ($M_{f_k s} \Phi$) operates in the opposite direction when heat raises the marginal productivity of frontier inputs through compensatory mechanisms ($M_{f_k s} > 0$).

Which channel dominates depends on the crop-specific response of marginal frontier productivity to heat as well as on the decline in the surviving output share. Phase overlap affects these forces but does not determine the sign of q_k by itself. Section 2.5 maps the resulting regime to each Quebec crop.

Because pesticides and frontier inputs are jointly chosen, cross-input complementarity can modify these best-response effects; under standard own-curvature dominance conditions,

[4] A sufficient condition for $q_k < 0$ is $M_{f_k} |\Phi_s| > M_{f_k s} \Phi$; see Appendix A.2.

the signs carry through to the joint equilibrium for a representative frontier-input margin (Appendix A.3).

The joint signs of q_p and q_k characterize four theoretically distinct adaptation regimes, three of which are observed empirically in Quebec; Section 2.5 maps each regime to the corresponding crop.

Prediction 2.1 (Intensive-margin adaptation: pesticides)**.** *For crops with $q_p > 0$ (damage-control channel dominates), the best-response force on pesticide demand is positive. When additionally $q_k \geq 0$, the joint equilibrium response is unambiguously positive. When $q_k < 0$, the positive equilibrium response additionally requires own-curvature dominance (Appendix A.3). When $q_p < 0$, pesticide retrenchment is unambiguous when also $q_k < 0$. When $q_p < 0$ and $q_k > 0$, the negative equilibrium response additionally requires own-curvature dominance, as characterized in Appendix A.3.*

Prediction 2.2 (Frontier input adjustment)**.** *The best-response direction for frontier input k satisfies* $\text{sign}(\partial f_k^{\text{BR}}/\partial \theta^{pest}) = \text{sign}(q_k)$ *(Lemma 2.2). Under own-curvature dominance, this sign carries through to the joint equilibrium for a representative frontier-input margin (Appendix A.3).*

When $q_k < 0$, frontier investment contracts: unambiguously when additionally $q_p < 0$, and under own-curvature dominance when $q_p > 0$. When $q_k > 0$, frontier investment expands: unambiguously when $q_p > 0$, and under own-curvature dominance when $q_p < 0$.

2.5 Climate Adaptation Regimes

The joint sign of (q_p, q_k) defines four theoretically distinct adaptation regimes, three of which are observed empirically in Quebec; Table 1 summarizes the mapping. Full comparative statics are in Appendix A.3.

Regime 1 ($q_p > 0$, $q_k < 0$): The damage-control channel raises pesticide demand, while the falling survival fraction Φ contracts the expected return to frontier inputs: the two input

classes move in opposite directions. Barley follows this regime; oats occupies its boundary ($q_p \approx 0$, $q_k < 0$), with frontier inputs disengaging sharply but no significant pesticide response.

Regime 2 ($q_p > 0$, $q_k > 0$): Both the damage-control and recovery-demand channels are positive, so all inputs expand simultaneously, a double-burden regime. Corn follows: pesticide, fertilizer, and seed expenditures all rise with growing-season heat.

Regime 3 ($q_p < 0$, $q_k < 0$): The yield-value channel dominates both margins, driving rational aggregate retrenchment: the farmer optimally reduces total input use precisely when pest pressure is highest. None of the four crops in our sample occupies this regime, though it constitutes a theoretically relevant benchmark.

Regime 4 ($q_p < 0$, $q_k > 0$): The yield-value channel collapses pesticide demand, while the recovery-demand channel expands frontier investment, again in opposite directions but with the sign pattern reversed relative to Regime 1. Soybeans follow: pesticide expenditure falls sharply while fertilizer and seed expenditures rise, consistent with the collapse of the damage-control return alongside the recovery of the yield frontier.

Testable predictions are summarized in Table 13.

2.6 Incomplete Adaptation: The Central Proposition

Lemmata 2.1 and 2.2 characterize how farmers rationally adjust input use in response to pest-signal heat, with the direction of each adjustment depending on crop-specific biological conditions. The central question is whether any input response (whether intensification, reallocation, or retrenchment) can fully offset the direct damage to yield. Adaptation is complete if the total effect of $\theta^{\text{yield}}_{c,t}$ on realized yield, accounting for the endogenous input adjustment, is zero. Whether this holds depends on the relative magnitudes of direct damage and adaptive response; the proposition below gives the condition under which adaptation is incomplete.

Proposition 1 (Incomplete adaptation)**.** *Suppose assumptions (A1)–(A4) hold. The total effect of yield-critical heat on realized yield decomposes as:*

$$\frac{d\, y^*_{c,t}}{d\, \theta^{\text{yield}}_{c,t}} = \underbrace{\left.\frac{\partial\, y^*_{c,t}}{\partial\, \theta^{\text{yield}}_{c,t}}\right|_{\mathbf{x}^*}}_{\text{direct damage, } <0} + \underbrace{\nabla_{\mathbf{x}}\, y_{c,t} \cdot \frac{d\mathbf{x}^*_{c,t}}{d\, \theta^{\text{yield}}_{c,t}}}_{\text{adaptation response}}. \tag{2.11}$$

Adaptation is incomplete whenever the second term fails to offset the first.

Proof. By (A1), the direct effect of yield-critical heat at fixed inputs is negative. Applying the chain rule gives equation (2.11). Hence the total yield effect remains negative whenever the endogenous adaptation term fails to offset this direct damage. Formal details are provided in Appendix A.2. □

The economic intuition of Proposition 1 is as follows. Heat stress during the yield-critical phase directly reduces potential yield $A_{c,t}$ through the frontier technology; at a given pest-signal s, this operates independently of the survival fraction Φ. Input adjustment can partially offset this damage: pesticide intensification raises the survival fraction Φ, and frontier investment raises A, but in this model such adjustment is triggered by the pest-phase signal $\theta^{\text{pest}}_{c,t}$, not directly by $\theta^{\text{yield}}_{c,t}$. When the two phases are imperfectly correlated, the adaptive response can be attenuated. Whether the resulting total yield effect remains negative is an empirical question, evaluated by the reduced-form yield specification below. If the estimated total yield effect (without conditioning on input choices) remains negative and significant, adaptation is incomplete. This result generalizes the observation in Deschênes and Greenstone (2007) and Fisher et al. (2012) that farm-level adaptation does not eliminate the aggregate yield response to temperature, and provides a microfoundation for the documented residual damage in panel estimates (Bareille and Chakir, 2024; Schlenker and Roberts, 2009).

Prediction 2.3 (Incomplete adaptation: central test)**.** *The total effect of yield-relevant heat stress on realized output is negative and significant. The empirical test excludes endogenous input controls so that the regression coefficient captures the net reduced-form response of yield*

to heat exposure, including endogenous behavioral adjustment.

2.7 Unit-Cost Implications

The joint effect of heterogeneous input responses and incomplete yield protection exerts upward pressure on the cost of production per unit of output. Whether the farmer intensifies defensively, reallocates the budget, or retrenches, the yield denominator falls; the net effect on unit cost depends on whether this output contraction outpaces any reduction in expenditure. Define production cost per tonne as:

$$c_{c,t} \equiv \frac{\text{Total expenditure}_{c,t}}{y_{c,t}} = \frac{\sum_{k=1}^{K} w_{k,t}\, x^{*}_{k,c,t}\, S_{c,t} + \text{Fixed costs}_{c,t}}{y_{c,t}\, S_{c,t}}. \tag{2.12}$$

Proposition 2 (Unit-cost implications of incomplete adaptation)**.** *Under Proposition 1,*

$$\frac{d \ln c_{c,t}}{d\theta^{\text{yield}}} = \frac{d \ln \text{Dep}_{c,t}}{d\theta^{\text{yield}}} - \frac{d \ln y_{c,t}}{d\theta^{\text{yield}}}.$$

Since $d \ln y_{c,t}/d\theta^{\text{yield}} < 0$*, unit cost rises automatically when total expenditure is non-decreasing. If total expenditure falls, unit cost rises whenever*

$$\left| \frac{d \ln y_{c,t}}{d\theta^{\text{yield}}} \right| > \left| \frac{d \ln \text{Dep}_{c,t}}{d\theta^{\text{yield}}} \right|. \tag{2.13}$$

Proof. See Appendix A.2. □

Proposition 2 and its decomposition (equation (A.2) in Appendix A.2) admit a natural accounting interpretation. Under complete yield adaptation, the yield term would vanish; unit cost would then move one-for-one with the expenditure channel, which need not be zero. Under incomplete adaptation, the denominator (output) shrinks because of residual yield damage that equilibrium input adjustment does not fully offset. The expenditure channel (positive for crops in a double-burden regime, negative for crops in a retrenchment regime) modulates the magnitude of the net cost change.

Prediction 2.4 (Rising unit cost of production)**.** *Under the condition in Proposition 2, heat stress raises unit cost through the yield channel alone or through the combination of yield loss and expenditure adjustment. The expenditure channel is crop-specific and need not share the sign of any individual input response.*

2.8 Phase Mismatch: A Mechanism for Structural Maladaptation

Proposition 1 characterizes the conditions under which endogenous adjustment fails to fully offset direct heat damage. A deeper form of incomplete adaptation, structural maladaptation, can arise when the phenological phases τ^{pest} and τ^* are imperfectly aligned.

Define the informativeness of the pest-phase signal as:[5]

$$\rho_c \;\equiv\; \text{Corr}\Big(\theta_{c,t}^{\text{yield}},\, \theta_{c,t}^{\text{pest}}\Big) \;\in\; [0,1]. \tag{2.14}$$

When $\tau^{\text{pest}} = \tau^*$, the phases coincide and $\rho_c = 1$: the observable signal perfectly predicts the damage-relevant shock. When $\rho_c < 1$, the farmer's information set is an imperfect proxy for the yield-critical shock, and input adjustment is attenuated relative to the response that would occur under perfect phase alignment.

Proposition 3 (Phase mismatch attenuates adaptive response)**.** *Suppose: (a) the pest-response and yield-critical phases are imperfectly correlated ($\rho_c < 1$); and (b) equilibrium pesticide adjustment responds positively to the observable pest signal ($dx^*_{pest}/d\theta^{pest} > 0$). Then:*

(i) The pesticide response to yield-critical heat is attenuated:

$$\frac{dx^*_{pest}}{d\theta^{yield}} = \lambda_c \frac{dx^*_{pest}}{d\theta^{pest}}, \qquad \lambda_c = \rho_c \frac{\sigma_s}{\sigma_r},$$

and is therefore strictly smaller than its value at $\rho_c = 1$.[6]

[5] The restriction $\rho_c \in [0,1]$ follows from Cauchy–Schwarz and the agronomic prior $\text{Cov}(\theta^{\text{yield}}, \theta^{\text{pest}}) \geq 0$. Appendix A.2 formalizes the signal structure using linear conditional expectations.

[6] $\sigma_s \equiv \text{Std}(\theta_{c,t}^{\text{pest}})$ and $\sigma_r \equiv \text{Std}(\theta_{c,t}^{\text{yield}})$ are the standard deviations of the pest-phase and yield-critical heat

(ii) If additionally the net within-season input adjustment is yield-protective, $\Gamma \equiv \nabla_{\mathbf{x}} y \cdot d\mathbf{x}^*/d\theta^{pest} > 0$*, the adaptive component of the yield response to yield-critical heat,* $\lambda_c \Gamma$*, is attenuated relative to perfect phase alignment.*

(iii) As $\rho_c \to 0$*, adjustment triggered by the pest-phase signal provides no systematic adaptive response to variation in the yield-critical shock.*

Proof. See Appendix A.2. □

Application to barley (*orge*) in Quebec. Barley provides the clearest empirical instance of phase mismatch in the sample. The yield specification retains sowing-phase KDD_{25}, whereas growing-phase KDD_{25} is retained for pesticide expenditure. The two phase-specific heat measures are imperfectly correlated in the joint specification ($\hat{\rho} = 0.737$ after fixed-effect residualization), and barley yield is more sensitive per KDD to sowing-phase than to growing-phase heat. This discordance motivates the phase-mismatch diagnostic in equation (4.8).

Prediction 2.5 (Phase mismatch: structural maladaptation)**.** *For crops in which* $\tau^{pest} \neq \tau^*$*, the pest-phase signal is an imperfect proxy for yield-critical damage (*$\rho_c < 1$*, Proposition 3). The farmer's within-season pesticide response to yield-critical heat is attenuated, so the adaptive component of the yield response is smaller than under perfect phase alignment. The empirical counterpart of this mechanism is a deterioration in the conditional pesticide–yield association as yield-critical heat rises, tested in Section 5 via the augmented yield regression (equation (4.8)).*

3 Data

The analysis combines four data sources. Daily climate data are obtained from the NASA Prediction Of Worldwide Energy Resources (POWER) Daily API (NASA Langley Research Center POWER Project, 2026) over 1991–2022. POWER meteorological variables are derived

shocks, respectively. The ratio σ_s/σ_r scales the attenuation: if pest-phase heat varies less than yield-critical heat, $\lambda_c < \rho_c$.

primarily from NASA's MERRA-2 reanalysis at a native grid of 0.5° latitude × 0.625° longitude. For each agricultural region, daily maximum temperature, minimum temperature, and precipitation are extracted at one or more representative coordinates and averaged to a regional daily series; these records are used to construct the phase-specific KDD measures, precipitation controls, and moisture-deficit indices. NASA POWER temperature products perform well for daily minimum and maximum temperatures in validated comparisons against ground-station observations (Rodrigues and Braga, 2021; White et al., 2008). Phenological phase boundaries are based on the provincial crop-insurance calendar administered by the *Financière agricole du Québec* (FADQ). The crop-specific upper damage thresholds ($\bar{T}_c \in \{25, 30\}$ °C) are agronomically motivated: 25 °C for cool-season cereals and 30 °C for warm-season crops, following the crop-physiology literature (Hatfield et al., 2011; Porter and Gawith, 1999). Annual crop yields (t/ha) and harvested area (ha) at the administrative region level are sourced from Institut de la statistique du Québec (ISQ) and Statistics Canada (2023), compiling Statistics Canada agricultural surveys for nine Quebec regions over the full 1991–2022 period. Finally, farm-level input expenditure data (pesticides, fertilizers, and seeds, in $/ha) are drawn from triennial surveys (Centre for Studies on Agricultural Production Costs (CECPA), 2021) for survey years 2009, 2014, and 2019.

The analysis draws on two estimation samples. The primary panel covers nine agricultural regions of Quebec over 1991–2022 and is used for yield and sown-area regressions. Yield observations with zero harvested area are excluded because yield per hectare is undefined in those cells. The input-expenditure regressions use a second, smaller sample drawn from three CECPA survey waves (2009, 2014, 2019), merged with three-year backward-looking moving averages of climate variables ending at each survey year.[7]

[7]The nine regions are: Bas-Saint-Laurent, Saguenay–Lac-Saint-Jean, and Abitibi-Témiscamingue (aggregated into one northern unit); Capitale-Nationale and Mauricie; Estrie; Lanaudière; Laurentides; Outaouais; Chaudière-Appalaches; Montérégie; and Centre-du-Québec. For the northern aggregate (Bas-Saint-Laurent/Gaspésie–Îles-de-la-Madeleine, Saguenay–Lac-Saint-Jean/Côte-Nord, and Abitibi-Témiscamingue/Nord-du-Québec), the NASA POWER climate series is averaged across three representative coordinates; all other regions use one coordinate each. Full extraction coordinates are available upon request.

3.1 Growing Calendar and Phenological Phases

The construction of phase-specific KDD requires mapping each calendar day to one of the three phenological windows $\tau \in \mathcal{T} = \{\text{prep, sow, grow}\}$ introduced in Section 2.2. Phase boundaries are taken from the administrative calendars used by Quebec's crop-insurance program. Because these calendars are fixed independently of realized within-season heat, they provide a predetermined mapping of calendar days into phenological phases, which addresses the endogeneity concern in the assignment of days to phases.

The phase boundaries are prescribed by the FADQ crop-insurance calendar. The crop-specific damage thresholds are agronomically motivated: cool-season crops (oats, barley) are assigned $\bar{T}_c = 25\,^{\circ}\text{C}$, and warm-season crops (corn, soybeans) $\bar{T}_c = 30\,^{\circ}\text{C}$, consistent with their higher thermal optima (Hatfield et al., 2011; Porter and Gawith, 1999). The phase calendars are as follows.

The preparation phase (PP) spans the weeks preceding planting, during which soil preparation and weed management occur, and the farmer begins to observe early-season heat signals. The sowing phase (GS1) covers germination and seedling emergence; the growth and maturation phase (GS2) runs from the close of GS1 through harvest and captures later vegetative and reproductive development as well as much of in-season pest pressure. The yield-critical phase is crop-specific and need not coincide with GS2.

Table 2 summarizes the phase boundaries and damage thresholds for the four crops retained in the analysis.

The phase boundaries are fixed by the FADQ crop-insurance calendar, whereas the crop-specific damage thresholds are selected ex ante from the crop-physiology literature rather than estimated from the data. Both components are therefore predetermined with respect to the estimation. Appendix D.1 assesses robustness to alternative threshold choices.

3.2 Killing Degree Days

Phase-specific extreme heat exposure is measured using Killing Degree Days (KDD). For crop c, phenological phase τ, and year t, KDD is defined as the accumulated daily thermal exposure above the crop-specific physiological damage threshold $\bar{T}_c$:

$$\theta_{c,\tau,t} = \sum_{d \in \mathcal{W}_{c,\tau}} h\big(T_{d,t}^{\max},\, T_{d,t}^{\min};\, \bar{T}_c\big), \tag{3.1}$$

where $h(T_{d,t}^{\max}, T_{d,t}^{\min}; \bar{T}_c)$ is the degree-days above threshold $\bar{T}_c$ on day d of year t, computed by single-sine interpolation from daily maximum and minimum temperatures following Schlenker and Roberts (2009); the exact piecewise formula is reported in Appendix B.1. The set $\mathcal{W}_{c,\tau}$ is the set of calendar days belonging to phase τ for crop c as described in Section 3.1. Daily temperature records are extracted from the NASA POWER grid cell corresponding to the representative coordinate(s) of each agricultural region and averaged where multiple coordinates are used.

Unlike standard Growing Degree Days, which aggregate heat beneficial to development between a lower and an upper threshold, KDD capture only the damaging tail of the temperature distribution. Above $\bar{T}_c$, additional heat impairs cellular processes: it denatures key metabolic enzymes, disrupts protein synthesis, reduces photosynthetic efficiency, and compromises pollen viability (Hatfield et al., 2011; Lobell et al., 2013). These physiological injuries are largely irreversible and accumulate nonlinearly with exposure duration (Schlenker and Roberts, 2009). The empirical superiority of this measure over average seasonal temperature is well established: Schlenker and Roberts (2009) show that above-threshold heat accumulation predicts US crop yield losses far more precisely than seasonal means, because it captures the nonlinear and asymmetric character of the temperature–yield relationship.

A critical feature of equation (3.1) is the phase subscript τ. The agronomic literature establishes that crop sensitivity to heat stress is strongly stage-dependent (Hatfield et al., 2011): the same thermal dose inflicts qualitatively different injuries depending on whether it occurs

during germination, vegetative growth, or reproductive development. This heterogeneity is the empirical basis for the phase-mismatch mechanism documented in Section 4 and for the KDD selection strategy described in Appendix B.1.

Descriptive statistics of KDD

Table 3 reports phase-specific KDD across the nine Quebec agricultural regions over 1991–2022. Heat exposure is strongly uneven across the crop calendar. For cool-season cereals, exposure above 25 °C is negligible during preparation, moderate during sowing, and substantially larger during the growing phase. For corn and soybeans, exposure above 30 °C is concentrated almost entirely in the growing phase. Warm-crop KDD values are consequently much smaller in absolute magnitude, reflecting the higher physiological threshold; raw KDD levels, and the magnitudes of coefficients expressed per KDD, should therefore not be compared mechanically across crop groups.

Table 3 also highlights the cross-phase contrast central to the empirical analysis. Barley is the only crop for which the phase retained for yield differs from the phase retained for pesticide expenditure: sowing-phase KDD_{25} is yield-relevant, whereas growing-phase KDD_{25} is most strongly associated with pesticide expenditure. For all other crops, the two retained phases coincide. The phase-selection procedure and supporting regressions are reported in Appendix B.1 and introduced formally in Section 4. Beyond these phase-specific differences, Figure 1 documents substantial temporal and spatial variation in the KDD measures used in the yield regressions. For each crop, the retained yield-relevant KDD series fluctuates sharply from year to year, with episodic heat spikes rather than a smooth increase over the sample period. The shaded regional ranges also reveal substantial within-year spatial dispersion in exposure. None of the four series exhibits a statistically significant linear trend over 1991–2022 (all $p > 0.33$). The different vertical scales across panels reflect the crop-specific physiological thresholds: exceedances of 30 °C are substantially rarer than exceedances of 25 °C in Quebec. The resulting variation across regions and years provides the underlying

variation for the panel analysis; the fixed-effects specifications identify the heat response from within-region year-to-year deviations after absorbing permanent regional differences and common year shocks.

3.3 Crop Yields and Area

Regional outcome data provide annual marketable production, harvested area, and sown area for each crop-region-year over 1991–2022. Yield is defined as marketable production divided by harvested area. Region-crop-year observations with zero harvested area are excluded because yield per hectare is undefined in those cells. The extensive-margin analysis (Section 5.2) instead uses sown area, which records the area committed at planting and therefore precedes within-season heat exposure and crop losses.

Table 4 reports summary statistics for yields and harvested area. Corn has the highest average yield in the sample, at 7.60 tonnes per hectare, compared with 3.01 for barley, 2.66 for soybeans, and 2.46 for oats, and also exhibits the largest unconditional dispersion across region-year observations. Average harvested area differs substantially across crops, with corn and soybeans occupying considerably larger regional areas than oats and barley. These level differences are absorbed by the crop-specific region fixed effects used in the empirical specifications.

3.4 Variable Inputs: Pesticides, Fertilizers, and Seeds

Table 5 reports expenditure distributions and sample sizes for pesticides, fertilizers, and seeds across the four crops. The CECPA sample covers conventional farms, and sample sizes differ across input categories within crops because of item non-response. Corn has by far the highest average fertilizer and seed expenditures, at $383 and $225 per hectare respectively, whereas soybeans have the highest average pesticide expenditure ($87/ha) and the lowest fertilizer expenditure ($56/ha). Barley and oats generally occupy lower expenditure levels. These differences in baseline input composition are useful for interpreting the crop-specific

responses estimated below, but the empirical coefficients capture changes in expenditure with heat exposure rather than differences in expenditure levels across crops.

3.5 Additional Weather Controls

The empirical specifications include three groups of additional weather controls, all constructed at the phenological-phase level to help separate above-threshold heat exposure from correlated variation in water availability.

Phase-specific precipitation. Total precipitation (mm) is accumulated separately for each phenological phase from daily records provided by NASA Langley Research Center POWER Project (2026). The yield regressions include precipitation for the preparation, sowing, and growing phases. The extensive-margin specification includes preparation- and sowing-phase precipitation, while the CECPA specifications use three-year backward-looking averages of the corresponding phase-specific measures.

Phase-specific moisture deficit. A water-deficit index is constructed at the phase level as the difference between phase-accumulated precipitation and phase-accumulated reference evapotranspiration, both drawn from the NASA POWER API (NASA Langley Research Center POWER Project, 2026). The yield regressions include moisture-deficit controls for the sowing and growing phases. The extensive-margin specification includes preparation- and sowing-phase measures. The CECPA specifications use three-year backward-looking averages of the corresponding phase-specific deficit measures. A negative deficit signals that atmospheric evaporative demand exceeds available soil moisture, a condition that compounds the physiological stress associated with above-threshold temperatures.

Preparation-phase frost. The yield regressions additionally control for the number of preparation-phase days with minimum temperature below 0 °C. This variable captures early-season cold exposure that is distinct from above-threshold heat stress.

4 Empirical Strategy

The empirical strategy translates the theoretical predictions of Section 2 into three complementary empirical designs. First, a long regional panel is used to estimate the effects of phase-specific heat on crop yields and sown area. Second, a three-wave farm-level CECPA sample is used to characterize input-expenditure responses, expenditure composition, and unit production costs. Third, an augmented farm-level production specification interacts pesticide expenditure with yield-relevant heat to test the phase-mismatch prediction (Prediction 2.5).

Implementing these tests requires mapping the phase-specific KDD measures into the two climate objects of the theoretical model: θ^{yield}, the heat realization relevant for yield formation, and θ^{pest}, the heat signal relevant for pest pressure and pesticide adjustment.

Yield-critical phase (τ^*). The yield-critical phase is assigned ex ante using agronomic evidence on crop-specific thermal sensitivity windows and the institutional phenological calendar prescribed by Quebec's Crop Insurance Programme (Section 3.1). For oats and barley, this evidence identifies the sowing and establishment phase as the relevant window; for corn and soybeans, the corresponding window is the growing and reproductive phase. Phase-by-phase yield regressions reported in Table 14 provide supporting econometric evidence for these agronomic assignments. The designated phases generate yield responses consistent with the expected direction and biological timing, while the full set of candidate-phase estimates is reported for transparency. These regressions serve as supporting evidence rather than as a mechanical selection rule.

Pesticide-response phase (τ^{pest}). The phase relevant for observed pesticide adjustment is evaluated using both agronomic plausibility and the within-region, within-wave explanatory power of phase-specific KDD for pesticide expenditure. For each candidate phase, pesticide expenditure and KDD are residualised on region and survey-wave fixed effects, and the resulting slopes and partial R^2 values are reported in Table 14. These diagnostics are used as supporting evidence rather than as a mechanical selection rule. The full input-expenditure specification in equation (4.3) additionally conditions on weather controls and log output

price.

The resulting phase mapping reveals the contrast central to the empirical design. For barley, sowing-phase KDD_{25} is yield-relevant, whereas growing-phase KDD_{25} is the phase most closely associated with pesticide adjustment. Barley is therefore the only crop for which the retained yield-relevant and pesticide-response phases differ. For oats, corn, and soybeans, the two retained phases coincide. This discordance creates the phase contrast required by the phase-mismatch mechanism of Section 2.8 and motivates the targeted test in Section 4.5.

4.1 Yield Panel

The primary yield specification is

$$\ln y_{r,c,t} = \beta_c \, \theta_{r,c,\tau^*(c),t} + \mathbf{X}'_{r,c,t} \boldsymbol{\gamma} + \alpha_r + \lambda_t + \varepsilon_{r,c,t}, \tag{4.1}$$

where $y_{r,c,t}$ is yield per hectare (t/ha) in region r, crop c, and year t; $\theta_{r,c,\tau^*(c),t}$ is KDD accumulated during the crop-specific yield-critical phase $\tau^*(c)$, as described in Section 3.2; and $\mathbf{X}_{r,c,t}$ contains the additional weather controls described in Section 3.5. Region fixed effects α_r absorb time-invariant regional heterogeneity within each crop-specific regression, while year fixed effects λ_t absorb shocks common to all regions in a given year.

Equation (4.1) is estimated separately for each crop using the within-group fixed-effects estimator, with region as the panel identifier. Standard errors are clustered at the region level, allowing for arbitrary serial correlation and heteroskedasticity within regions. Tables 17 and 18 report complementary specifications allowing for differential regional trends and pooling crops within thermal groups, respectively.

Input expenditures are excluded from Equation (4.1) because they are themselves endogenous responses to heat. The coefficient β_c is therefore interpreted as the reduced-form effect of heat on yield, incorporating both direct physiological damage and behavioral adjustment through inputs. Conditioning on realized input choices would block part of this adaptation

channel and alter the estimand of interest. This corresponds to Proposition 1, which predicts a negative net yield response whenever behavioral adjustment fails to fully offset direct heat damage.

4.2 Extensive Margin: Heat Exposure and Sown Area

Beyond input adjustment, producers may also respond to climate conditions through the extensive margin by changing the area allocated to a crop. Because sown area is determined at planting, the timing of the relevant climate information differs from that governing within-season input adjustment. I distinguish between current early-season thermal conditions and retrospective information from the preceding growing season.

The specification is

$$\ln S_{r,c,t} = \delta_{1,c}\,\theta_{r,c,\text{prep},t} + \delta_{2,c}\,\theta_{r,c,\text{sow},t} + \delta_{3,c}\,\theta_{r,c,\text{grow},t-1} + \delta_{4,c}\,\ln y_{r,c,t-1} + \mathbf{X}'_{r,c,t}\,\boldsymbol{\gamma} + \alpha_r + \lambda_t + \varepsilon_{r,c,t}, \tag{4.2}$$

where $S_{r,c,t}$ is total sown area for crop c in region r and year t. The current preparation-phase KDD, $\theta_{r,c,\text{prep},t}$, captures thermal conditions observed before planting. Sowing-phase KDD, $\theta_{r,c,\text{sow},t}$, captures thermal conditions around the planting and establishment window; because this measure is contemporaneous with the planting commitment, its coefficient is interpreted as an early-season association rather than as a fully pre-determined planting signal. The retrospective variables are previous-year growing-season KDD, $\theta_{r,c,\text{grow},t-1}$, and lagged log yield, $\ln y_{r,c,t-1}$, which summarize the most recent heat exposure and realized crop productivity in that region (Burke and Emerick, 2016).

The vector $\mathbf{X}_{r,c,t}$ contains precipitation and moisture-deficit controls for the current preparation and sowing phases. Equation (4.2) is estimated separately by crop with region and year fixed effects and standard errors clustered at the region level.

This specification therefore separates information available before planting from conditions

realized around the planting window and from retrospective signals inherited from the previous season. In particular, the lagged growing-season KDD tests for a retrospective adjustment channel in which current crop area responds to heat experienced during the preceding growing season.

4.3 Input-Expenditure Regressions

The three-wave CECPA sample is used to estimate how farm-level input expenditures respond to phase-specific heat. For each input $k \in \{\text{pesticides, fertilizers, seeds}\}$,

$$\ln x^{k}_{i,r,c,m} = \eta_{k,c}\,\bar{\theta}_{r,c,\tau^{\text{pest}}(c),m} + \mathbf{W}'_{r,c,m}\,\boldsymbol{\gamma}_k + \mathbf{Q}'_{i,r,c,m}\,\boldsymbol{\mu}_k + \phi_r + \psi_m + \nu^{k}_{i,r,c,m}, \tag{4.3}$$

where i indexes farms, r regions, c crops, and $m \in \{2009, 2014, 2019\}$ survey years. The climate regressor is the three-year backward-looking average of KDD during the crop-specific pesticide-response phase $\tau^{\text{pest}}(c)$ defined above:

$$\bar{\theta}_{r,c,\tau^{\text{pest}}(c),m} = \frac{1}{3}\sum_{\ell=0}^{2} \theta_{r,c,\tau^{\text{pest}}(c),m-\ell}. \tag{4.4}$$

$\mathbf{W}_{r,c,m}$ contains the corresponding three-year averages of regional weather controls, and $\mathbf{Q}_{i,r,c,m}$ contains log output price. Region fixed effects ϕ_r absorb time-invariant regional heterogeneity; survey-wave fixed effects ψ_m absorb shocks common to all farms in a given survey year. Equation (4.3) is estimated separately by crop and input category, with standard errors clustered at the region level. Organic farms are excluded.

For all three input categories, the climate regressor is constructed from the crop-specific pest-signal phase $\tau^{\text{pest}}(c)$ identified above: sowing for oats and growing for barley, corn, and soybeans. The three-year average smooths idiosyncratic annual weather variation and captures recent climate conditions relevant for input budgeting. Because CECPA records expenditures rather than physical quantities, $\eta_{k,c}$ is the semi-elasticity of expenditure on input k with respect to one additional KDD and is interpreted as an expenditure response

rather than a physical input-demand response.

Identification comes from within-region variation across survey waves in $\bar{\theta}_{r,c,\tau^{\text{pest}}(c),m}$, after absorbing region and survey-wave fixed effects and conditioning on observed weather controls. The identifying assumption is that the remaining variation in KDD is unrelated to other time-varying determinants of input expenditure. Predetermined phenological calendars and biologically motivated damage thresholds limit researcher discretion in constructing the heat measure, while phase-specific precipitation and moisture-deficit controls account for correlated water stress. Appendix C.1 discusses the identifying assumptions in detail.

4.4 Unit-Cost Effects and Decomposition

Proposition 2 links incomplete adaptation to the unit cost of production. Unit cost can respond to heat through two accounting channels: changes in expenditure per hectare and changes in output per hectare. Define farm-level cost per tonne as

$$\ln\left(\frac{C_{i,r,c,m}}{y_{i,r,c,m}}\right) = \ln C_{i,r,c,m} - \ln y_{i,r,c,m}, \tag{4.5}$$

where $C_{i,r,c,m}$ is total production expenditure per hectare and $y_{i,r,c,m}$ is output per hectare for farm i, region r, crop c, and survey year m. Both variables are observed in the CECPA sample.

The reduced-form unit-cost response is estimated directly from

$$\ln\left(\frac{C_{i,r,c,m}}{y_{i,r,c,m}}\right) = \kappa_c\, \bar{\theta}_{r,c,\tau^*(c),m} + \mathbf{W}'_{r,c,m}\,\boldsymbol{\gamma} + \mathbf{Q}'_{i,r,c,m}\,\boldsymbol{\mu} + \phi_r + \psi_m + u_{i,r,c,m}, \tag{4.6}$$

using the same region and survey-wave fixed effects and controls as the input-expenditure regressions. The coefficient κ_c measures the semi-elasticity of cost per tonne with respect to one additional KDD.

Using the same right-hand-side specification and the same yield-critical KDD, I also estimate a companion regression with $\ln C_{i,r,c,m}$ as the dependent variable; the corresponding

heat coefficient is denoted η_c^C.

Conceptually, the log identity in Equation (4.5) implies that, within a common sample and specification,

$$\frac{\partial \ln(C/y)}{\partial \bar{\theta}} = \underbrace{\eta_c^C}_{\text{expenditure channel}} - \underbrace{\beta_c}_{\text{yield channel}}, \tag{4.7}$$

where η_c^C denotes the response of total expenditure per hectare to heat. Proposition 2 predicts a positive unit-cost response whenever the proportional output loss exceeds any proportional reduction in expenditure.

For the unit-cost exercise, both total expenditure and cost per tonne are indexed to the yield-critical heat measure $\bar{\theta}_{r,c,\tau^*(c),m}$, so that the expenditure and output channels refer to the same underlying heat realization.[8]

In the empirical application, κ_c is estimated directly from the three-wave CECPA sample rather than imposing $\kappa_c = \eta_c^C - \beta_c$ across datasets. The long regional panel is used to estimate β_c because the CECPA sample contains only three survey waves. The decomposition therefore provides an accounting interpretation of the expenditure and output channels, whereas the direct CECPA regression provides the empirical estimate of the unit-cost response.

4.5 Maladaptation Regression: Yield with Pesticide Interaction

To test Prediction 2.5, I augment the farm-level yield specification with pesticide expenditure and its interaction with yield-relevant heat. For each crop c,

[8]This differs from the input-category regressions in Section 4.3, which use the pest-signal phase $\tau^{\text{pest}}(c)$. The distinction matters for barley, where $\tau^*(c) =$ sowing while $\tau^{\text{pest}}(c) =$ growing. Accordingly, η_c^C measures how total expenditure responds to yield-relevant heat, whereas $\eta_{k,c}$ measures how individual inputs respond to the pest-relevant signal.

$$
\begin{aligned}
\ln y_{i,r,c,m} = {} & \beta_{1,c}\,\bar{\theta}^{\text{yield}}_{r,c,m} + \beta_{2,c}\,\ln x^{\text{pest}}_{i,r,c,m} \\
& + \beta_{3,c}\big(\bar{\theta}^{\text{yield}}_{r,c,m} \times \ln x^{\text{pest}}_{i,r,c,m}\big) \\
& + \mathbf{W}'_{r,c,m}\,\boldsymbol{\gamma} + \mathbf{Q}'_{i,r,c,m}\,\boldsymbol{\mu} + \phi_r + \psi_m + \nu_{i,r,c,m},
\end{aligned}
\tag{4.8}
$$

where $\bar{\theta}^{\text{yield}}_{r,c,m}$ is the three-year moving average of KDD during the crop-specific yield-critical phase $\tau^*(c)$: $\text{KDD}_{25,\text{sow}}$ for oats and barley and $\text{KDD}_{30,\text{grow}}$ for corn and soybeans. $\ln x^{\text{pest}}_{i,r,c,m}$ is log pesticide expenditure per hectare; $\mathbf{W}_{r,c,m}$ contains the regional weather controls; and $\mathbf{Q}_{i,r,c,m}$ contains log output price. Equation (4.8) is estimated separately by crop with region and survey-wave fixed effects and standard errors clustered at the region level.

The coefficient of interest is $\beta_{3,c}$, which measures how the conditional association between pesticide expenditure and yield changes with yield-relevant heat. For a phase-mismatched crop, the empirical diagnostic associated with Prediction 2.5 is $\beta_{3,c} < 0$: the pesticide–yield association deteriorates as yield-critical heat increases because the signal guiding defensive adjustment is imperfectly aligned with the phase in which yield damage occurs. For phase-aligned crops, the phase-mismatch mechanism does not predict a negative interaction; estimates around zero or positive are consistent with the absence of a timing mismatch.

Barley provides the central test because it is the only crop for which the retained phases differ: $\tau^*(c) =$ sowing while $\tau^{\text{pest}}(c) =$ growing. The interaction therefore tests whether the conditional pesticide–yield association deteriorates when heat exposure is high in the earlier yield-critical phase. Oats, corn, and soybeans provide within-sample comparisons because their yield-critical and pest-signal phases coincide.

5 Results

The results are presented in six steps that mirror the empirical design and the theoretical predictions. First, the yield estimates from the long regional panel document phase-specific

heat effects on output. Second, the sown-area regressions examine adjustment along the extensive margin. Third, the input-expenditure regressions characterize intensive-margin adjustment. Fourth, the budget-share estimates examine the composition of expenditure. Fifth, the unit-cost analysis evaluates the cost implications of these adjustments under Proposition 2. Finally, the pesticide–yield interaction provides the central empirical test of the phase-mismatch mechanism in Prediction 2.5.

5.1 Yield effects: heat exposure depresses output across all crops

Table 6 reports the within-group fixed-effects estimates of equation (4.1). The main finding is a consistent negative yield response across all four crops in the sample: one additional phase-relevant KDD lowers yield by roughly 1.5 to 2.6 percent, depending on the crop and the relevant damage threshold. The estimates remain precisely estimated after including region and year fixed effects, phase-specific precipitation and moisture-deficit controls, and preparation-phase frost exposure.

Two patterns deserve particular emphasis. First, the yield penalties are similar for the two warm-season crops: corn and soybeans display semi-elasticities of approximately -0.025. This pattern is consistent with evidence that extreme heat becomes particularly damaging for warm crops during reproductive development and grain filling. For corn, yield responses become sharply negative at temperatures around 30 °C (Lobell et al., 2013; Schlenker and Roberts, 2009). Soybean yields display a similar nonlinear response, with temperatures around and above 30 °C reducing reproductive success and seed formation (Boote et al., 2018; Schlenker and Roberts, 2009). The estimates are therefore consistent with the established sensitivity of both crops to above-threshold heat during their reproductive phase.

Second, the cold-crop estimates, based on the 25 °C threshold, are smaller in absolute value but remain negative and statistically significant. This pattern is consistent with the greater sensitivity of cool-season cereals to unusually warm conditions during early development. Evidence from northern cereal systems shows that cool early-season conditions are favorable

to both oats and barley and that barley is particularly sensitive to warmer conditions during early growth (Hakala et al., 2020). Consistent with this timing, the phase-selection exercise retains sowing-phase KDD for both crops, consistent with early-season heat exposure being particularly informative for subsequent yield losses.

More broadly, the yield estimates are consistent with Prediction 2.3 and the incomplete-adaptation mechanism developed in Section 2.6. After conditioning on phase-specific weather controls, region fixed effects, and year fixed effects, yield remains negatively and significantly associated with yield-critical heat for all four crops. This persistence is consistent with evidence that observed agricultural adjustment does not fully eliminate the productivity effects of extreme heat (Burke and Emerick, 2016). In the framework of Proposition 1, the negative reduced-form yield response is consistent with behavioral adjustment failing to fully offset direct heat damage.

Appendix D.1 confirms robustness to alternative KDD thresholds (R1), phase switching (R2), crop-specific regional trends (R3), and within-group pooling (R4, Table 18).

5.2 Extensive margin: heterogeneous area responses to heat

Table 7 reports the estimates of equation (4.2). For oats and barley, neither current preparation-phase nor sowing-phase KDD is statistically significant. For barley, however, lagged growing-season KDD enters negatively and significantly, indicating that sown area is lower following a year with greater growing-season heat exposure. This pattern is consistent with a retrospective adjustment channel: because the forthcoming growing-season heat shock is not observed at planting, previous-season conditions provide information that may enter current planting decisions. For oats, neither current nor lagged KDD is statistically distinguishable from zero.

Corn displays a different pattern. Both preparation- and sowing-phase KDD enter positively and significantly, with the preparation-phase coefficient providing the cleaner pre-planting signal because it is observed before the planting commitment. This extensive-margin

association is not inconsistent with the negative yield effect of yield-critical growing-season heat: the two estimates refer to distinct phenological windows and distinct margins of adjustment.

Soybean area shows no statistically significant association with any of the three KDD measures or with lagged yield. Taken together, the estimates point to crop-specific extensive-margin responses rather than a uniform acreage response to heat: barley exhibits a retrospective contraction following greater growing-season heat in the previous year, corn area is positively associated with current early-season warmth, and oats and soybeans show no statistically detectable response.

5.3 Input-expenditure responses are crop- and input-specific

Table 8 reports the input-expenditure estimates from equation (4.3). Responses differ sharply across crops and input categories. For oats and barley, the dominant pattern is contraction in frontier-expanding inputs. Fertilizer expenditure falls significantly for both oats ($\hat{\eta}_{\text{fert,oats}} = -0.493$, $p = 0.002$) and barley ($\hat{\eta}_{\text{fert,barley}} = -0.034$, $p < 0.001$), while seed expenditure also declines ($\hat{\eta}_{\text{seed,oats}} = -0.208$, $p = 0.050$, and $\hat{\eta}_{\text{seed,barley}} = -0.020$, $p < 0.001$). These responses are consistent with frontier disengagement in Lemma 2.2: as the surviving output share contracts, the marginal return to frontier investment can fall. The pesticide margin, however, differs across the two crops. For barley, pesticide expenditure rises significantly with growing-phase KDD ($\hat{\eta}_{\text{pest,barley}} = 0.050$, $p < 0.001$), consistent with the defensive-intensification channel in Lemma 2.1. Oats, by contrast, show no statistically significant pesticide response. Barley therefore exhibits the Regime 1 configuration of pesticide intensification combined with frontier disengagement, while oats lies near its boundary with a negligible pesticide response. The positive barley pesticide response is also consistent with evidence of increased pesticide expenditure under within-season weather shocks among French cereal producers (Bareille et al., 2024).

Corn displays a different adjustment regime. Pesticide ($\hat{\eta}_{\text{pest,corn}} = 0.617$, $p = 0.004$),

fertilizer ($\hat{\eta}_{\text{fert,corn}} = 0.193$, $p = 0.014$), and seed expenditure ($\hat{\eta}_{\text{seed,corn}} = 0.081$, $p < 0.001$) all increase with growing-phase heat. The pesticide response is consistent with defensive intensification, while the simultaneous increase in fertilizer and seed expenditure indicates that frontier investment also expands rather than contracts. Corn therefore corresponds to Regime 2, in which the damage-control and recovery-demand channels generate adjustment in the same direction.

Soybeans exhibit the reverse cross-input configuration. Pesticide expenditure falls sharply ($\hat{\eta}_{\text{pest,soy}} = -0.300$, $p < 0.001$), while fertilizer ($\hat{\eta}_{\text{fert,soy}} = 0.089$, $p < 0.001$) and seed expenditure ($\hat{\eta}_{\text{seed,soy}} = 0.169$, $p < 0.001$) both increase. This is selective reallocation rather than broad retrenchment: the damage-control margin contracts while frontier expenditure expands. In the model, this sign configuration corresponds to Regime 4: the yield-value channel dominates pesticide adjustment while the recovery-demand channel supports frontier expenditure. Soybean nitrogen physiology provides one possible interpretation of the fertilizer response: high temperatures can impair biological nitrogen fixation, potentially increasing the value of purchased nutrient inputs (Hungria and Vargas, 2000; Salvagiotti et al., 2008). Because the CECPA regressions identify expenditure responses rather than the underlying agronomic adjustment itself, this mechanism should be interpreted as consistent with the estimated fertilizer response rather than as directly identified in the data.

Taken together, the estimates show that input adaptation to heat is not uniform across crops or input categories. Oats primarily reduce frontier expenditure, barley combines defensive intensification with frontier disengagement, corn expands all three input categories, and soybeans reallocate expenditure away from pesticides and toward frontier inputs. These crop-specific sign patterns correspond closely to the adaptation regimes derived in Section 2.5.

5.4 Budget shares: heat reallocates the composition of expenditure

Table 9 complements the expenditure-level results by examining each input as a share of total production expenditure. Changes in these shares therefore capture shifts in expenditure

composition rather than proportional changes in all spending categories.

For oats, the pesticide (-0.012, $p = 0.004$), fertilizer (-0.085, $p < 0.001$), and seed (-0.041, $p = 0.020$) shares all decline with sowing-phase heat. The expenditure contraction documented in Section 5.3 is therefore accompanied by a decline in the budget shares of all three measured variable inputs.

Barley exhibits a different composition response. The pesticide share rises ($+0.005$, $p < 0.001$), the fertilizer share declines (-0.001, $p < 0.001$), and the seed share increases ($+0.003$, $p < 0.001$). The budget therefore shifts away from fertilizer and toward both pesticides and seeds. In particular, the increase in the pesticide share reinforces the expenditure-level evidence of defensive intensification, while the decline in the fertilizer share is consistent with frontier disengagement.

For corn, all three input expenditures increase in levels, but their shares move in different directions. The pesticide share rises ($+0.011$, $p = 0.021$), whereas the fertilizer (-0.030, $p = 0.012$) and seed (-0.034, $p < 0.001$) shares decline. Heat is therefore associated not only with broad input intensification but also with a relative shift in expenditure composition toward pesticides.

Soybeans display the opposite compositional pattern. The pesticide share falls sharply (-0.055, $p < 0.001$), while the fertilizer ($+0.025$, $p < 0.001$) and seed ($+0.007$, $p = 0.003$) shares increase. This mirrors the expenditure-level results: soybean adjustment involves a relative shift away from the damage-control margin and toward frontier inputs.

Taken together, the share regressions sharpen the expenditure-level evidence: heat changes not only the level of input spending but also its composition, and the direction of that reallocation differs across crops.

5.5 Cost per tonne: expenditure and output channels

Table 10 reports the expenditure and unit-cost responses to yield-critical heat. The estimated unit-cost coefficient is positive for all four crops and statistically significant for oats, barley,

and soybeans; the corn estimate is positive but imprecise. Conceptually, the accounting identity $\ln(C/y) = \ln C - \ln y$ separates the unit-cost response into an expenditure channel and an output channel. Because these channels are estimated using different samples, this decomposition is used for interpretation rather than imposed econometrically: the unit-cost coefficient $\hat{\kappa}_c$ is estimated directly in the CECPA sample, while the long regional panel provides the corresponding yield response.

For oats, total expenditure rises significantly with sowing-phase heat ($\hat{\eta}^C_{\text{oats}} = 0.289$, $p < 0.001$), while the long-panel yield estimate is negative. Both accounting channels therefore point toward higher cost per tonne, consistent with the positive direct unit-cost estimate ($\hat{\kappa}_{\text{oats}} = 0.395$, $p = 0.010$).

For barley, total expenditure rises sharply with yield-critical sowing-phase heat ($\hat{\eta}^C_{\text{barley}} = 0.776$, $p < 0.001$), while the long-panel yield response to the same yield-critical phase is negative. Both accounting channels therefore point toward higher unit cost, consistent with the positive and statistically significant direct effect ($\hat{\kappa}_{\text{barley}} = 0.227$, $p < 0.001$). This expenditure response should be distinguished from the pesticide response in Section 5.3, which is associated with growing-phase heat, the retained pest-signal phase.

For corn, total expenditure rises significantly ($\hat{\eta}^C_{\text{corn}} = 0.280$, $p = 0.042$), while the long-panel yield response is negative. Both accounting channels therefore point toward higher unit cost. The directly estimated unit-cost coefficient is positive ($\hat{\kappa}_{\text{corn}} = 0.137$, $p = 0.603$) but statistically imprecise, so the reduced-form unit-cost effect cannot be distinguished from zero.

For soybeans, the total-expenditure coefficient is positive but statistically indistinguishable from zero ($\hat{\eta}^C_{\text{soy}} = 0.245$, $p = 0.103$), whereas the long-panel yield response is negative and statistically significant. The direct unit-cost estimate is positive and significant ($\hat{\kappa}_{\text{soy}} = 0.488$, $p = 0.009$), consistent with the output channel playing an important role in the increase in cost per tonne.

Taken together, the estimates indicate positive unit-cost responses for all four crops, with statistically clear effects for oats, barley, and soybeans and a positive but imprecise estimate

for corn. The pattern is consistent with Proposition 2: observed expenditure adjustment does not necessarily prevent persistent heat-related yield losses from translating into higher production costs per unit of output. This incomplete protection is consistent with evidence that agricultural adaptation attenuates but does not fully eliminate the productivity effects of adverse climate conditions (Bareille and Chakir, 2024; Burke and Emerick, 2016).

5.6 Phase mismatch: a significant negative interaction only for barley

Table 11 reports $\hat{\beta}_{3,c}$ from equation (4.8), which measures how the conditional association between pesticide expenditure and yield changes with yield-relevant heat. The negative interaction is statistically significant only for barley, the sole crop for which the retained yield-critical and pesticide-response phases differ.

For barley, the interaction between yield-relevant sowing-phase KDD and log pesticide expenditure is negative and statistically significant ($\hat{\beta}_{3,\text{barley}} = -0.090$, $p = 0.015$). Thus, as sowing-phase heat increases, the conditional pesticide–yield association becomes less favorable. This pattern is consistent with the phase-mismatch diagnostic in Prediction 2.5. The retained yield-critical measure for barley is $\text{KDD}_{25,\text{sow}}$, whereas pesticide expenditure is most strongly associated with $\text{KDD}_{25,\text{grow}}$. The negative interaction therefore appears precisely for the crop in which the heat signal associated with defensive expenditure is imperfectly aligned with the heat phase most strongly associated with yield damage.

Figure 2 illustrates the yield side of this phase contrast in the long regional panel. Residualized barley yield declines more strongly with sowing-phase KDD than with growing-phase KDD, with slopes of approximately -0.015 and -0.005, respectively. The figure therefore complements the interaction regression by showing that the phase most strongly associated with pesticide expenditure is not the phase to which barley yield is most sensitive per degree-day.

The cross-crop comparison reinforces the selectivity of the result. For oats and corn,

the interaction estimates are negative but statistically indistinguishable from zero ($\hat{\beta}_{3,\text{oats}} = -0.087$, $p = 0.467$; and $\hat{\beta}_{3,\text{corn}} = -0.080$, $p = 0.471$). For soybeans, the interaction is positive and significant ($\hat{\beta}_{3,\text{soy}} = 0.034$, $p = 0.052$). In all three crops, the retained pesticide-response and yield-critical phases coincide. The theory imposes no sign restriction on $\beta_{3,c}$ for phase-aligned crops. Accordingly, the statistically null estimates for oats and corn and the positive estimate for soybeans provide no evidence of the systematic negative interaction associated with the phase-mismatch mechanism. A significant negative interaction appears only for barley, where $\tau^{\text{pest}} \neq \tau^{*}$.

The barley result is also robust to alternative thermal specifications and to tests designed to discriminate between the two phases. Using a higher heat threshold strengthens the negative interaction, and replacing the three-year KDD average with annual KDD leaves the estimate essentially unchanged, as shown in Appendix Table 20. By contrast, replacing the sowing-phase interaction with the growing-phase interaction reduces the coefficient to approximately zero ($\hat{\beta}'_3 = -0.002$, $p_{\text{WCB}} = 0.769$). When both phase interactions enter jointly, the sowing-phase interaction remains negative ($\hat{\delta}_s = -0.145$, $p_{\text{WCB}} = 0.070$), whereas the growing-phase interaction is small and statistically insignificant ($\hat{\delta}_g = 0.015$, $p_{\text{WCB}} = 0.278$). The equality test yields $p = 0.041$ under conventional cluster-robust inference and $p = 0.063$ under wild-cluster bootstrap inference; full results appear in Appendix Table 21. Finally, all six leave-one-region-out barley estimates remain negative, and five remain statistically significant, including an estimate of -0.295 ($p = 0.014$) when the dominant barley region is excluded, as documented in Appendix Table 19. Taken together, these checks support the interpretation that the barley interaction is specific to the sowing–growing phase discordance rather than to a generic negative relationship between pesticide expenditure and yield.

6 Robustness

The main results are robust to alternative definitions of heat exposure, phase specification, temporal aggregation, and controls for regional heterogeneity. Full estimates are reported in Appendix D.1 for the yield regressions and Appendix D.2 for the phase-mismatch analysis.

For the long regional panel, the negative yield responses are not sensitive to the particular KDD specification. Raising the damage thresholds from 25 to 27 °C for cold crops and from 30 to 31 °C for warm crops leaves all four yield estimates negative and statistically significant. Replacing each crop's retained phenological phase with the alternative phase weakens the yield relationship, supporting the phase-selection procedure. The estimates also remain stable after adding crop-specific regional linear trends and when cold- and warm-season crops are pooled within their respective groups. These exercises are reported in Appendix Tables 15–18.

The phase-mismatch result is similarly insensitive to alternative empirical choices. As summarized in Section 5.6, the negative barley interaction persists under a higher KDD threshold and when annual rather than three-year average heat exposure is used. More importantly, the coefficient is not statistically distinguishable from zero when the interaction is indexed to growing-phase heat, whereas the sowing-phase interaction remains negative when the two phase interactions enter jointly. Leave-one-region-out estimates are also uniformly negative. Full results, including the joint phase-discrimination test and wild-cluster bootstrap inference, are reported in Appendix D.2.

Taken together, these exercises indicate that the main findings are not driven by a particular temperature threshold, temporal aggregation rule, phase assignment, regional trend, or single region. The phase-specific contrast underlying the barley result is especially stable: the negative interaction is associated with the yield-critical sowing phase rather than with the growing-phase signal that is most strongly associated with pesticide expenditure.

7 Climate Projections

Phase-specific heat stress is associated with lower yields, heterogeneous input-expenditure adjustment across crops, and positive projected unit-cost changes. The underlying unit-cost estimates are statistically significant for oats, barley, and soybeans, but imprecise for corn. Combining the estimated yield and unit-cost semi-elasticities with CanDCS-M6 projections (Sobie et al., 2024) under SSP2–4.5, I quantify the near-term implications of additional heat exposure over 2025–2030. The projected increments in the yield-critical heat measures are $\Delta\text{KDD}_{25,\text{sow}} = 1.52$ for oats and barley and $\Delta\text{KDD}_{30,\text{grow}} = 2.65$ for corn and soybeans. The analysis focuses on this near-term horizon to limit extrapolation relative to more distant climate scenarios.

Table 12 reports the corresponding point projections. For oats and barley, projected yield declines are 3.4 and 2.3 percent, respectively, while projected unit costs increase by 82 and 41 percent. These patterns are consistent with the accounting evidence in Section 5.5: for both crops, total expenditure rises with yield-critical heat while output falls, so the expenditure and output channels point in the same direction.

For corn, the point projection implies a 6.1 percent yield decline and a 44 percent unit-cost increase. The latter should be interpreted cautiously, however, because the underlying unit-cost coefficient is statistically imprecise. For soybeans, the projected yield decline is 6.6 percent and the unit-cost point projection is 264 percent. The soybean estimate does not imply a corresponding increase in out-of-pocket expenditure: pesticide expenditure falls while fertilizer and seed expenditures rise, and the estimated total-expenditure response is statistically indistinguishable from zero. The large unit-cost projection therefore reflects an important output-side contribution to the increase in cost per tonne rather than a statistically identified increase in total expenditure.

Extensive-margin projections are not reported because equation (4.2) combines current and lagged climate signals with lagged yield, so projecting sown area would require additional assumptions about the dynamic paths of climate exposure and agricultural outcomes.

8 Conclusion

This paper studies how phase-specific heat stress affects agricultural production and adaptation across four major crops in Quebec. Combining a long regional panel of yields and sown area with farm-level expenditure data, I embed phase-specific climate risk in a damage-control production framework that distinguishes pesticides from frontier-expanding inputs. The central insight is that adaptation depends not only on whether farmers adjust, but also on whether the signal guiding that adjustment is aligned with the phenological phase in which yield damage occurs. When these phases differ, adjustment associated with one climate signal may provide limited protection against heat exposure in the phase most strongly associated with output.

The empirical results first document substantial residual heat damage. Yield falls significantly with yield-critical KDD for all four crops, with semi-elasticities ranging from approximately -0.015 to -0.026 per degree-day. The relevant phase differs across crop groups: sowing-phase heat is retained for oats and barley, whereas growing-phase heat is retained for corn and soybeans. These negative reduced-form responses persist after conditioning on region and year fixed effects and phase-specific weather controls, and are robust to alternative thresholds, phase definitions, and regional trends. In the framework of Proposition 1, this persistence is consistent with observed adjustment failing to fully offset direct heat-related damage.

Adjustment itself is highly heterogeneous. Oats reduce fertilizer and seed expenditure without a statistically detectable pesticide response. Barley combines higher pesticide expenditure with lower fertilizer and seed expenditure, corresponding to the defensive-intensification and frontier-disengagement configuration of Regime 1. Corn instead increases expenditure on all three inputs, consistent with broad intensification. Soybeans display the reverse cross-input configuration: pesticide expenditure falls while fertilizer and seed expenditure rise, indicating selective reallocation away from the damage-control margin and toward frontier inputs. The budget-share estimates reinforce these differences in expenditure

composition. Taken together, these results show that adaptation cannot be characterized by a uniform increase or decrease in input use; both the direction and composition of adjustment depend on the crop and input margin.

The extensive margin is similarly crop-specific. Barley sown area is lower following greater growing-season heat in the previous year, consistent with a retrospective adjustment channel. Corn area is positively associated with current early-season warmth, with preparation-phase heat providing the cleanest pre-planting signal. Oats and soybeans show no statistically detectable area response. These results suggest that extensive-margin adjustment provides an additional, but heterogeneous, response to climate conditions rather than a common mechanism operating across crops.

Despite these behavioral responses, the estimated unit-cost coefficient is positive for all four crops and statistically significant for oats, barley, and soybeans; the corn estimate is positive but imprecise. For oats and barley, total expenditure rises with yield-critical heat while yield falls, so the expenditure and output channels both point toward higher cost per tonne. For soybeans, the total-expenditure response is statistically indistinguishable from zero while the yield response remains negative, consistent with the output channel playing an important role in the increase in unit cost. These findings are consistent with Proposition 2: input-expenditure adjustment need not prevent residual output losses from translating into higher production costs per unit of output.

The strongest evidence for the phase-mismatch mechanism comes from barley. It is the only crop for which the retained pesticide-response phase differs from the yield-critical phase: pesticide expenditure is most strongly associated with growing-phase KDD, whereas yield is most sensitive to sowing-phase KDD. Correspondingly, barley is the only crop with a statistically significant negative interaction between yield-critical heat and pesticide expenditure ($\hat{\beta}_{3,\text{barley}} = -0.090$, $p = 0.015$). As sowing-phase heat increases, the conditional pesticide–yield association therefore becomes less favorable. Oats and corn exhibit negative but statistically insignificant interactions, while the soybean interaction is positive and

significant; all three are phase-aligned crops. The result is also phase-specific within barley: replacing the sowing-phase interaction with growing-phase heat reduces the coefficient to approximately zero, and a joint specification finds a substantially more negative interaction for sowing than for growing heat. This selectivity is consistent with the diagnostic prediction of Proposition 3: the negative pesticide–yield interaction is statistically evident precisely where the signal associated with defensive expenditure is imperfectly aligned with the phase most strongly associated with yield damage.

The policy implication is therefore not simply that farmers should use more or fewer inputs. The results instead emphasize the value of matching adaptation decisions to crop-specific biological timing. Where the yield-critical and defensive-signal phases are aligned, adjustment can take very different forms across crops without generating the negative phase-mismatch diagnostic observed for barley. Where they are misaligned, policies focused solely on expanding input access may overlook the timing and information problem. Phase-specific forecasts and agronomic decision support are therefore natural interventions to investigate, although the present estimates do not identify the productivity or welfare effects of such policies directly.

The near-term climate projections illustrate the potential economic relevance of these estimated responses. Applying the yield and unit-cost semi-elasticities to five-model CanDCS-M6 ensemble-mean changes under SSP2–4.5 for 2025–2030 yields projected yield declines of 2.3 to 6.6 percent across the four crops. The corresponding unit-cost point projections range from 41 percent for barley to 264 percent for soybeans. These figures are scenario-based transformations of the estimated semi-elasticities rather than probabilistic forecasts, and the corn unit-cost projection should be interpreted particularly cautiously because its underlying coefficient is statistically imprecise. The large soybean projection likewise reflects an important output-side contribution to cost per tonne rather than a statistically identified increase in total expenditure.

Several limitations define the scope for future work. The CECPA data contain only

three survey waves, limiting the ability to trace adjustment dynamically within farms and to study slow-moving responses such as variety substitution, crop rotation, or capital investment. Annual expenditure data also do not reveal the exact within-season timing of pesticide applications, so the phase-mismatch interpretation relies on the contrast between the climate phase most strongly associated with pesticide expenditure and the phase most strongly associated with yield damage. In addition, the analysis is confined to Quebec, and the magnitude of these relationships may differ in warmer agricultural regions. Barley observations are geographically concentrated, although the negative interaction remains stable in leave-one-region-out exercises.

Higher-frequency farm data would allow the timing mechanism to be tested more directly by linking individual input applications to realized phenological conditions. A complementary structural model of crop choice under phase-specific climate uncertainty could quantify how farmers form beliefs about future heat exposure and how those beliefs affect planting and within-season input decisions. More broadly, extending the analysis across climates and production systems would establish whether the phase mismatch documented here is specific to barley in Quebec or represents a more general constraint on agricultural adaptation. The broader lesson is that adaptation should be evaluated not only by whether producers respond to climate shocks, but by whether those responses are directed toward the biological processes that ultimately determine production.

References

Bareille, F. and Chakir, R. (2024). Structural identification of weather impacts on crop yields: disentangling agronomic from adaptation effects. *American Journal of Agricultural Economics*, 106(3):989–1019.

Bareille, F., Chakir, R., and Keles, D. (2024). Weather shocks and pesticide purchases. *European Review of Agricultural Economics*, 51(2):309–353.

Barnett, J. and O'Neill, S. (2010). Maladaptation. *Global Environmental Change*, 20(2):211–213.

Blanc, E. and Schlenker, W. (2017). The use of panel models in assessments of climate impacts on agriculture. *Review of Environmental Economics and Policy*, 11(2):258–279.

Boote, K. J., Prasad, P. V. V., Allen, L. H., Singh, P., and Jones, J. W. (2018). Modeling sensitivity of grain yield to elevated temperature in the DSSAT crop models for peanut, soybean, dry bean, chickpea, sorghum, and millet. *European Journal of Agronomy*, 100:99–109.

Burke, M. and Emerick, K. (2016). Adaptation to climate change: Evidence from US agriculture. *American Economic Journal: Economic Policy*, 8(3):106–140.

Cameron, A. C., Gelbach, J. B., and Miller, D. L. (2008). Bootstrap-based improvements for inference with clustered errors. *Review of Economics and Statistics*, 90(3):414–427.

Carpentier, A. and Letort, E. (2012). Accounting for heterogeneity in multicrop micro-econometric models: implications for variable input demand modeling. *American Journal of Agricultural Economics*, 94(1):209–224.

Centre for Studies on Agricultural Production Costs (CECPA) (2021). Surveys on production costs of conventional and organic crops in Quebec: 2009, 2014, 2019. Data collected through farm-level surveys conducted in various administrative regions of Quebec, covering both conventional and organic crop production.

Deschênes, O. and Greenstone, M. (2007). The economic impacts of climate change: evidence from agricultural output and random fluctuations in weather. *American Economic Review*, 97(1):354–385.

Fisher, A. C., Hanemann, W. M., Roberts, M. J., and Schlenker, W. (2012). The economic impacts of climate change: evidence from agricultural output and random fluctuations in weather: comment. *American Economic Review*, 102(7):3749–3760.

Hakala, K., Jauhiainen, L., Rajala, A. A., Jalli, M., Kujala, M., and Laine, A. (2020). Different responses to weather events may change the cultivation balance of spring barley and oats in the future. *Field Crops Research*, 259:107956.

Hatfield, J. L., Boote, K. J., Kimball, B. A., Ziska, L. H., Izaurralde, R. C., Ort, D., Thomson, A. M., and Wolfe, D. (2011). Climate impacts on agriculture: implications for crop production. *Agronomy Journal*, 103(2):351–370.

Hatfield, J. L. and Prueger, J. H. (2015). Temperature extremes: effect on plant growth and development. *Weather and Climate Extremes*, 10:4–10.

Hungria, M. and Vargas, M. A. T. (2000). Environmental factors affecting N2 fixation in grain legumes in the tropics, with an emphasis on Brazil. *Field Crops Research*, 65(2-3):151–164.

Institut de la statistique du Québec (ISQ) and Statistics Canada (2023). Crop yields per hectare, by grouped administrative regions, Quebec, 1990–2022. https://www.stat.gouv.qc.ca. Data compiled by ISQ from Statistics Canada sources.

Jagnani, M., Barrett, C. B., Liu, Y., and You, L. (2021). Within-season producer response to warmer temperatures: Defensive investments by Kenyan farmers. *The Economic Journal*, 131(633):392–419.

Kawasaki, K. (2023). Impact of climate change on crop pests and diseases: ensemble modeling of time-varying weather effects. *Journal of the Association of Environmental and Resource Economists*, 10(6):1515–1543.

Lichtenberg, E. and Zilberman, D. (1986). The econometrics of damage control: why specification matters. *American Journal of Agricultural Economics*, 68(2):261–273.

Lobell, D. B., Hammer, G. L., McLean, G., Messina, C., Roberts, M. J., and Schlenker, W. (2013). The critical role of extreme heat for maize production in the United States. *Nature Climate Change*, 3(5):497–501.

MacKinnon, J. G. and Webb, M. D. (2017). Wild bootstrap inference for wildly different cluster sizes. *Journal of Applied Econometrics*, 32(2):233–254.

NASA Langley Research Center POWER Project (2026). POWER CERES/MERRA-2 Native Resolution Daily Data. https://power.larc.nasa.gov. Daily maximum temperature, minimum temperature, precipitation, and reference evapotranspiration extracted via the NASA POWER Daily API for representative coordinates in each Quebec agricultural region, 1991–2022. Accessed March 2026.

Porter, J. R. and Gawith, M. (1999). Temperatures and the growth and development of wheat: a review. *European Journal of Agronomy*, 10(1):23–36.

Rezaei, E. E., Webber, H., Gaiser, T., Naab, J., and Ewert, F. (2015). Heat stress in cereals: mechanisms and modelling. *European Journal of Agronomy*, 64:98–113.

Rodrigues, G. C. and Braga, R. P. (2021). Evaluation of NASA POWER reanalysis products to estimate daily weather variables in a hot summer Mediterranean climate. *Agronomy*, 11(6):1207.

Salvagiotti, F., Cassman, K. G., Specht, J. E., Walters, D. T., Weiss, A., and Dobermann, A. (2008). Nitrogen uptake, fixation and response to fertilizer N in soybeans: A review. *Field Crops Research*, 108(1):1–13.

Schlenker, W. and Roberts, M. J. (2009). Nonlinear temperature effects indicate severe damages to US crop yields under climate change. *Proceedings of the National Academy of Sciences*, 106(37):15594–15598.

Sexton, S. E., Lei, Z., and Zilberman, D. (2007). The economics of pesticides and pest control. *International Review of Environmental and Resource Economics*, 1(3):271–326.

Sobie, S. R., Ouali, D., Curry, C. L., and Zwiers, F. W. (2024). Multivariate Canadian downscaled climate scenarios for CMIP6 (CanDCS-M6). *Geoscience Data Journal*, 11:806–824.

White, J. W., Hoogenboom, G., Stackhouse, P. W., and Hoell, J. M. (2008). Evaluation of NASA satellite- and assimilation model-derived long-term daily temperature data over the continental US. *Agricultural and Forest Meteorology*, 148(10):1574–1584.

A Theoretical Appendix

A.1 Summary of Testable Predictions

The theoretical framework generates five testable predictions that map directly onto the empirical architecture of the paper. The extensive-margin analysis of Section 5.2 is an additional empirical contribution not formalised as a separate theoretical prediction. These are summarized in Table 13.

A.2 Proofs of Lemmas and Propositions

The proofs proceed in order of appearance in the text. The shorthand used is $H_p \equiv \partial H/\partial x_{\text{pest}}$, $H_\theta \equiv \partial H/\partial \theta^{\text{pest}}$, and $H_{p\theta} \equiv \partial^2 H/\partial x_{\text{pest}}\,\partial \theta^{\text{pest}}$. By (A2): $H_p < 0$ and $\partial^2 H/\partial x_{\text{pest}}^2 \geq 0$. By (A3): $H_{p\theta} \geq 0$.

Proof of Lemma 2.1 (Defensive intensification).

In compact notation, the pesticide FOC (2.5) is:

$$\mathcal{F}(x_{\text{pest}},\, s) \;\equiv\; P \cdot M(\mathbf{f}, s) \cdot \Phi_x(x_{\text{pest}}, s) - w_{\text{pest}} \;=\; 0,$$

where $s = \theta^{\text{pest}}$, $M(\mathbf{f}, s) = \mathbb{E}[A \mid s] > 0$, and $\Phi_x = D'(N)|H_p| > 0$. By the implicit function theorem:

$$\frac{dx^*_{\text{pest}}}{ds} = -\frac{\partial \mathcal{F}/\partial s}{\partial \mathcal{F}/\partial x_{\text{pest}}}.$$

Step 1: Sign of the denominator $\partial \mathcal{F}/\partial x_{pest}$.

$$\frac{\partial \mathcal{F}}{\partial x_{\text{pest}}} = P \cdot M \cdot \Phi_{xx}, \quad \text{where} \quad \Phi_{xx} = -D''(N)H_p^2 - D'(N)\frac{\partial^2 H}{\partial x_{\text{pest}}^2} \leq 0,$$

since $D'' > 0$ (A4), $H_p^2 > 0$, $D' > 0$, and $\partial^2 H/\partial x_{\text{pest}}^2 \geq 0$ (A2). At least one term is strictly negative (SOC), so $\partial \mathcal{F}/\partial x_{\text{pest}} = P \cdot M \cdot \Phi_{xx} < 0$.

Step 2: Sign of the numerator $\partial \mathcal{F}/\partial s$.

Differentiating the product $M(f, s) \cdot \Phi_x(x, s)$ with respect to s:

$$\frac{\partial \mathcal{F}}{\partial s} = P\left[\underbrace{M_s \cdot \Phi_x}_{\text{yield-value channel}} + \underbrace{M \cdot \Phi_{xs}}_{\text{damage-control channel}}\right] = P\, q_p,$$

where $M_s \equiv \partial \mathbb{E}[A \mid s]/\partial s$ and $\Phi_{xs} \equiv D''(N)H_\theta|H_p| - D'(N)H_{p\theta}$.

Damage-control channel. Since $D'' > 0$, $H_\theta > 0$, $|H_p| > 0$, the first term in Φ_{xs} is strictly positive. The second term $D'(N)H_{p\theta} \geq 0$ by (A3). Under the sufficient condition $D''H_\theta|H_p| >$

$D'H_{p\theta}$, $\Phi_{xs} > 0$ and therefore $M\Phi_{xs} > 0$.

Yield-value channel. $M_s = \partial\mathbb{E}[A(\theta^{\text{yield}}, \mathbf{f}) \mid \theta^{\text{pest}} = s]/\partial s$. When $\rho_c \approx 0$ (phases uncorrelated), the pest signal carries no information about θ^{yield}, so $M_s \approx 0$. When $\rho_c > 0$, the maintained positive cross-phase association implies that higher pest-signal heat predicts higher yield-critical heat. Combined with (A1), this lowers expected potential yield, so $M_s \leq 0$, and the yield-value channel pulls q_p negative with force $|M_s|\Phi_x$.

Step 3: Conclusion.

$$\text{sign}\left(\frac{dx^*_{\text{pest}}}{ds}\right) = -\frac{\text{sign}(P\, q_p)}{\text{sign}(P\, M\, \Phi_{xx})} = \text{sign}(q_p).$$

Under the sufficient condition ($\Phi_{xs} > 0$ and $M_s \geq 0$), $q_p > 0$ and $dx^*/ds > 0$. When $M_s < 0$ (correlated phases) and the yield-value channel dominates ($|M_s|\Phi_x > M\Phi_{xs}$), $q_p < 0$ and retrenchment follows. □

Proof of Lemma 2.2 (Frontier disengagement).

In compact notation, the frontier-input FOC (2.6) is:

$$\Psi(f_k,\, s) \;\equiv\; P \cdot M_{f_k}(\mathbf{f}, s) \cdot \Phi(x_{\text{pest}}, s) - w_k \;=\; 0,$$

where $s = \theta^{\text{pest}}$, $M_{f_k} \equiv \partial M/\partial f_k = \mathbb{E}[\partial A/\partial f_k \mid s] > 0$ under (A5), and $\Phi(x_{\text{pest}}, s) = 1 - D(H(s, x_{\text{pest}}, \mathbf{z})) > 0$. By the implicit function theorem:

$$\frac{df_k^*}{ds} = -\frac{\partial\Psi/\partial s}{\partial\Psi/\partial f_k}.$$

Step 1: Sign of the denominator $\partial\Psi/\partial f_k$.

$$\frac{\partial\Psi}{\partial f_k} = P \cdot M_{f_k f_k} \cdot \Phi \;\leq\; 0,$$

since $M_{f_k f_k} = \mathbb{E}[\partial^2 A/\partial f_k^2 \mid s] \leq 0$ under (A5) and $\Phi > 0$. Strictly negative when $\partial^2 A/\partial f_k^2 <$

0.

Step 2: Sign of the numerator $\partial\Psi/\partial s$.

Differentiating Ψ with respect to s yields two channels:

$$\frac{\partial \Psi}{\partial s} = P\left[\underbrace{M_{f_k s}\,\Phi}_{\text{recovery-demand channel}} + \underbrace{M_{f_k}\,\Phi_s}_{\text{survival-share channel}}\right] = P\,q_k,$$

where $M_{f_k s} \equiv \partial M_{f_k}/\partial s = \partial\,\mathbb{E}[A_{f_k}(r,\mathbf{f}) \mid s]/\partial s$ captures how the pest-signal heat updates beliefs about the yield-relevant state, thereby affecting the expected marginal return to frontier inputs, and $\Phi_s = -D'(N)H_\theta < 0$ (since $D'(N) > 0$ by definition of D and $H_\theta > 0$ under (A3)).

The survival-share term $M_{f_k}\Phi_s < 0$ is always negative: high pest-signal heat shrinks the fraction of potential yield that survives, reducing the return to frontier investment. The recovery-demand term $M_{f_k s}\Phi$ is positive when heat raises the marginal product of frontier inputs through compensatory mechanisms ($M_{f_k s} > 0$, a case permitted by the unrestricted cross-partial in (A5)).

Step 3: Conclusion.

By the implicit function theorem,

$$\operatorname{sign}\left(\frac{df_k^*}{ds}\right) = \operatorname{sign}(q_k) = \operatorname{sign}(M_{f_k s}\,\Phi + M_{f_k}\,\Phi_s)\,.$$

Frontier disengagement ($df_k^*/ds < 0$) holds when $q_k < 0$, i.e., when the survival-share channel dominates the recovery-demand channel ($M_{f_k}|\Phi_s| > M_{f_k s}\Phi$). The recovery-demand regime ($df_k^*/ds > 0$) holds when $q_k > 0$, as empirically observed for corn and soybeans (Section 5.3). □

Proof of Proposition 1 (Incomplete adaptation).

Step 1: Chain-rule decomposition.

Along the equilibrium input rule, the total derivative of realized yield with respect to $\theta^{\text{yield}}_{c,t}$ decomposes as:

$$\frac{dy^*_{c,t}}{d\,\theta^{\text{yield}}_{c,t}} = \left.\frac{\partial y_{c,t}}{\partial\,\theta^{\text{yield}}_{c,t}}\right|_{\mathbf{x}^*} + \nabla_{\mathbf{x}}\, y_{c,t} \cdot \frac{d\mathbf{x}^*_{c,t}}{d\,\theta^{\text{yield}}_{c,t}}.$$

The first term is strictly negative (Step 2 below). Whether the second term offsets it depends on how well the observable pest-phase signal tracks the yield-critical shock; that relationship is the subject of Proposition 3.

Step 2: Sign of $\partial y/\partial\theta^{yield}$ at fixed inputs.

Differentiating the production function (2.2) with respect to $\theta^{\text{yield}}_{c,t}$ at fixed $\mathbf{x} = \mathbf{x}^*_{c,t}$:

$$\left.\frac{\partial y_{c,t}}{\partial\,\theta^{\text{yield}}_{c,t}}\right|_{\mathbf{x}^*} = \underbrace{\frac{\partial A_{c,t}}{\partial\,\theta^{\text{yield}}_{c,t}} \cdot \left[1 - D(N^*_{c,t})\right]}_{<\,0 \text{ by (A1) and } \Phi>0}. \tag{A.1}$$

$\partial A/\partial\theta^{\text{yield}} < 0$ by (A1) and $[1 - D(N^*)] > 0$ by definition of D, so this term is strictly negative. Note that $N^* = H(\theta^{\text{pest}}, \mathbf{x}^*, \mathbf{z})$ does not depend on θ^{yield} when the two phases are distinct; the direct effect operates exclusively through $A(\theta^{\text{yield}}, \mathbf{f})$. Therefore:

$$\left.\frac{\partial y_{c,t}}{\partial\,\theta^{\text{yield}}_{c,t}}\right|_{\mathbf{x}^*} < 0.$$

Step 3: Conclusion.

The direct damage (Step 2) is strictly negative. The proposition holds whenever the endogenous adaptation term fails to offset this direct negative effect, which is the condition stated in Proposition 1 and is evaluated empirically via Prediction 2.3. □

Proof of Proposition 2 (Rising unit cost of production).

Step 1: Log-decomposition of unit cost.

From the definition (2.12), $c_{c,t} = \text{Dep}_{c,t}/y_{c,t}$ where $\text{Dep}_{c,t}$ denotes total expenditure per

hectare. Taking logs and differentiating with respect to $\theta_{c,t}^{\text{yield}}$:

$$\frac{d \ln c_{c,t}}{d \theta^{\text{yield}}} = \underbrace{\frac{d \ln \text{Dep}_{c,t}}{d \theta^{\text{yield}}}}_{\text{expenditure channel, sign varies by crop}} - \underbrace{\frac{d \ln y_{c,t}}{d \theta^{\text{yield}}}}_{< 0 \text{ (Proposition 1)}}. \tag{A.2}$$

Step 2: Sign of $d \ln c_{c,t}/d\theta^{yield}$.

By Proposition 1, $d \ln y_{c,t}/d\theta^{\text{yield}} < 0$, so the output channel $-d \ln y_{c,t}/d\theta^{\text{yield}}$ is strictly positive. The sign of the total-expenditure channel $d \ln \text{Dep}_{c,t}/d\theta^{\text{yield}}$ is not determined by the response of any single input, because total expenditure aggregates input margins that may move in opposite directions.

If $d \ln \text{Dep}_{c,t}/d\theta^{\text{yield}} \geq 0$, the expenditure and output channels both imply a positive unit-cost response. If total expenditure falls, unit cost still rises whenever the proportional yield contraction exceeds the proportional expenditure reduction, which is precisely condition (2.13). Hence $d \ln c_{c,t}/d\theta^{\text{yield}} > 0$ under the condition stated in Proposition 2. □

Proof of Proposition 3 (Phase mismatch).

Let $r = \theta_{c,t}^{\text{yield}}$ and $s = \theta_{c,t}^{\text{pest}}$. Under the linear conditional structure,

$$\mathbb{E}[s \mid r] = \mu_s + \lambda_c(r - \mu_r), \qquad \lambda_c \equiv \frac{\text{Cov}(r, s)}{\text{Var}(r)} = \rho_c \frac{\sigma_s}{\sigma_r}.$$

Because equilibrium input choices respond to s, not directly to r,

$$\frac{d\mathbf{x}^*}{dr} = \lambda_c \frac{d\mathbf{x}^*}{ds}.$$

Part (i). For pesticides,

$$\frac{dx_{\text{pest}}^*}{dr} = \lambda_c \frac{dx_{\text{pest}}^*}{ds}.$$

By condition (b), $dx_{\text{pest}}^*/ds > 0$. Since λ_c is strictly increasing in ρ_c, imperfect phase correlation attenuates the pesticide response relative to perfect phase alignment, establishing

part (i).

Part (ii). Define the yield effect of endogenous within-season input adjustment as

$$\Gamma \;\equiv\; \nabla_{\mathbf{x}} y \cdot \frac{d\mathbf{x}^*}{ds}.$$

When $\Gamma > 0$, the adaptive component associated with variation in the yield-critical shock is

$$\nabla_{\mathbf{x}} y \cdot \frac{d\mathbf{x}^*}{dr} = \lambda_c \, \Gamma.$$

Since λ_c is strictly increasing in ρ_c, this adaptive component is strictly smaller under imperfect phase alignment than at $\rho_c = 1$, establishing part (ii).

Part (iii). As $\rho_c \to 0$, $\lambda_c \to 0$, so

$$\frac{d\mathbf{x}^*}{dr} \to 0 \qquad \text{and} \qquad \lambda_c \, \Gamma \to 0.$$

Adjustment triggered by the pest-phase signal therefore provides no systematic adaptive response to variation in the yield-critical shock. □

A.3 Joint Comparative Statics and the Four-Regime Classification

Setup. Using $M(\mathbf{f}, s)$ and $\Phi(x_{\text{pest}}, s)$ as defined in Section 2.4, the per-hectare expected-profit function is

$$\Pi(x_{\text{pest}}, f_k;\, s) \;=\; P \cdot M(\mathbf{f}, s) \cdot \Phi(x_{\text{pest}}, s) \;-\; w_{\text{pest}}\, x_{\text{pest}} \;-\; w_k\, f_k.$$

For a representative frontier input f_k, the 2×2 Hessian evaluated at the interior optimum (x^*, f_k^*) is

$$J \;=\; \begin{pmatrix} \Pi_{xx} & \Pi_{x f_k} \\ \Pi_{f_k x} & \Pi_{f_k f_k} \end{pmatrix},$$

with

$$\Pi_{xx} = P \cdot M \cdot \Phi_{xx} < 0 \quad \text{(concavity of damage control in } x_{\text{pest}}\text{, Assumptions (A2)–(A4))},$$

$$\Pi_{f_k f_k} = P \cdot M_{f_k f_k} \cdot \Phi < 0 \quad \text{(concavity of } A \text{ in frontier input } k\text{, Assumption (A5))},$$

$$\Pi_{x f_k} = P \cdot M_{f_k} \cdot \Phi_x > 0 \quad \text{(complementarity of pesticide survival and frontier value)}.$$

The second-order conditions require $\det(J) = \Pi_{xx}\Pi_{f_k f_k} - (\Pi_{x f_k})^2 > 0$.

Cramer's rule. Differentiating both first-order conditions totally with respect to $s \equiv \theta^{\text{pest}}$ and collecting terms yields the linear system

$$J \begin{pmatrix} dx^*/ds \\ df_k^*/ds \end{pmatrix} = \begin{pmatrix} -P\, q_p \\ -P\, q_k \end{pmatrix},$$

where q_p and q_k are defined in equations (2.7) and (2.9). Cramer's rule then gives

$$\frac{dx^*}{ds} = \frac{P\big(-q_p\, \Pi_{f_k f_k} + q_k\, \Pi_{x f_k}\big)}{\det(J)}, \tag{A.3}$$

$$\frac{df_k^*}{ds} = \frac{P\big(-q_k\, \Pi_{xx} + q_p\, \Pi_{x f_k}\big)}{\det(J)}. \tag{A.4}$$

Since $P > 0$ and $\det(J) > 0$, the signs are determined by the respective numerators. Setting $A \equiv -\Pi_{f_k f_k} > 0$ and $B \equiv -\Pi_{xx} > 0$ and $C \equiv \Pi_{x f_k} > 0$:

$$\operatorname{sign}\left(\frac{dx^*}{ds}\right) = \operatorname{sign}(A\, q_p - C\, |q_k| \cdot \mathbf{1}_{q_k<0} + C\, q_k \cdot \mathbf{1}_{q_k>0}), \tag{A.5}$$

or more simply: $\operatorname{sign}(dx^*/ds) = \operatorname{sign}(A\, q_p + C\, q_k)$ and $\operatorname{sign}(df_k^*/ds) = \operatorname{sign}(B\, q_k + C\, q_p)$.

Regime classification.

Regime 2 ($q_p > 0$, $q_k > 0$; empirically: corn). Both terms in each numerator are strictly positive: $A\, q_p + C\, q_k > 0$ and $B\, q_k + C\, q_p > 0$. Both inputs increase unambiguously: $dx^*/ds > 0$ and $df_k^*/ds > 0$. Empirically, corn exhibits defensive intensification in pesticides

and recovery-demand expansion in fertilizer and seeds (Section 5.3).

Regime 3 ($q_p < 0$, $q_k < 0$; theoretically possible, no observed crop). Both terms in each numerator are strictly negative: $A\,q_p + C\,q_k < 0$ and $B\,q_k + C\,q_p < 0$. Both inputs fall unambiguously: $dx^*/ds < 0$ and $df_k^*/ds < 0$. This is a broad-retrenchment regime; the decline in pesticide demand overturns the standard defensive-intensification prediction.

Regime 1 ($q_p > 0$, $q_k < 0$; empirically: barley). The two terms in each numerator have opposite signs, so unambiguous predictions require a diagonal-dominance condition:

$$A\,|q_p| > C\,|q_k| \quad \text{and} \quad B\,|q_k| > C\,|q_p|. \tag{A.6}$$

When (A.6) holds, $dx^*/ds > 0$ (defensive intensification) and $df_k^*/ds < 0$ (frontier disengagement). Under the damage-control technology, $C = P \cdot M_{f_k} \cdot \Phi_x$ is a product of first-order terms, while A and B grow with the curvature of damage and the concavity of the production frontier. Condition (A.6) therefore holds whenever the curvature of each instrument's own objective is sufficiently strong relative to their cross- interaction, a standard regularity condition in damage-control models. Empirically, barley exhibits defensive intensification in pesticides and disengagement in fertilizer and seeds (Section 5.3), consistent with (A.6).

Regime 4 ($q_p < 0$, $q_k > 0$; empirically: soybeans). The terms have opposite signs. For pesticides: $A\,q_p + C\,q_k < 0$ when $A\,|q_p| > C\,q_k$ (yield-value channel dominates). For frontier inputs: $B\,q_k + C\,q_p > 0$ when $B\,q_k > C\,|q_p|$ (recovery-demand channel dominates). Pesticide demand falls while frontier input demand rises: the two margins move in opposite directions. Empirically, soybeans exhibit pesticide retrenchment and frontier-input expansion (Section 5.3), consistent with this regime.

B Data Appendix

B.1 KDD Phase Selection

Single-sine computation. For each day d, the threshold exceedance $h(T_{d,t}^{\max}, T_{d,t}^{\min}; \bar{T}_c)$ is computed by single-sine interpolation (Schlenker and Roberts, 2009). Let $\bar{t} = (T^{\max} + T^{\min})/2$ and $a = (T^{\max} - T^{\min})/2$. Three cases arise:

- $T_{d,t}^{\min} \geq \bar{T}_c$: entire day above threshold, $h = \bar{t} - \bar{T}_c$.
- $T_{d,t}^{\max} \leq \bar{T}_c$: no exceedance, $h = 0$.
- $T_{d,t}^{\min} < \bar{T}_c < T_{d,t}^{\max}$: partial crossing, $h = \frac{1}{\pi}\left[a \sin\alpha + (\bar{t} - \bar{T}_c)\,\alpha\right]$, where $\alpha = \arccos\left(\frac{\bar{T}_c - \bar{t}}{a}\right)$.

Phase assignment follows two conceptually distinct criteria, matching the two theoretical objects: $r = \theta^{\text{yield}}$ (yield-critical realization, a biological object determined before input decisions) and $s = \theta^{\text{pest}}$ (the observable pest-signal, a behavioral object that governs input adjustment).

Yield-critical phase (τ^*). The yield-critical phase is assigned ex ante using two sources of information: (i) the crop-specific thermal sensitivity windows established in the agronomic literature (Hakala et al., 2020; Hatfield et al., 2011; Lobell et al., 2013), and (ii) the institutional phenological calendar prescribed by Quebec's Crop Insurance Programme, which fixes the phase boundaries listed in Table 2. For cold cereals (oats, barley), the agronomic evidence motivates the sowing and establishment window (GS1: May 16–June 15) as the candidate yield-critical phase. Table 14 reports phase-by-phase yield regressions as supporting evidence: for both oats and barley, the sowing-phase KDD carries the strongest negative yield association among the retained candidate windows (column 3), consistent with the agronomic designation.

Pest-response phase (τ^{pest}). The pesticide-response phase is assigned using agronomic timing together with the empirical association between phase-specific KDD and pesticide

expenditure. For transparency, Table 14 reports the residualized slope $\hat{b}_\tau$ and partial R^2 for each candidate phase:

$$R^2_\tau(c,v) = 1 - \frac{\mathrm{Var}(\tilde{v} - \hat{b}_\tau \tilde{\theta}_\tau)}{\mathrm{Var}(\tilde{v})}, \qquad \hat{b}_\tau = \frac{\mathrm{Cov}(\tilde{\theta}_\tau, \tilde{v})}{\mathrm{Var}(\tilde{\theta}_\tau)}, \tag{B.1}$$

where tildes denote residuals after the relevant projection. For the pesticide-expenditure diagnostics, both pesticide expenditure and KDD are residualised on region and survey-wave fixed effects. The yield-panel diagnostics additionally partial out the weather controls used in the yield specification. These statistics are used as supporting diagnostics rather than as a mechanical selection criterion, and no universal sign restriction is imposed because the model permits both positive and negative pesticide responses depending on the adaptation regime.

For oats, barley, and corn, the retained phase is consistent with the agronomic timing of crop protection and is also supported by the relative explanatory power of the phase-specific KDD measures. For soybeans, the sowing-phase estimate is extremely large relative to the empirical support of sowing-phase KDD, while the growing phase corresponds to the crop's main in-season pest-pressure window and is therefore retained on agronomic grounds. Table 14 reports all candidate-phase estimates so that the phase assignment is fully transparent.

B.2 Distribution of Phase-Specific KDD by Crop

Figure 3 shows the full empirical distribution of phase-specific KDD for each crop over 1991–2022. The pronounced difference in scale between cold-crop panels (KDD_{25}, up to 60 degree-days) and warm-crop panels (KDD_{30}, up to 6 degree-days) reflects the agronomic design of the measure: Quebec summers regularly exceed 25 °C, so cold crops accumulate many degree-days above that threshold, whereas temperatures above 30 °C remain rare at the province's latitude. This scale contrast reflects the differing frequency of threshold exceedances across latitudes and should not be interpreted as a direct comparison of damage severity across crop groups.

C Identification

C.1 Identification Strategy

The region×crop fixed effects $\alpha_{r,c}$ absorb all time-invariant differences across cells: soil quality, local agronomy, baseline pest pressure, and average climatic conditions. Year fixed effects λ_t absorb aggregate shocks common to all cells in a given year, including input price movements, policy changes, and technology diffusion. Identification therefore comes entirely from within-cell, year-to-year variation in phase-specific KDD, conditional on weather controls.

Three features of the design strengthen the interpretation of $\hat{\beta}_c$ as a reduced-form heat response. First, the phase calendars are pre-determined by the crop insurance authority; neither farmers nor I select which days contribute to $\theta_{c,\tau,t}$, eliminating researcher discretion in the assignment of days to phases. Second, the damage thresholds $\bar{T}_c \in \{25, 30\}$ °C are selected ex ante from the crop-physiology literature and predetermined with respect to the estimation, limiting the scope of the threshold-selection critique raised by Schlenker and Roberts (2009). Third, including phase-specific precipitation and moisture deficit controls

helps distinguish the heat channel from correlated water stress.

The phase mapping for barley provides a targeted diagnostic for the phase-mismatch mechanism. Sowing-phase KDD is retained for the yield specification, whereas growing-phase KDD is retained for pesticide expenditure. Because these phase assignments are based on distinct biological and behavioral objects, their discordance motivates the interaction and falsification tests in Sections 4.5 and 5.6.

D Robustness Checks

D.1 Yield Robustness Checks

Four robustness checks supplement the main yield estimates in Table 6. R1 replaces the main KDD thresholds (KDD_{25} for cold crops, KDD_{30} for warm crops) with alternative thresholds (KDD_{27} and KDD_{31}, respectively), following Schlenker and Roberts (2009). R2 replaces each crop's retained phenological phase with the alternative candidate phase; weaker and less precise estimates provide a robustness check on the agronomically motivated phase assignment of Appendix B.1. R3 adds crop-specific regional linear time trends to control for slow-moving confounders such as variety adoption, drainage investment, and regional agro-climatic drift. R4 pools cold-season crops (oats + barley) and warm-season crops (corn + soybeans) within a single regression per group, replacing region fixed effects with region×crop fixed effects and adding a crop-specific linear time trend; the pooled estimates lie between the individual-crop estimates and are more precisely estimated, confirming that the main results are not driven by outlier dynamics in either crop.

D.2 Maladaptation Robustness Tables

Leave-one-region-out: barley maladaptation coefficient. Table 19 reports the leave-one-region-out exercise for $\hat{\beta}_3$. Each row re-estimates equation (4.8) on the barley CECPA sample after dropping the indicated administrative region. The result is geographically stable and uniformly negative across all six exclusions. Among the five peripheral-region exclusions, the estimates range from -0.088 to -0.106, and four remain statistically significant. Excluding the dominant Region 1 (43 of 61 observations) reduces the sample to 18 observations; $\hat{\beta}_3$ deepens to -0.295 ($p = 0.014$) (SE $= 0.059$). Overall, all six leave-one-region-out estimates remain negative and five are statistically significant. The leave-one-out exercise for the main yield equation confirms similar stability: dropping any single region leaves $\hat{\beta}_c$ (barley, $\text{KDD}_{25,\text{sow}}$) in the range $[-0.010, -0.018]$, all significant at conventional levels.

Robustness of $\hat{\beta}_3$ across specifications. Table 20 compares $\hat{\beta}_3$ for barley across three robustness checks. R1 uses an alternative KDD threshold (KDD_{27} instead of KDD_{25}); R2 uses annual KDD in the interaction term instead of the three-year moving average; R3 is the falsification test replacing the interaction with $KDD_{grow} \times \ln(pest)$ (the phase to which pesticide expenditure responds) in place of $KDD_{sow} \times \ln(pest)$. If the negative interaction is specific to mismatch with the yield-critical sowing phase, the growing-phase interaction should be close to zero. Table 20 is consistent with this implication.

E Additional Results

E.1 Climate Projections (Illustrative)

The estimated semi-elasticities are translated into near-term projected yield and unit-cost changes under SSP2–4.5 for 2025–2030. The exercise uses an arithmetic mean across five CanDCS-M6 models (CNRM-CM6-1, CNRM-ESM2-1, IPSL-CM6A-LR, MIROC-ES2L, and MIROC6; Sobie et al., 2024) under SSP2–4.5. For each phase, $\Delta\text{KDD} = \overline{\text{KDD}}_{\text{scenario}} - \overline{\text{KDD}}_{\text{baseline}}$, where the baseline is the 1990–2024 five-model CanDCS-M6 ensemble mean and the projected values are the 2025–2030 period averages from the same ensemble. Cold crops (oats, barley) use $\Delta\text{KDD}_{25,\text{sow}}$; warm crops (corn, soybeans) use $\Delta\text{KDD}_{30,\text{grow}}$. Cost and yield changes follow $\Delta c/c_0 = \exp(\hat{\kappa}_c \cdot \Delta\text{KDD}) - 1$ and $\Delta y/y_0 = \exp(\hat{\beta}_c \cdot \Delta\text{KDD}) - 1$, where $\hat{\kappa}_c$ and $\hat{\beta}_c$ are from Tables 10 and 6.

Three caveats apply. First, the log-linear functional form extrapolates outside the historical sample range. Longer-horizon projections are not reported because they require substantially greater extrapolation beyond the historical climate conditions used for estimation. Second, the CanDCS-M6 projections and the historical NASA POWER series differ in data-generating system and spatial resolution. The former are statistically downscaled gridded climate-model projections; the latter are reanalysis-based meteorological series extracted at representative regional coordinates. These differences may affect comparability and reduce confidence in the absolute magnitudes of the projections. Third, the projections assume no technological adaptation, crop switching, or behavioral response.

The 2025–2030 horizon is the only one reported because it requires substantially less extrapolation than the longer-horizon scenarios. Given differences between the historical and CMIP6 climate products, the projections should be interpreted as illustrative transformations of the estimated semi-elasticities rather than as forecasts within a precisely defined common support.

Tables

Table 1: Climate adaptation regimes by sign of (q_p, q_k)

Regime	Name	q_p	q_k	Quebec crops
1	Defensive intensification + frontier disengagement	> 0	< 0	Barley; oats (boundary, $q_p \approx 0$)
2	Broad intensification	> 0	> 0	Corn
3	Broad retrenchment	< 0	< 0	(none observed)
4	Pesticide retrenchment + frontier recovery	< 0	> 0	Soybeans

Table 2: Crop Calendars and KDD Damage Thresholds

Crop	$\bar{T}_c$	**Preparation (PP)**	**Sowing (GS1)**	**Growth & Harvest (GS2)**
Oats	25 °C	Apr 1 – May 15	May 16 – Jun 15	Jun 16 – Oct 25
Barley	25 °C	Apr 1 – May 15	May 16 – Jun 15	Jun 16 – Oct 25
Corn	30 °C	Mar 15 – May 1	May 2 – Jun 1	Jun 2 – Nov 25
Soybeans	30 °C	Apr 1 – May 15	May 16 – Jun 15	Jun 16 – Nov 15

Note: Phase boundaries are based on the Government of Quebec Crop Insurance Programme. The upper physiological damage thresholds are agronomically motivated and fixed ex ante from the crop-physiology literature. $\bar{T}_c$ is the threshold used to compute KDD (equation (3.1)).

Table 3: Descriptive Statistics of Phase-Specific KDD by Crop, Quebec 1991–2022

Crop	Phase	Obs	Mean	Std. Dev.	Min	Max
Cool-season crops (threshold $\bar{T}_c = 25$ °C)						
Oats	Preparation	266	0.462	0.881	0.000	5.875
	Sowing*†	266	3.300	3.110	0.000	16.269
	Growth	266	23.144	14.464	1.713	62.188
Barley	Preparation	266	0.387	0.682	0.000	4.893
	Sowing*	266	3.244	2.842	0.000	14.547
	Growth†	266	23.280	14.548	1.713	68.382
Warm-season crops (threshold $\bar{T}_c = 30$ °C)						
Corn	Preparation	243	0.000	0.002	0.000	0.027
	Sowing	243	0.081	0.275	0.000	1.923
	Growth*†	243	1.249	1.559	0.000	7.432
Soybeans	Preparation	208	0.005	0.043	0.000	0.542
	Sowing	208	0.157	0.350	0.000	1.923
	Growth*†	208	1.253	1.589	0.000	7.432

Notes: KDD computed from daily maximum and minimum temperatures from NASA Langley Research Center POWER Project (2026) using the single-sine interpolation method of Schlenker and Roberts (2009) and the phase calendars described in Section 3.1. Sample period: 1991–2022; 9 agricultural regions of Quebec. * Phase retained in yield regressions. † Phase retained in input-expenditure regressions. For barley, the yield regression uses Sowing-phase KDD_{25} while the pesticide regression uses Growth-phase KDD_{25}, generating a mismatch between the yield-relevant and pesticide-response phases.

Table 4: Descriptive Statistics of Crop Yields and Harvested Area, Quebec 1991–2022

Crop	Obs	Yield (t/ha)		Harvested area (ha)	
		Mean	Std. Dev.	Mean	Std. Dev.
Oats	265	2.46	0.40	7,377	5,099
Barley	266	3.01	0.54	8,785	8,110
Corn	243	7.60	1.43	34,735	42,998
Soybeans	208	2.66	0.41	25,767	23,567

Note: Sample: 9 agricultural regions of Quebec, 1991–2022; unbalanced panel (region-crop-year cells with zero reported area are excluded). Area figures are at the region level; province-level totals are approximately ten times these values. Source: Institut de la statistique du Québec (ISQ) and Statistics Canada (2023).

Table 5: Descriptive Statistics of Input Expenditures ($/ha), Quebec, CECPA Surveys 2009, 2014, 2019

Crop	Input	Obs	Mean	Std. Dev.	P25	P75
Oats	Pesticides	95	29.8	22.4	15.0	39.6
	Fertilizers	89	126.7	56.7	100.0	170.0
	Seeds	96	98.1	42.7	70.1	117.0
Barley	Pesticides	64	46.1	34.5	18.0	65.5
	Fertilizers	66	161.7	72.1	106.0	193.0
	Seeds	66	120.4	42.4	90.0	146.0
Corn	Pesticides	136	58.7	26.3	37.5	74.3
	Fertilizers	136	383.2	88.5	325.2	435.4
	Seeds	136	225.0	39.7	194.1	255.8
Soybeans	Pesticides	148	87.1	50.3	47.5	122.1
	Fertilizers	96	56.3	47.3	16.6	87.5
	Seeds	148	172.7	54.9	131.7	213.7

Notes: Nominal expenditures in Canadian dollars per hectare. P25 and P75 denote the 25th and 75th percentiles. Sample sizes differ across input types within crops due to item non-response. Climate regressors in input-expenditure equations use three-year backward-looking moving averages of phase-specific KDD ending at each survey year (2009, 2014, 2019). Source: Centre for Studies on Agricultural Production Costs (CECPA) (2021).

Table 6: Yield panel estimates: Effect of heat exposure

	Oats $\text{KDD}_{25,\text{sow}}$	*Barley* $\text{KDD}_{25,\text{sow}}$	*Corn* $\text{KDD}_{30,\text{grow}}$	*Soybeans* $\text{KDD}_{30,\text{grow}}$
Heat exposure ($\hat{\beta}_c$)	-0.023^{***}	-0.015^{**}	-0.024^{**}	-0.026^{***}
	(0.007)	(0.005)	(0.007)	(0.007)
Precipitation (3 phases)	✓	✓	✓	✓
Hydric deficit (sowing, growing)	✓	✓	✓	✓
Frost days (preparation)	✓	✓	✓	✓
Region FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
N	265	266	243	208
R^2_{within}	0.573	0.596	0.798	0.642

Notes: Each column reports the coefficient on the crop-specific phase KDD from a separate within-group regression of log yield per hectare estimated by the within-group (LSDV) fixed-effects estimator. The panel runs from 1991 to 2022 across nine Quebec agricultural regions; panel identifier is the region cell, and within-variation in annual KDD conditional on region and year fixed effects provides identification. The climate regressor is $\text{KDD}_{25,\text{sow}}$ for oats and barley and $\text{KDD}_{30,\text{grow}}$ for corn and soybeans, assigned using agronomic evidence and supported by the phase-specific yield estimates in Appendix B.1. Controls include phase-specific precipitation (preparation, sowing, growth), hydric deficit indicators for the sowing and growth phases, and frost days during land preparation. Standard errors (in parentheses) are clustered at the region level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 7: Extensive margin estimates: Effect of heat exposure on log sown area

	Oats	*Barley*	*Corn*	*Soybeans*
$\text{KDD}_{\tau,\text{prep}}$ (current)	−0.047	−0.081	8.670**	−0.098
	(0.094)	(0.100)	(3.055)	(0.265)
$\text{KDD}_{\tau,\text{sow}}$ (current)	−0.015	−0.008	0.349**	0.024
	(0.029)	(0.027)	(0.138)	(0.089)
$\text{KDD}_{\tau,\text{grow}}$ (lagged, $t-1$)	−0.004	−0.010**	0.019	−0.016
	(0.005)	(0.004)	(0.019)	(0.023)
ln(yield) (lagged, $t-1$)	0.234	−0.409	0.359	−0.142
	(0.160)	(0.356)	(0.438)	(0.210)
Precipitation (preparation, sowing)	✓	✓	✓	✓
Hydric deficit (preparation, sowing)	✓	✓	✓	✓
Region FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
N	248	247	229	198
R^2_{within}	0.387	0.756	0.354	0.852

Notes: Each column reports estimates from a separate within-group fixed-effects regression of log sown area on: (i) current-season phase-specific KDD for the preparation and sowing phases; (ii) lagged growing-season KDD (year $t-1$); (iii) lagged log yield (year $t-1$); and (iv) controls. Estimated by the within-group (LSDV) fixed-effects estimator on the panel of nine Quebec agricultural regions, 1991–2022. KDD_τ denotes KDD_{25} (threshold 25 °C) for oats and barley, and KDD_{30} (threshold 30 °C) for corn and soybeans. Regional controls include phase-specific precipitation and hydric deficit indicators for the preparation and sowing phases. Region and year fixed effects are included in all specifications. Standard errors (in parentheses) are clustered at the region level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 8: Input-expenditure estimates (CECPA): Effect of heat exposure

	Pesticide expenditure	Fertilizer expenditure	Seed expenditure
Oats			
$\text{KDD}_{25,\text{sow}}$ ($\hat{\eta}_c$)	0.052	-0.493^{***}	-0.208^{**}
	(0.071)	(0.083)	(0.081)
Barley†			
$\text{KDD}_{25,\text{grow}}$ ($\hat{\eta}_c$)	0.050^{***}	-0.034^{***}	-0.020^{***}
	(0.003)	(0.001)	(0.000)
Corn			
$\text{KDD}_{30,\text{grow}}$ ($\hat{\eta}_c$)	0.617^{***}	0.193^{**}	0.081^{***}
	(0.137)	(0.057)	(0.004)
Soybeans			
$\text{KDD}_{30,\text{grow}}$ ($\hat{\eta}_c$)	-0.300^{***}	0.089^{***}	0.169^{***}
	(0.027)	(0.008)	(0.009)
Precipitation	✓	✓	✓
Log output price	✓	✓	✓
Region FE	✓	✓	✓
Survey-wave FE	✓	✓	✓

† For barley, all three input-expenditure columns use growing-phase $\text{KDD}_{25,\text{grow}}$, the phase identified by the behavioral criterion (Appendix B.1). This differs from the sowing-phase $\text{KDD}_{25,\text{sow}}$ used in the barley yield regression, generating the phase mismatch that underlies the maladaptation identification.

Notes: Each cell reports $\hat{\eta}_c$ from equation (4.3), the KDD coefficient from a separate regression of log input expenditure per hectare on the crop-specific KDD, estimated with region and survey-wave fixed effects. The sample covers CECPA survey years 2009, 2014, and 2019 across nine Quebec agricultural regions. Climate regressors are three-year moving averages of the phase-specific KDD: $\text{KDD}_{25,\text{sow}}$ for oats; $\text{KDD}_{25,\text{grow}}$ for barley; $\text{KDD}_{30,\text{grow}}$ for corn and soybeans. Controls include log output price, phase-specific precipitation, and hydric deficit indicators. No input-mix controls (fertilizer and seed expenditure excluded from all columns). Organic farms excluded. Standard errors (in parentheses) clustered at the region level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 9: Budget shares: Effect of heat exposure on input expenditure shares

	Pesticide share	Fertilizer share	Seed share
Oats			
$\text{KDD}_{25,\text{sow}}$ $(\hat{\beta})$	-0.012^{***}	-0.085^{***}	-0.041^{**}
	(0.002)	(0.005)	(0.012)
Barley†			
$\text{KDD}_{25,\text{grow}}$ $(\hat{\beta})$	0.005^{***}	-0.001^{***}	0.003^{***}
	(0.000)	(0.000)	(0.000)
Corn			
$\text{KDD}_{30,\text{grow}}$ $(\hat{\beta})$	0.011^{**}	-0.030^{**}	-0.034^{***}
	(0.004)	(0.008)	(0.005)
Soybeans			
$\text{KDD}_{30,\text{grow}}$ $(\hat{\beta})$	-0.055^{***}	0.025^{***}	0.007^{***}
	(0.002)	(0.000)	(0.002)
Climate controls	✓	✓	✓
Log output price	✓	✓	✓
Region FE	✓	✓	✓
Survey-wave FE	✓	✓	✓

† For barley, all three columns use $\text{KDD}_{25,\text{grow}}$ (same phase selection as Table 8).

Notes: Each cell reports the KDD coefficient $\hat{\beta}$ from a separate regression of the input budget share (input expenditure as a fraction of total production expenditure, including machinery, labour, land, and fixed overhead recorded by CECPA) on the crop-specific phase KDD, estimated with region and survey-wave fixed effects. CECPA survey years 2009, 2014, 2019. Climate regressors are three-year moving averages of the pest-signal KDD for each crop (same phase selection as Table 8). Controls include log output price, phase-specific precipitation, and hydric deficit indicators. Organic farms excluded. Standard errors (in parentheses) clustered at the region level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 10: Cost-per-tonne decomposition: Effect of heat exposure

	Total expenditure ($\hat{\eta}_c^C$)	Cost per tonne ($\hat{\kappa}_c$)
Oats		
$\text{KDD}_{25,\text{sow}}$	0.289^{***}	0.395^{**}
	(0.025)	(0.098)
Barley		
$\text{KDD}_{25,\text{sow}}$	0.776^{***}	0.227^{***}
	(0.007)	(0.005)
Corn		
$\text{KDD}_{30,\text{grow}}$	0.280^{**}	0.137
	(0.112)	(0.253)
Soybeans		
$\text{KDD}_{30,\text{grow}}$	0.245	0.488^{***}
	(0.131)	(0.137)
Climate controls	✓	✓
Log output price	✓	✓
Region FE	✓	✓
Survey-wave FE	✓	✓

Notes: Column (1) reports the coefficient from a regression of log total input expenditure per hectare on the crop-specific phase KDD; Column (2) reports the coefficient from a regression of log cost per tonne (log expenditure minus log production) on the same KDD. Both regressions are estimated with region and survey-wave fixed effects, using three-year backward-looking moving averages of KDD ending at the CECPA survey years (2009, 2014, 2019). Column (2) is estimated directly from the CECPA data as the coefficient on KDD in a regression of $\ln(C/y)$ on KDD; it therefore captures the reduced-form effect of heat on unit cost without imposing the decomposition identity $\hat{\kappa}_c = \hat{\eta}_c^C - \hat{\beta}_c$ across datasets. The identity holds within a single sample, but the three-wave CECPA panel lacks the temporal variation needed to recover a reliable within-sample yield response; the decomposition of the unit-cost effect into its expenditure and output channels is therefore discussed conceptually using $\hat{\beta}_c$ from the long-run yield panel in Table 6, which is the appropriate source for yield elasticities. Controls include log output price, phase-specific precipitation, and hydric deficit indicators. Organic farms excluded. Standard errors (in parentheses) clustered at the region level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 11: Maladaptation test: interaction of yield-relevant KDD and log pesticide expenditure

	Oats $\text{KDD}_{25,\text{sow}}$	*Barley*† $\text{KDD}_{25,\text{sow}}$	*Corn* $\text{KDD}_{30,\text{grow}}$	*Soybeans* $\text{KDD}_{30,\text{grow}}$
KDD × ln(pesticides) ($\hat{\beta}_3$)	−0.087	−0.090**	−0.080	0.034*
	(0.111)	(0.022)	(0.104)	(0.014)
Climate controls	✓	✓	✓	✓
ln(output price)	✓	✓	✓	✓
Region FE	✓	✓	✓	✓
Survey-wave FE	✓	✓	✓	✓

† For barley, the interaction uses $\text{KDD}_{25,\text{sow}}$ (yield-critical phase), while pesticide expenditure is associated with $\text{KDD}_{25,\text{grow}}$ (pest-signal phase): the phase mismatch identified in Appendix B.1.

Notes: Each column reports $\hat{\beta}_3$, the coefficient on the KDD × log-pesticide interaction from a regression of log production per hectare on: (i) the three-year moving average of the yield-relevant KDD ($\bar{\theta}^{\text{yield}}$), (ii) log pesticide expenditure, (iii) their interaction, and (iv) controls. The regression is estimated with region and survey-wave fixed effects; CECPA survey years 2009, 2014, 2019. The yield-relevant KDD is $\text{KDD}_{25,\text{sow}}$ for oats and barley and $\text{KDD}_{30,\text{grow}}$ for corn and soybeans. Controls include log output price, phase-specific precipitation, and hydric deficit indicators. Fertilizer and seed expenditures are excluded from this specification to avoid conditioning on endogenous input-mix responses to the same climate shock. Organic farms excluded. Standard errors (in parentheses) clustered at the region level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 12: Projected yield and unit-cost changes under SSP2–4.5, 2025–2030

	Oats	*Barley*	*Corn*	*Soybeans*
ΔKDD	1.52	1.52	2.65	2.65
Δ yield (%)	-3.4	-2.3	-6.1	-6.6
Δ unit cost (%)	$+82$	$+41$	$+44^{\dagger}$	$+264$

Notes: Five-model CanDCS-M6 ensemble mean (Sobie et al., 2024) under SSP2–4.5, 2025–2030 horizon. ΔKDD refers to the yield-critical phase: $\Delta\text{KDD}_{25,\text{sow}}$ for oats and barley; $\Delta\text{KDD}_{30,\text{grow}}$ for corn and soybeans. Δ yield (%) $= (\exp(\hat{\beta}_c \cdot \Delta\text{KDD}) - 1) \times 100$; $\hat{\beta}_c$ from Table 6. Δ unit cost (%) $= (\exp(\hat{\kappa}_c \cdot \Delta\text{KDD}) - 1) \times 100$; $\hat{\kappa}_c$ from Table 10. Unit-cost changes refer to cost per tonne of output, not total expenditure per hectare. These are scenario-based point projections obtained by applying the estimated semi-elasticities to the five-model CanDCS-M6 ensemble-mean KDD changes and should not be interpreted as probabilistic forecasts. † The corn unit-cost projection is based on a positive but statistically insignificant estimate of $\hat{\kappa}_{\text{corn}}$.

Table 13: Testable predictions of the theoretical framework

Prediction	Mechanism	Expected sign / direction
2.1	Pest-signal heat $\to \uparrow$ pesticide demand when $q_p > 0$ (Lemma 2.1); $\downarrow$ when yield-value channel dominates ($q_p < 0$, soybeans)	$\partial x^*_{\text{pest}} / \partial \theta^{\text{pest}} \gtrless 0$
2.2	Pest-signal heat $\to \downarrow$ frontier input demand when $q_k < 0$ (Lemma 2.2); $\uparrow$ in recovery-demand regime (corn, soybeans)	$\partial f^*_k / \partial \theta^{\text{pest}} \lessgtr 0$
2.3	Inputs do not offset yield loss (Prop. 1)	$dy^* / d\theta^{\text{yield}} < 0$
2.4	Unit cost rises if expenditure is non-decreasing; if expenditure falls, yield loss must exceed the proportional expenditure reduction (Prop. 2); expenditure direction is crop-specific	$d \ln c / d\theta > 0$
2.5	Phase mismatch attenuates adaptive response (Section 2.8)	$\beta_3 < 0$ as diagnostic for phase-mismatched crops; no sign restriction for phase-aligned crops

Notes: θ^{pest} denotes heat stress during the phase associated with pest pressure and input adjustment; θ^{yield} denotes heat stress during the phenologically critical phase for yield formation. For most crops these coincide ($\tau^{\text{pest}} = \tau^*$); for barley (*orge*), the pest-relevant phase is the vegetative growth window while the yield-critical phase is the sowing and germination window. β_3 is the coefficient on $\theta^{\text{yield}} \times \ln x_{\text{pest}}$ in the augmented yield regression (equation (4.8)).

Table 14: Phase Selection: Partial R^2 and Slope by Crop, Phase, and Outcome

		Yield panel		Pesticide expenditure	
Crop	**Phase**	$\hat{b}_\tau$	R^2_τ	$\hat{b}_\tau$	R^2_τ
Cold crops (threshold 25 °C)					
Oats	Preparation	+0.012	0.002	−0.190	0.002
	Sowing ★	−0.023***	0.057 ★	+0.168	0.005 ★
	Growing	−0.004	0.026	+0.021	0.004
Barley	Preparation	+0.046**	0.023	−1.164	0.061
	Sowing ★	−0.015**	0.020 ★	+1.261*	0.106
	Growing	−0.005*	0.024	+0.112**	0.136 ★
Warm crops (threshold 30 °C)					
Corn	Sowing	+0.041	0.006	−0.456	0.000
	Growing ★	−0.024**	0.036 ★	+0.117	0.003 ★
Soybeans	Sowing	−0.042*	0.008	+5.218**	0.018
	Growing ★	−0.026***	0.038 ★	+0.134	0.003 ★

Notes: $\hat{b}_\tau$ is the OLS slope of the outcome on KDD after residualising both variables on their respective fixed effects (Frisch-Waugh-Lovell); yield-panel columns additionally partial out the precipitation and frost controls listed in Appendix B.1, while pesticide-expenditure columns use fixed effects only. R^2_τ is the squared correlation between those residuals. ★ marks the phase retained in the main estimation. Retention combines agronomic timing with the diagnostic evidence reported in the table. Yield panel: region×crop fixed effects and year dummies. Pesticide expenditure: region and survey-wave fixed effects. Standard errors clustered at the region level. For soybeans, the sowing-phase pesticide-expenditure coefficient ($\hat{b} = 5.22$) is implausible relative to the empirical support of sowing-phase KDD; the growing phase is retained on agronomic grounds as the crop's main in-season pest-pressure window (Boote et al., 2018; Schlenker and Roberts, 2009). The residualized phase-selection slope for soybeans growing-phase pesticide expenditure ($\hat{b} = +0.134$) differs in sign from the fully controlled estimate in Table 8 (-0.300^{***}). Table 14 documents relative phase-specific explanatory content rather than reproducing the coefficient from the main expenditure specification; the latter additionally conditions on the full set of controls in equation (4.3). $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 15: R1 Yield estimates: alternative KDD thresholds (KDD_{27} cold crops / KDD_{31} warm crops)

	Oats	*Barley*	*Corn*	*Soybeans*
$KDD_{27,sow}$ $(\hat{\beta}_c)$	-0.036^{***}	-0.025^{***}		
	(0.007)	(0.007)		
$KDD_{31,grow}$ $(\hat{\beta}_c)$			-0.031^{**}	-0.038^{**}
			(0.013)	(0.015)
N	265	266	243	208
R^2_{within}	0.568	0.596	0.795	0.640

Notes: Same specification as Table 6 with alternative thresholds KDD_{27} (cold crops) and KDD_{31} (warm crops). All four estimates remain negative and significant, confirming robustness to threshold choice. Standard errors clustered at the region level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 16: R2 Yield estimates: alternative phenological phase

	Oats	*Barley*	*Corn*	*Soybeans*
$KDD_{25,grow}$ $(\hat{\beta}_c)$	-0.004	-0.005^{*}		
	(0.003)	(0.002)		
$KDD_{30,sow}$ $(\hat{\beta}_c)$			0.041	-0.042^{*}
			(0.034)	(0.022)
N	265	266	243	208
R^2_{within}	0.559	0.598	0.792	0.631

Notes: Same specification as Table 6 with the alternative candidate phase: growing instead of sowing for cold crops and sowing instead of growing for warm crops. Oats and corn lose statistical significance, while the barley and soybean coefficients become substantially smaller in magnitude and only marginally significant. Within-R^2 declines for oats, corn, and soybeans and is essentially unchanged for barley. Taken together with the agronomic evidence, these results support the retained phase assignments without relying on fit as a mechanical selection criterion. Standard errors clustered at the region level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 17: R3 Yield estimates: with crop-specific regional linear time trends

	Oats	*Barley*	*Corn*	*Soybeans*
$\text{KDD}_{25,\text{sow}}$ $(\hat{\beta}_c)$	-0.024^{***}	-0.014^{*}		
	(0.006)	(0.006)		
$\text{KDD}_{30,\text{grow}}$ $(\hat{\beta}_c)$			-0.029^{***}	-0.027^{**}
			(0.007)	(0.009)
Regional time trends	✓	✓	✓	✓
N	265	266	243	208
R^2_{within}	0.639	0.640	0.833	0.692

Notes: Same specification as Table 6 augmented with crop-specific regional linear time trends to control for gradual heterogeneous trends in technology, variety adoption, and agro-climatic conditions. All four estimates remain negative and statistically significant with magnitudes very close to the main estimates, confirming that the yield effects are not artefacts of region-specific trends. Standard errors clustered at the region level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 18: R4 Yield estimates: pooled within-group specifications

	Cold Pool	*Warm Pool*
	$\text{KDD}_{25,\text{sow}}$	$\text{KDD}_{30,\text{grow}}$
Heat exposure $(\hat{\beta})$	-0.019^{***}	-0.024^{***}
	(0.004)	(0.006)
Region×crop FE	✓	✓
Crop time trend	✓	✓
Year FE	✓	✓
N	531	451
R^2_{within}	0.534	0.667

Notes: Cold Pool pools oats and barley; Warm Pool pools corn and soybeans. Pooled specifications replace region fixed effects with region×crop fixed effects and add a crop-specific linear time trend to allow for differential productivity growth across crops in the same group. The pooled coefficient lies between the two crop-specific estimates in each group (see Table 6) and is more precisely estimated. Weather controls as in Table 6. Standard errors clustered at the region level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 19: Leave-one-region-out: $\hat{\beta}_3$ (barley maladaptation)

Region excluded	$\hat{\beta}_3$	SE	N
Region 1	-0.295^{**}	(0.059)	18
Region 2	-0.088^{**}	(0.023)	58
Region 3	-0.106^{***}	(0.022)	56
Region 4	-0.090^{**}	(0.022)	61
Region 6	-0.090^{**}	(0.022)	59
Region 8	-0.095	(0.082)	53
None (full sample)	-0.090^{**}	(0.022)	61

Notes: Each row reports $\hat{\beta}_3$ from equation (4.8) estimated on the barley CECPA sample after dropping the indicated administrative region. The estimation sample contains 61 farm-level observations (one singleton dropped by the fixed-effects estimator) across three CECPA survey waves (2009, 2014, 2019). Six administrative regions contain barley observations. All six leave-one-out estimates are negative; five are statistically significant. Among the five peripheral-region exclusions, the estimates range from -0.088 to -0.106, and four are statistically significant. Excluding Region 1 (dominant barley region, 43 of 61 observations) leaves $N = 18$; $\hat{\beta}_3$ deepens to -0.295 ($p = 0.014$) (SE $= 0.059$), confirming that the result is not driven by the dominant region. The full-sample estimate (-0.090, SE $= 0.022$) is consistent with Table 11. Standard errors clustered at the region level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 20: Robustness of the maladaptation coefficient $\hat{\beta}_3$: barley only

	(1) Main	(2) R1: KDD_{27}	(3) R2: Annual	(4) R3: Falsification
ln(pesticides) ($\hat{\beta}_2$)	0.261^{***}	0.216^{***}	0.261^{***}	0.121^{*}
	(0.021)	(0.004)	(0.021)	(0.062)
$KDD_{25,sow} \times \ln(pest)$ ($\hat{\beta}_3$)	-0.090^{**}			
	(0.022)			
$KDD_{27,sow} \times \ln(pest)$ ($\hat{\beta}_3$, R1)		-0.238^{***}		
		(0.036)		
$KDD^{ann}_{25,sow} \times \ln(pest)$ ($\hat{\beta}_3$, R2)			-0.090^{**}	
			(0.022)	
$KDD_{25,grow} \times \ln(pest)$ ($\hat{\beta}'_3$, R3)				-0.002
				(0.009)
N	61	61	61	61

Notes: Barley CECPA sample only ($N = 61$). The outcome is log production per hectare. All specifications include region and survey-wave fixed effects, log output price, and phase-specific precipitation and hydric deficit controls. The interaction term in column (1) uses the three-year moving average of $KDD_{25,sow}$. Column (2) uses the alternative threshold KDD_{27}; column (3) replaces the moving average with the annual KDD in the interaction only. Column (4) is the falsification specification: the interaction uses $KDD_{25,grow}$ (the phase to which pesticide expenditure is associated) rather than $KDD_{25,sow}$. A near-zero and insignificant $\hat{\beta}'_3 = -0.002$ in column (4) is consistent with pesticide spending being better aligned with the growing-phase signal. The direct KDD coefficient ($\hat{\beta}_1$) is omitted; see Tables 6 and 15. Standard errors clustered at the region level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 21: Joint phase-discrimination test: barley, sowing vs. growing phase

	Coefficient (cluster-robust SE)	Bootstrap p (wild-cluster, Webb weights)
$\text{KDD}_{25,\text{sow}} \times \ln(\text{pest})$ ($\hat{\delta}_s$)	-0.145^{*} (0.025)	0.070
$\text{KDD}_{25,\text{grow}} \times \ln(\text{pest})$ ($\hat{\delta}_g$)	$+0.015$ (0.018)	0.278
$H_0\colon \delta_s = \delta_g$ (Wald p)	0.041	0.063
$\hat{\rho}(\text{KDD}_{\text{sow}}, \text{KDD}_{\text{grow}} \mid \text{FE})$		0.737
N		61

Notes: Barley CECPA sample only ($N = 61$). Both phase interactions enter simultaneously. The outcome is log production per hectare; controls include log output price, phase-specific precipitation, and hydric deficit indicators; region and survey-wave fixed effects included. Cluster-robust standard errors (in parentheses) clustered at the region level ($G = 5$). Bootstrap p-values use the wild-cluster bootstrap with Webb weights (Cameron et al., 2008; MacKinnon and Webb, 2017). $\hat{\rho}$ is the within-FE partial correlation between the two KDD regressors; the resulting collinearity attenuates the power of the joint test. Given the small number of clusters ($G = 5$), inference is based on wild-cluster bootstrap p-values using Webb weights. Significance stars reflect bootstrap inference. $^{*}p_{\text{WCB}} < 0.10$, $^{**}p_{\text{WCB}} < 0.05$, $^{***}p_{\text{WCB}} < 0.01$.

Table 22: Illustrative projected yield and unit-cost changes under SSP2–4.5 (five-model CanDCS-M6 ensemble mean)

Period		*Oats*	*Barley*	*Corn*	*Soybeans*
2025–2030	ΔKDD	1.52	1.52	2.65	2.65
	Δ yield (%)	-3.4	-2.3	-6.1	-6.6
	Δ unit cost (%)‡	$+82$	$+41$	$+44$	$+264$

Horizons 2030–2100 not reported: they require substantially greater extrapolation beyond historical conditions.

Notes: $\Delta\text{KDD} = \overline{\text{KDD}}_{\text{scenario}} - \overline{\text{KDD}}_{\text{baseline (1990–2024)}}$ from the five-model CanDCS-M6 ensemble mean (Sobie et al., 2024) under SSP2–4.5. Cold crops use $\Delta\text{KDD}_{25,\text{sow}}$ (KDD above 25 °C during the sowing phase); warm crops use $\Delta\text{KDD}_{30,\text{grow}}$ (KDD above 30 °C during the growing phase). $\hat{\kappa}_c$ and $\hat{\beta}_c$ from Tables 10 and 6. ‡ Δ unit cost $= (\exp(\hat{\kappa}_c \cdot \Delta\text{KDD}) - 1) \times 100$: the proportional change in cost per tonne of output, not in total expenditure per hectare. Longer-horizon projections are not reported because they require substantially greater extrapolation beyond historical climate conditions.

Figures

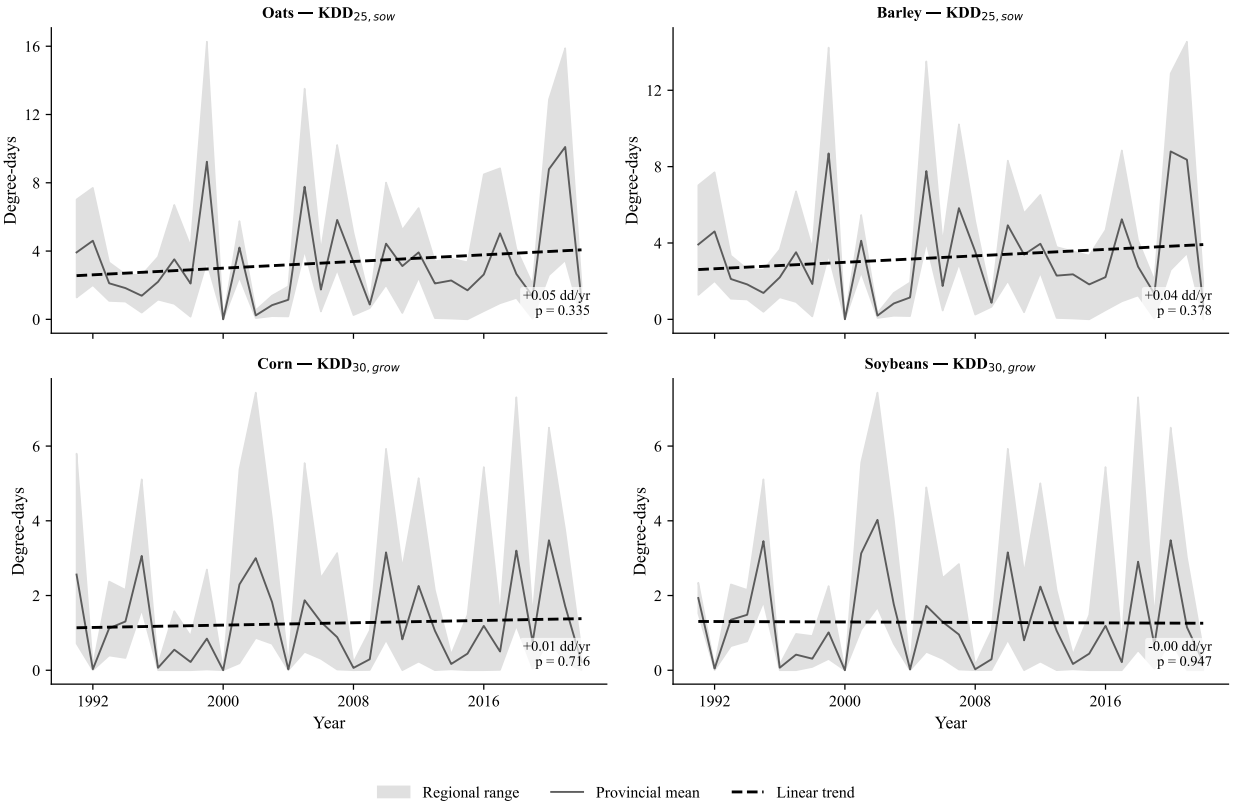


Figure 1: Trends in Phase-Specific Killing Degree Days by Crop, Quebec 1991–2022

Notes. Each panel plots the provincial mean of the retained phase-specific KDD for one crop over 1991–2022: $\text{KDD}_{25,\text{sow}}$ for oats and barley; $\text{KDD}_{30,\text{grow}}$ for corn and soybeans. The solid line is the annual provincial mean; the shaded area marks the range across regions. The dashed line is an OLS linear trend. For all four crops, the estimated trend is economically small and statistically insignificant, consistent with KDD variation being dominated by episodic year-to-year fluctuations rather than a deterministic linear trend. Panel scales differ because the crop-specific KDD measures use different physiological thresholds (Schlenker and Roberts, 2009).

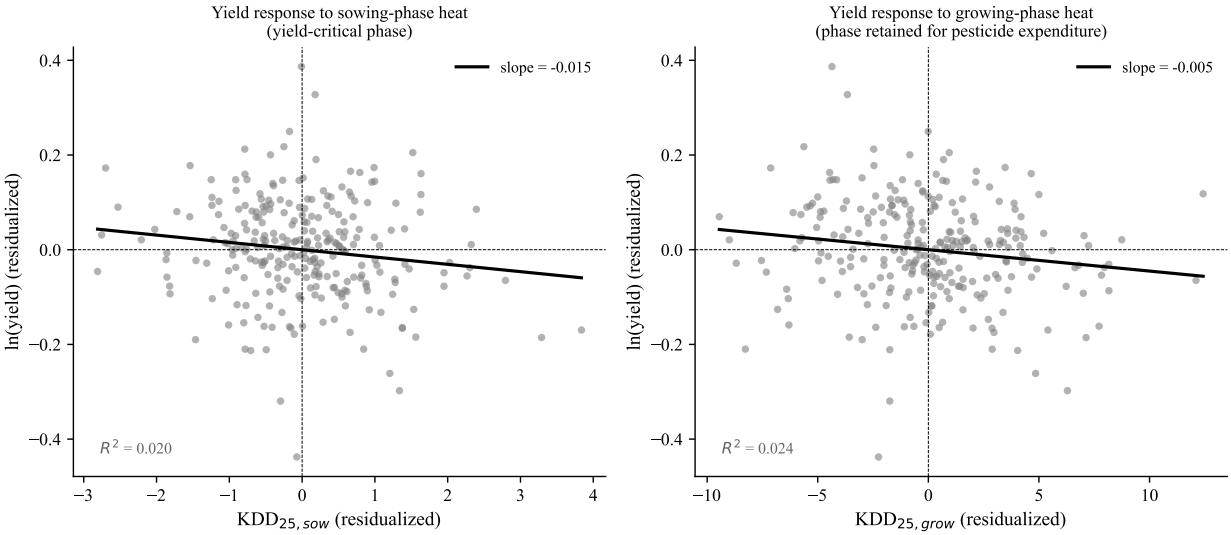


Figure 2: Barley Phase Discordance: Yield Sensitivity to Sowing- and Growing-Phase Heat

Notes. Barley only ($N = 266$ region-year observations, 1991–2022). Both panels plot residualized log yield against phase-specific KDD after partialling out region and year fixed effects and the weather controls used in the phase-selection yield regressions (precipitation, moisture deficit, and frost days by phase). *Left:* residualized log yield on $\text{KDD}_{25,\text{sow}}$, the yield-critical phase retained in the main specification; the negative slope shows that barley yield is more negatively associated with sowing-phase heat. *Right:* the same log yield residuals on $\text{KDD}_{25,\text{grow}}$, the phase retained for pesticide expenditure in Table 14; the unit effect (-0.005 per degree-day) is substantially smaller than the sowing-phase sensitivity, despite the growing phase exhibiting much greater thermal variance. The yield association is more negative per KDD in the sowing phase than in the growing phase. Together with the pesticide-expenditure phase evidence in Table 14, this contrast illustrates the phase discordance that motivates the maladaptation test.

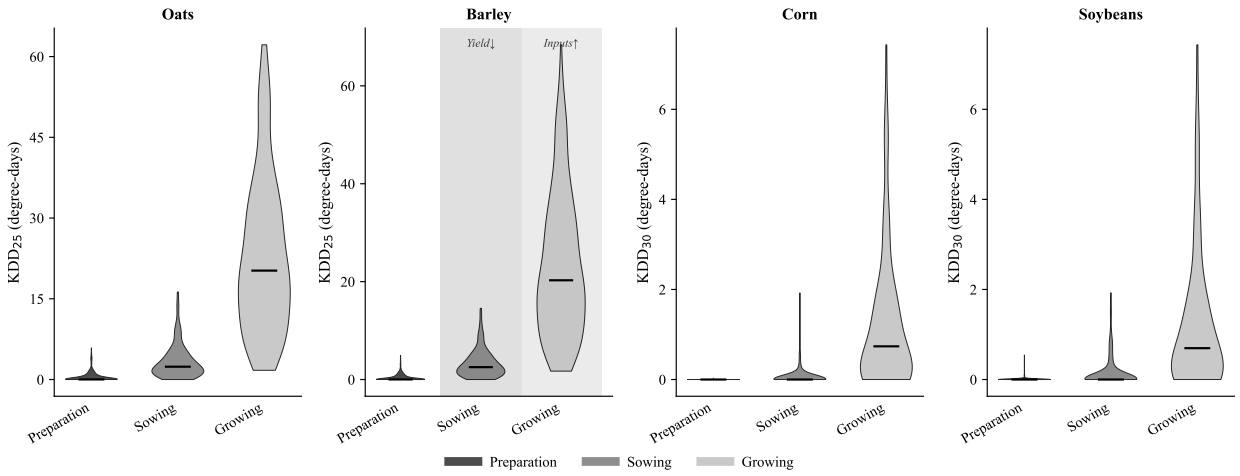


Figure 3: Distribution of Phase-Specific Killing Degree Days by Crop, Quebec 1991–2022

Notes. Each violin shows the kernel density of phase-specific KDD accumulated over the 1991–2022 panel. KDD_{25} is used for oats and barley; KDD_{30} for corn and soybeans. Horizontal lines indicate the median. Note the difference in vertical scale across crop groups: cold crops (left two panels) accumulate up to 60 degree-days above 25 °C per phase, while warm crops (right two panels) accumulate at most 5–6 degree-days above 30 °C, reflecting Quebec's position at the northern boundary of the warm-crop zone. For barley, the shaded bands identify the sowing phase (phase retained for yield, in darker grey) and the growing phase (phase retained for pesticide expenditure, in lighter grey), illustrating the empirical discordance that motivates Proposition 3.