Crop Insurance and the Pricing of Risk-Reducing Practices: Evidence from Rotational Complexity

Eunchun Park[1], Lawson Connor[1], Victor Funes-Leal[1], Frederi Viens[2], Gina Pizzo[2], Tyler Bagwell[2], Aria McLauchlan[3], and Harley Cross[3]

[1]Department of Agricultural Economics and Agribusiness, University of Arkansas

[2]Department of Statistics, Rice University

[3]Land Core

Abstract

Under the Federal Crop Insurance Program's Actual Production History–based Continuous Rating Formula (FCIP–APH–CRF), county benchmark rates are adjusted to the unit level primarily using average yield history, rather than separately pricing unit-level downside tail risk. We study this question using rotational complexity as an empirical application. We develop a mean–tail decomposition and define a *pricing gap* as the difference between expected indemnities implied by the full predictive yield distribution and APH-based CRF premiums. Using field-level Illinois corn data and posterior predictive yield distributions, we show that, in this application, changes relevant for insurance pricing arise primarily through the downside tail rather than through large mean shifts. Because CRF premiums scale with the APH mean while remaining insensitive to unit-specific tail risk, much of this risk reduction is not reflected in premium differences. A two-feature Shapley decomposition shows that the resulting change in the pricing gap is driven primarily by the tail channel rather than by small mean shifts. The results suggest that APH-only pricing weakly rewards the main risk-management benefit of more complex crop rotations. More generally, they illustrate how insurance schedules that recognize mean productivity but not downside-risk reduction can under-reward risk-reducing agronomic practices.

Keywords: crop insurance; continuous rating; crop rotational complexity; tail risk; expected shortfall

JEL Codes: Q14; Q18; G22; C58

1 Introduction

Crop insurance is the backbone of on-farm risk management in the United States. Under the federal crop insurance program's Continuous Rating Formula (FCIP-CRF), the premium paid by a producer is determined primarily by the unit's own yield history (actual production history, APH) relative to county benchmarks, with county loss–cost ratios, coverage level, and unit–structure adjustments layered on top (Tsiboe et al., 2025). This architecture is transparent to differences in expected yield, but it does not price farm–specific variance or downside tails since premiums primarily scale with the mean liability implied by the yield history.

The central question of the paper is: *When the rating formula recognizes average productivity but not practice-specific reductions in downside tail risk, do farmers still have a premium-based incentive to adopt risk-reducing practices?* We study this broader insurance-pricing question through rotational complexity, which provides a useful empirical application in settings where rotation diversification is associated with greater resilience under adverse conditions. We combine a pricing identity consistent with the CRF and field-level predictive yield distributions to quantify how much of this risk improvement goes unpriced in producer terms when farmers adopt more complex rotations.

We develop a unified mean–tail framework for evaluating how APH-based crop insurance pricing rewards, or fails to reward, downside-risk reduction. We define a *pricing gap*—the wedge between distribution-sensitive expected indemnities and mean-based CRF premiums—and characterize its comparative statics in a low-dimensional mean–tail space. The associated zero-gap boundary provides an empirically implementable object for evaluating practice-induced risk changes. We apply this framework to within-field variation in rotational complexity across weather regimes and use an Aumann–Shapley/finite-difference Shapley allocation to decompose changes in the pricing gap into mean and tail components.[1]

[1] As a policy illustration, Appendix E shows that a one-parameter volatility tilt can close a meaningful share of the producer-paid pricing gap at low–intermediate coverage without altering the CRF's multiplicative structure.

Across long-term experiments in North America, more diversified rotations are associated with greater resilience under adverse conditions, especially drought, primarily through reductions in downside yield risk rather than large increases in mean yields (Bowles et al., 2020; Socolar et al., 2021). More complex rotations also improve internal nitrogen cycling, consistent with buffered stress responses (Bowles et al., 2020). At the same time, rotational complexity remains limited in practice: monocultures and two-crop rotations account for most acreage nationally, and more complex rotations are least prevalent on biophysically advantaged land (Burchfield et al., 2024; Socolar et al., 2021). If crop insurance pricing recognizes mean yields but not downside-risk reduction, an important private return to diversification may remain unrewarded.

Under the actual CRF, within-field transitions in the Rotational Complexity Index (RCI, Socolar et al., 2021) yield three main findings. Each transition is represented in a shortfall–mean plane defined by the Aumann–Shapley zero-gap boundary. First, moving from low to high RCI reduces shortfall measures with only modest changes in mean yield, indicating that rotational complexity operates primarily through the lower tail of the yield distribution. Second, because CRF premiums are tied to the mean and not to unit-specific tail risk, these improvements generate only limited reductions in producer-paid premiums. Third, Shapley decompositions show that most of the resulting change in the pricing gap is attributable to the tail channel, with the strongest effects under dry conditions and at higher coverage levels.

Empirically, the analysis builds on recent spatial Bayesian work on rotational complexity and drought response. We adapt the posterior-predictive framework in Manski et al. (2024), which yields field-level predictive yield distributions by rotation and weather regime, and use those distributions to recover the means, shortfalls, and pricing gaps required for the analysis. Details of the posterior construction are in Manski et al. (2024), and Pizzo (2025) extends the validation to four additional Midwestern states with heterogeneous agroecological conditions. Appendix D reports implementation details and diagnostics. [2]

The remainder of the paper is organized as follows. Section 2 develops the theoretical framework and derives the pricing-gap identity and zero-gap boundary. Section 3 describes the data, predictive setup, and the mapping from predictive yield distributions to expected indemnities and premiums. Section 4 presents the empirical results. Section 5 discusses policy implications and sensitivity analyses. Section 6 concludes.

2 Theoretical Framework

Let i index farmer (unit), j county, t year (or rating epoch), and u unit structure. The *rate* (τ_{ijtu}, percentage) under the CRF is

$$\tau_{ijtu} = \alpha_{jt} \left(\frac{\bar{y}_{ijt}}{\bar{R}_{jt}}\right)^{\beta_{jt}} \cdot \Lambda(\theta_{it}, u) \cdot (1 + \delta_{jt}), \tag{1}$$

where α_{jt} is the county reference rate (from loss cost ratios), $\bar{y}_{ijt}$ is the unit's APH, $\bar{R}_{jt}$ is the county reference yield, and $\beta_{jt} < 0$ typically[3], $\Lambda(\theta, u)$ adjusts for coverage level θ and unit structure u, and δ_{jt} adds catastrophic loads.[4]

The *premium in dollars* (Π_{ijtu}) is taken to be the *producer-paid* premium (i.e., after subsidy). Let $\zeta(\theta_{it}, u) \in [0, 1)$ denote the statutory premium subsidy rate as a function of

[2]Nothing in our results hinges on the specific predictive engine or sampling capabilities: when only reported percentiles are available, if sufficiently descriptive of the overall predictive yield distribution, we recover shortfall via monotone interpolation and numerical integration, and the pricing-gap conclusions remain.

[3]β_{jt} is the *log elasticity* of the CRF *rate* with respect to the APH ratio $\bar{y}_{ijt}/\bar{R}_{jt}$. Thus a higher APH relative to the county benchmark lowers the *rate* when $\beta_{jt} < 0$ (the empirical norm). Because dollar premium equals rate $\times$ liability (and liability is linear in APH), the (gross or producer-paid) premium's log-elasticity with respect to APH is $1 + \beta_{jt}$: the producer subsidy $\zeta(\theta, u)$ does not depend on APH, so it does not affect this elasticity. See Tsiboe et al. (2025) for CRF calibration details.

[4]See Tsiboe et al. (2025) for construction and calibration details.

coverage and unit structure.[5] Then

$$\Pi_{ijtu} = (1 - \zeta(\theta_{it}, u))\, \tau_{ijtu} \cdot \text{Liability}_{ijtu} = (1 - \zeta(\theta_{it}, u))\, \tau_{ijtu} \cdot \Big[P_{it} \cdot \theta_{it} \cdot \bar{y}_{ijt} \cdot A_i\Big], \tag{2}$$

with price election P_{it} and acres A_i. Note that Π scales linearly with $\bar{y}_{ijt}$ via liability, while τ scales with $(\bar{y}_{ijt}/\bar{R}_{jt})^{\beta_{jt}}$; the subsidy factor $(1-\zeta)$ is multiplicative and depends only on (θ, u).[6]

2.1 Incentives under CRF

Let the farmer's yield be Y_i with cumulative distribution function (CDF) F_i, affected by a practice choice $a \in \mathcal{A}$ (e.g., an RCI level). Denote $\sigma_i^2 = \text{Var}(Y_i)$. Under an APH-based shortfall contract, the *indemnity* payoff is

$$I_i(Y_i; \bar{y}_i, \theta, P) \;=\; P\,(\theta \bar{y}_i - Y_i)^+. \tag{3}$$

The expected indemnity is a functional of the entire yield distribution (support in $[0, \infty)$):

$$\mathbb{E}[I_i] = P \int_0^{\theta \bar{y}_i} (\theta \bar{y}_i - y)\, dF_i(y) = P \int_0^{\theta \bar{y}_i} F_i(y)\, dy.^7 \tag{4}$$

Equivalently, for the relative yield $Z_i \equiv Y_i/\bar{y}_i$ with CDF G_i, using z as the integration argument,

$$\mathbb{E}[I_i] = P\, \bar{y}_i \int_0^{\theta} (\theta - z)\, dG_i(z) \;\equiv\; P\, \bar{y}_i\, S(\theta; G_i), \tag{5}$$

[5] Within-field contrasts in our empirical design hold θ and u fixed, so $\zeta(\theta, u)$ is locally constant. If a counterfactual changes coverage or unit structure, interpret ζ pointwise.

[6] If needed for comparison, the corresponding gross premium is $\Pi^{\text{gross}} = \Pi/(1 - \zeta(\theta, u))$.

[7] "Area" proof: for any $a > 0$, $(a - Y)^+ = \int_0^a \mathbf{1}\{Y \leq y\}\, dy$. Taking expectations and applying Fubini, $\mathbb{E}(a - Y)^+ = \int_0^a \Pr(Y \leq y)\, dy = \int_0^a F(y)\, dy$. Equivalently, by Riemann–Stieltjes integration by parts, $\int_0^a (a - y)\, dF(y) = [(a - y)F(y)]_0^a + \int_0^a F(y)\, dy = \int_0^a F(y)\, dy$ (since $F(0^-) = 0$ when $Y \geq 0$). Setting $a = \theta \bar{y}_i$ yields equation (4).

where $S(\theta; G) = \int_0^\theta (\theta - z)\, dG(z)$ summarizes downside mass below the guarantee (*shortfall*). Thus tail probability mass near/below the guarantee critically shapes $\mathbb{E}[I_i]$.

Under (1)–(2), holding county factors $(\alpha, \beta, \bar{R}, \delta)$, Λ, P, θ, and insured acres A fixed (so that the subsidy factor $(1 - \zeta(\theta, u))$ is locally constant with respect to a):

$$\frac{\partial \ln \Pi_i}{\partial \ln \bar{y}_i} = 1 + \beta, \qquad \text{(liability linear in } \bar{y}_i\text{; rate scales with } (\bar{y}_i/\bar{R})^\beta) \tag{6}$$

$$\frac{\partial \Pi_i}{\partial \sigma_i^2} = 0, \qquad \frac{\partial \Pi_i}{\partial S(\theta; G_i)} = 0, \tag{7}$$

because Π_i (interpreted as the *producer-paid* premium) is a function of the unit's APH $\bar{y}_i$—both through the rate $(\bar{y}_i/\bar{R}_{jt})^{\beta_{jt}}$ and linearly through liability—and of county-level rating parameters $(\alpha_{jt}, \beta_{jt}, \bar{R}_{jt}, \delta_{jt})$ and plan terms $\Lambda(\theta, u), P, \theta, A_i$, but it does not depend on the unit-specific variance σ_i^2 nor on the tail/shortfall functional $S(\theta; G_i)$.[8]

In contrast, working with relative yield $Z_i \equiv Y_i/\bar{y}_i$ gives

$$\mathbb{E}[I_i] = P\, \bar{y}_i\, \mathbb{E}\big[(\theta - Z_i)^+\big].$$

Since $(\theta - z)^+$ is convex, for any two relative yields with the same mean, if $Z \preceq_{cx} Z'$ (i.e., Z' is a mean-preserving spread of Z), then

$$\mathbb{E}[I_i \,|\, Z'] = P\, \bar{y}_i\, \mathbb{E}\big[(\theta - Z')^+\big] \;\geq\; P\, \bar{y}_i\, \mathbb{E}\big[(\theta - Z)^+\big] = \mathbb{E}[I_i \,|\, Z].$$

Equivalently, $S(\theta; G) = \mathbb{E}[(\theta - Z)^+]$ is increasing in the convex order; see Lemma 2.2, with proof in Appendix A.

Let practice a (e.g., an RCI level) shift the yield mean and distribution, $\big(\mu_i(a),\ F_i(\cdot \mid a)\big)$, where $\mu_i(a) \equiv \bar{y}_i(a)$. The farmer's period wealth with insurance is

$$W_i(a) = P\, Y_i(a) \;-\; C(a) \;-\; \Pi_i\big(\mu_i(a)\big) \;+\; I_i\big(Y_i(a); \mu_i(a), \theta, P\big),$$

[8] If one wishes to compare to the gross premium, divide by $(1 - \zeta(\theta, u))$. Our within-field contrasts keep θ and u fixed, hence $(1 - \zeta)$ is constant locally.

where $\Pi_i(\cdot)$ is the *producer-paid* premium defined in (2). With Constant Relative Risk Aversion (CRRA) utility $U(W) = W^{1-\eta}/(1-\eta)$, expected utility is $EU_i(a) = \mathbb{E}[U(W_i(a))]$. Differentiating and focusing on deterministic channels,[9]

$$\frac{d\,EU_i}{da} \propto \underbrace{P\,\frac{d\mu_i}{da}}_{\text{revenue}} - \underbrace{\frac{d\,\Pi_i}{da}}_{\text{premium (producer-paid)}} + \underbrace{\frac{d\,\mathbb{E}[I_i]}{da}}_{\text{indemnity}} - \underbrace{C'(a)}_{\text{cost}}. \tag{8}$$

To make the indemnity channel transparent, we use a stylized reduced-form notation in this subsection. Specifically, under the simplified pricing map used for exposition,

$$\mathbb{E}[I_i(a)] \;=\; P\,\mu_i(a)\,S\big(\theta; G_a\big), \qquad \Pi_i(a) \;=\; K(\theta,u)\,\mu_i(a)^{\,1+\beta},$$

where $S(\theta; G_a)$ captures the shape/tail of the yield distribution G_a (induced by $F_i(\cdot \mid a)$) and $\beta \in (-1, 0)$ in practice. Here $K(\theta, u)$ is understood to be the *producer-paid* pricing coefficient, i.e., it absorbs the factor $(1 - \zeta(\theta, u))$ from (2).[10] Then

$$\frac{d\,\mathbb{E}[I_i(a)]}{da} = \underbrace{P\,S(\theta; G_a)\,\frac{d\mu_i}{da}}_{\text{indemnity: mean base}} + \underbrace{P\,\mu_i(a)\,\frac{dS(\theta; G_a)}{da}}_{\text{indemnity: tail shape}}, \tag{9}$$

$$\frac{d\,\Pi_i}{da} = \Pi_i(a)\,(1+\beta)\,\frac{1}{\mu_i(a)}\,\frac{d\mu_i}{da}. \tag{10}$$

Equation (9) cleanly separates the indemnity effect into a mean component and a pure tail-shape component, while (10) shows that under the CRF the *producer-paid* premium responds only to the mean (liability) and is insensitive to tail improvements. All empirical calculations in Section 3.3, however, use the full county-level RMA baserate implementation, including the fixed-rate component δ_j.

Under the CRF, tail-only improvements holding μ_i fixed do not change the producer-paid premium ($d\Pi_i/da = 0$) but lower expected indemnity via the tail-shape term, attenuating

[9] Formally, $\frac{d}{da}\mathbb{E}[U(W_i(a))] = \mathbb{E}\big[U'(W_i(a))\{\cdots\}\big]$. We summarize signs by treating variation in $U'(W)$ as second order for small risk changes and by replacing dY_i/da with $d\mu_i/da$.

[10] Equivalently, $K^{\text{gross}}(\theta, u) = K(\theta, u)/(1 - \zeta(\theta, u))$.

the private value of variance reduction. If μ_i rises, dollar producer-paid premiums increase mechanically while rates fall (rate $\propto \mu_i^{\beta}$), making mean-channel incentives ambiguous. See Appendix B for the full derivation.

2.2 Mean–Tail Interactions under CRF

In the CRF, the guarantee base is the administrative APH $\bar{y}$. In the empirical application, unit-level APH is not observed, so we proxy $\bar{y}$ with a predictive mean $\mu \equiv \mathbb{E}[Y]$ recovered from field-level predictive distributions when mapping outcomes into the CRF (e.g., $\Pi^{\text{prod}}(\mu) = K\,\mu^{1+\beta}$, $I = P(\theta\mu - Y)^+$). See Section 3.1.

Let Y_i denote the yield and $\mu_i(a) \equiv \mathbb{E}[Y_i \mid a]$ the expected yield in practice $a \in \mathcal{A}$ (e.g., an RCI level). Define the relative yield $Z_i \equiv Y_i/\mu_i(a)$ with cdf G_a (so $\mathbb{E}[Z_i] = 1$). Under a yield contract with coverage $\theta \in (0, 1)$ and price P, the indemnity is

$$I_i \;=\; P\,(\theta\mu_i - Y_i)^+ \;=\; P\,\mu_i\,(\theta - Z_i)^+. \tag{11}$$

Hence expected indemnity decomposes as

$$\mathbb{E}[I_i \mid a] \;=\; P\,\mu_i(a)\;\underbrace{\mathbb{E}\big[(\theta - Z_i)^+\big]}_{S(\theta;G_a)} \;=\; P\,\mu_i(a)\int_0^{\theta}(\theta - z)\,dG_a(z), \tag{12}$$

where $S(\theta; G_a)$ depends only on the shape/tail of G_a (variance, skew, tail mass), not on the level μ_i.[11]

On a *producer-paid* basis, let

$$A_i(\theta) \;\equiv\; \frac{P\,\theta\,\alpha_j}{\bar{R}_j^{\beta_j}}.$$

[11] Because the contract is defined on a *relative* guarantee, let $R = Y/\mu$ with CDF G. Then, by the change of variables $y = \mu r$, $\mathbb{E}[I] = P\int_0^{\theta\mu}(\theta\mu - y)\,dF_Y(y) = P\,\mu\int_0^{\theta}(\theta - r)\,dF_R(r) = P\,\mu\,S(\theta; G)$. Hence $S(\theta; G) = \mathbb{E}[(\theta - R)^+]$ is a scale-free lower partial moment of R and depends only on the *shape/tail* of G, not on the level μ.

The CRF dollar premium used in the empirical implementation is

$$\Pi_i^{\text{prod}}\big(\mu_i\big) \;=\; (1-\zeta(\theta,u))\Big[A_i(\theta)\,\mu_i^{1+\beta_j} + P\,\theta\,\delta_j\,\mu_i\Big], \tag{13}$$

where β_j and δ_j are the county-level APH exponent and fixed-rate component from the RMA baserate files. For exposition, when the fixed-rate component is suppressed we sometimes write the premium in reduced form as $K_i\,\mu_i^{1+\beta}$.

We now define the *producer-paid pricing gap* as

$$\text{Gap}_i^{\text{prod}}(a) \;\equiv\; \mathbb{E}[I_i \mid a] - \Pi_i^{\text{prod}}(\mu_i(a)) \;=\; P\,\mu_i(a)\,S(\theta;G_a) - (1-\zeta(\theta,u))\Big[A_i(\theta)\,\mu_i(a)^{1+\beta_j} + P\,\theta\,\delta_j\,\mu_i(a)\Big]. \tag{14}$$

$\text{Gap}_i^{\text{prod}}(a)$ is the wedge between distribution-sensitive expected indemnities and producer-paid CRF premiums. A positive gap indicates underpricing for that field–practice configuration—expected losses exceed producer-paid premiums—so adopting a is privately subsidized; a negative gap indicates overpricing and discourages adoption. Because the first term responds to downside tails via $S(\theta;G_a)$ while the premium responds to the predictive mean through both a convex mean term and a linear fixed-rate component, pure tail improvements ($d\mu_i/da = 0$, $dS/da < 0$) reduce the gap but are not separately priced by the CRF. Accordingly, we track within-field changes $\Delta\text{Gap}_i^{\text{prod}}$ across $a^{\text{low}} \to a^{\text{high}}$ *within the same weather regime* and compare them to the zero-gap locus implied by equation (15).

Lemma (Monotonicity in downside risk). For any cdf's G_1, G_2 with $\mathbb{E}[Z] = 1$, if G_2 is a mean-preserving spread of G_1 *(equivalently, G_1 second-order stochastically dominates G_2, i.e. $\int_{-\infty}^{x}\big(G_1(t)-G_2(t)\big)\,dt \le 0$ for all x, with equality at $+\infty$)*, then

$$S(\theta;G_2) \;\ge\; S(\theta;G_1) \qquad \forall\,\theta \in (0,1),$$

because $z \mapsto (\theta - z)^+$ is convex in z. Thus SSD (Second-order Stochastic Dominates) improvements (less dispersion/thinner downside tails) imply $dS/da < 0$.[12]

Corollary (Pure tail improvement, producer-paid). If practice a changes only the shape of G_a while keeping $\mu_i(a)$ fixed (i.e., $\mu_i'(a) = 0$), then

$$\frac{d\,\mathrm{Gap}_i^{\mathrm{prod}}}{da} \;=\; P\,\mu_i(a)\,\frac{d\,S(\theta; G_a)}{da} \;<\; 0,$$

so the producer-paid gap strictly shrinks: CRF does not pass tail improvements into lower producer-paid premiums.

Proposition (Mean–tail interaction under CRF, producer-paid). Totally differentiating the producer-paid gap with respect to a yields

$$\begin{aligned}\frac{d\,\mathrm{Gap}_i^{\mathrm{prod}}}{da} &= \mu_i'(a)\Big[P\,S(\theta; G_a) - (1-\zeta)\Big(A_i(\theta)(1+\beta_j)\mu_i(a)^{\beta_j} + P\,\theta\,\delta_j\Big)\Big] + P\,\mu_i(a)\,\frac{d\,S(\theta; G_a)}{da} \\ &= \underbrace{\mu_i'(a)\Big[P\,S(\theta; G_a) - (1-\zeta)\Big(A_i(\theta)(1+\beta_j)\mu_i(a)^{\beta_j} + P\,\theta\,\delta_j\Big)\Big]}_{\text{mean (liability–rate) channel}} + \underbrace{P\,\mu_i(a)\,\frac{d\,S(\theta; G_a)}{da}}_{\text{distribution (tail) channel}}.\end{aligned} \tag{15}$$

Equation (15) is the *Aumann–Shapley* (marginal) allocation of $d\,\mathrm{Gap}^{\mathrm{prod}}$ across the mean and tail features of

$$f^{\mathrm{prod}}(\mu, S) = P\,\mu\,S - (1-\zeta)\Big[A_i(\theta)\mu^{1+\beta_j} + P\,\theta\,\delta_j\,\mu\Big].$$

For finite changes $a^{\mathrm{low}} \to a^{\mathrm{high}}$, we report the *two-feature Shapley* decomposition (average over the two orderings $\mu \to S$ and $S \to \mu$); see Methods, "Shapley decomposition of $\Delta\mathrm{Gap}^{\mathrm{prod}}$".

[12] Under SSD, expectations of every increasing concave function are ordered; a mean-preserving spread (MPS) makes all convex loss functionals weakly larger. See, e.g., Rothschild and Stiglitz (1970) or Hadar and Russell (1969).

Hence, when both mean and tail move, the sign of $\frac{d\,\mathrm{Gap}_i^{\mathrm{prod}}}{da}$ is governed by the trade-off between the producer-paid liability–rate bracket

$$P\,S(\theta;G_a)-(1-\zeta)\Big(A_i(\theta)(1+\beta_j)\mu_i(a)^{\beta_j}+P\,\theta\,\delta_j\Big)$$

and the tail change dS/da. A convenient sufficient condition for producer-paid gap reduction in the common case $\mu_i'(a)>0$ and $dS/da<0$ is

$$\underbrace{\frac{\mu_i'(a)}{\mu_i(a)}}_{\text{mean growth rate}} < \underbrace{\frac{-P\,\frac{dS(\theta;G_a)}{da}}{P\,S(\theta;G_a)-(1-\zeta)\Big(A_i(\theta)(1+\beta_j)\mu_i(a)^{\beta_j}+P\,\theta\,\delta_j\Big)}}_{\text{producer-paid liability–rate bracket}} \equiv \mathcal{T}_{\mathrm{AS}}(a,\theta), \quad (16)$$

where the numerator $-P\,\frac{dS(\theta;G_a)}{da}$ captures *tail improvement* and is positive whenever $dS/da<0$. When this inequality holds (and the bracket is positive), we have $\frac{d\,\mathrm{Gap}_i^{\mathrm{prod}}}{da}<0$. If the bracket is negative, the inequality reverses. Moreover, dividing the Aumann–Shapley marginal allocation in equation (15) by $\mu_i(a)>0$ and rearranging yields equation (16). For alternative sign configurations (e.g., $\Delta\mu/\mu<0$ or $\Delta S>0$) and the corresponding first-order boundary conditions, see Appendix C.

In the empirical application, we estimate $S(\theta;G_a)$ from predictive distributions (or from reported percentiles via monotone interpolation) and form within-field contrasts $a^{\mathrm{low}}\to a^{\mathrm{high}}$ to compute ΔS and $\Delta\mu_i/\mu_i$. Plotting these pairs in the $(\Delta\mu_i/\mu_i,\Delta S)$ plane against the *first-order (Aumann–Shapley) producer-paid zero-gap boundary*

$$\Big[P\,S-(1-\zeta)\Big(A_i(\theta)(1+\beta_j)\mu_i^{\beta_j}+P\,\theta\,\delta_j\Big)\Big]\cdot(\Delta\mu_i/\mu_i)+P\,\Delta S=0, \quad (17)$$

provides a transparent diagnostic: observations on the *negative side* of the boundary (i.e.

$$P\,\Delta S+\Big[P\,S-(1-\zeta)\Big(A_i(\theta)(1+\beta_j)\mu_i^{\beta_j}+P\,\theta\,\delta_j\Big)\Big](\Delta\mu/\mu)<0$$

) exhibit $\Delta\text{Gap}^{\text{prod}} < 0$.

While the RCI application often features higher means and lower downside tails, the framework does not require practice changes to move both margins in that direction. When means and tails move in other combinations, the sign of $\Delta\text{Gap}^{\text{prod}}$ is governed by the same boundary condition above.

3 Empirical Design

3.1 Data and Construction of Predictive Distributions

We use a field–year panel for Illinois corn over 2008–2020 by integrating four high-resolution layers: satellite-based predictive yields, crop-rotation histories, weather, and soils. Focusing on Illinois—a relatively homogeneous segment of the U.S. Corn Belt with dense corn acreage—facilitates use of fine spatial resolution while maintaining agro-climatic comparability.

Yields (SCYM). We use the Scalable Crop Yield Mapper (SCYM) framework (Lobell et al., 2015), which fuses satellite vegetation indices with process-based crop constraints to produce predictive yields at $\sim$30 m resolution; pixel predictions are aggregated to geospatial field polygons. SCYM outputs are used as predictive yields rather than as observed survey yields. Because SCYM reports above-ground biomass, we convert to grain using a standard harvest index of 0.5 (see Lobell et al. 2015).

Rotation Complexity Index (RCI). Using the USDA Cropland Data Layer (CDL), we construct the RCI on a rolling six-year window following Socolar et al. (2021). The index rises with the number and balance of distinct crops observed in the six-year history; at the lower bound, 1CROP denotes a monoculture history (the same crop in all six years), while at the upper bound, 6CROPS denotes six distinct crops appearing at least once over the window. Intermediate labels (2CROPS–5CROPS) indicate the count of distinct crops (e.g., a typical corn–soy sequence is 2CROPS). We decompose RCI into a between-field mean

$\mathrm{RCI}_i^{(m)}$ and a within-field annual deviation $\mathrm{RCI}_{it}^{(w)}$ to separate level and transitory practice effects.

Weather (Vapor Pressure Deficit, VPD). We summarize weather stress with vapor pressure deficit (VPD), the difference between saturation and actual vapor pressure; higher VPD indicates stronger atmospheric drying power, tighter stomatal regulation, and greater crop water stress. Meteorological inputs (temperature and humidity) are taken from standard gridded products (PRISM and TerraClimate), from which we compute VPD using a standard saturation curve and map the resulting rasters to field polygons by area weighting. For presentation and scenario design we work with three regimes indexed by $s \in \{dry, somewhat\ dry, normal\}$, anchored at representative values $\{22, 20, 18\}$ hPa (driest to least dry). When classifying historical observations, we use bins: *dry* > 21 hPa, *somewhat dry* $= 19$–21 hPa, *normal* < 19 hPa. These thresholds capture the tails that matter for shortfalls $S(\theta; G)$ while keeping Illinois counties' realizations well covered. Units are in hPa ($1\,\mathrm{kPa} = 10\,\mathrm{hPa}$).

Soils. Intrinsic land quality is measured using the gSSURGO National Commodity Crop Productivity Index (NCCPI), area-weighted to field polygons.

All layers are overlaid on field boundaries via GIS; covariates are standardized within year as needed. The final panel contains on the order of 10^5 field–year observations covering virtually all major corn counties. A compact summary of sources and variable construction appears in Table 1. For each scenario $c = (i, a, s)$ we estimate a predictive yield distribution $F_i^{(c)}$ and summarize it by $\mu_i^{(c)}$, percentiles, and shortfall functionals. Consistent with Section 2, we work with relative yields $R \equiv Y/\mu$ and proxy the APH base by the predictive mean. CRF premiums are then computed using county-level RMA baserate factors for Illinois corn, combining the county reference rate, APH exponent, reference yield, and fixed-rate component with the predictive mean proxy for APH for each field–RCI–weather scenario.

3.2 Predictive Setup and Recovery of (μ, S)

To recover the field-level yield distributions required for the pricing exercise, we adapt the predictive-posterior framework of Manski et al. (2024) and fit a parsimonious Bayesian mixed-effects model with field random intercepts and practice–weather–soil terms. The full equation, construction of scenario covariates, and diagnostics are provided in Appendix D.

For each field i and scenario $c = (i, a, s)$, we evaluate the fitted model at the scenario covariates, draw parameters and random effects from the joint posterior, and generate posterior predictive draws to form $F_i^{(c)}$. Let $G^{(c)}$ denote the CDF of the relative yield $R \equiv Y/\mu_i^{(c)}$. We then compute

$$S\big(\theta; G^{(c)}\big) \;=\; \int_0^{\theta} (\theta - r)\, dG^{(c)}(r) \;\approx\; \sum_{k:\, Q^{(c)}(p_k)/\mu_i^{(c)} \le \theta} \left(\theta - \frac{Q^{(c)}(p_k)}{\mu_i^{(c)}}\right) \Delta p_k, \tag{18}$$

using monotone interpolation on a fine percentile grid when only $\{Q^{(c)}(p_k)\}$ are stored. This yields (μ, S) for each field–RCI–weather scenario, which we map to expected indemnities $P\,\mu\,S$ and CRF premiums using the full county-level RMA baserate implementation in Section 3.3. For exposition, Section 2 sometimes writes the premium in reduced form as $K\,\mu^{1+\beta}$, but all empirical calculations include the fixed-rate component δ_j. Within-field contrasts pair the lowest and highest observed RCI levels under a common weather regime to form $(\Delta\mu/\mu, \Delta S)$ and the associated $\Delta\mathrm{Gap}^{\mathrm{prod}}$.

As a numerical robustness check, Appendix D.2 evaluates both the predictive-mean proxy and the shortfall recovery procedure under alternative feasible implementations. The predictive mean is nearly invariant across alternative APH-style mappings, and the paper's main result-space conclusions are highly stable across continuous shortfall interpolants (linear and monotone spline), although a discontinuous step interpolant is materially less stable and is therefore treated as a conservative stress test.

3.3 Mapping to Expected Indemnities and CRF Premiums

Because administrative APH at the unit level is not observed, we proxy $\bar{y}_i$ with the predictive mean μ_i recovered from the posterior predictive distribution for each field–RCI–weather cell (scenario). Fix coverage $\theta \in (0, 1)$ and price P. For each scenario $c = (i, a, s)$, expected indemnity is

$$\mathbb{E}[I_i \mid c] \; = \; P\, \mu_i^{(c)}\, S\big(\theta; G^{(c)}\big). \tag{19}$$

CRF-implied premiums are constructed from county-level RMA baserate factors. For county j, let α_j denote the reference rate, β_j the APH exponent, $\bar{R}_j$ the reference yield, and δ_j the fixed-rate component. Using $\mu_i^{(c)}$ as the empirical proxy for APH, we compute

$$\Pi_i^{(c)} = \frac{P\,\theta\,\alpha_j}{\bar{R}_j^{\beta_j}}\,(\mu_i^{(c)})^{1+\beta_j} + P\,\theta\,\delta_j\,\mu_i^{(c)}. \tag{20}$$

For exposition, we sometimes write the premium in reduced form as $\Pi_i(\mu) = K_i \mu^{1+\beta}$, where K_i collects county and contract-specific factors.

4 Results

Throughout this section, i indexes fields, a denotes RCI levels, and s denotes VPD regimes. Table 2 summarizes average premiums and expected indemnities across rotation complexity (RCI) and VPD regimes at fixed coverage $\theta = 0.75$. Premiums vary only modestly with RCI because the FCIP's APH-based rating prices off the mean. In contrast, expected indemnities fall sharply with rotation complexity—especially under *dry* VPD—so the pricing gap (EI$-$Premium) is less negative (often positive) in drier regimes and most negative under *normal*. This level pattern provides context for the within-field comparisons that follow. Figure 9 shows that increases in RCI reduce lower-tail shortfalls ($\Delta S < 0$), and Fig. 10 attributes most of the resulting change in the pricing gap to the tail channel rather than the mean.

4.1 Shortfall–mean plane.

For each within-field transition from a low to a high RCI level, we summarize the distributional change by two statistics: the percentage change in the mean, $\Delta\mu/\mu$, and the change in shortfall, $\Delta S(\theta; G)$. We plot these pairs in the *shortfall–mean plane* and overlay the dotted *zero-gap boundary* implied by the first-order expansion of the producer-paid pricing gap in equation (14) and its derivative decomposition in equation (15). Under the full county-level RMA implementation, the local approximation is

$$P\,\mu\,\Delta S \;+\; \Big[P\,S - (1-\zeta)\Big(A_i(\theta)(1+\beta_j)\mu^{\beta_j} + P\,\theta\,\delta_j\Big)\Big]\mu\,\frac{\Delta\mu}{\mu} \;\approx\; \Delta\mathrm{Gap}^{\mathrm{prod}}.$$

Points satisfying

$$P\,\mu\,\Delta S \;+\; \Big[P\,S - (1-\zeta)\Big(A_i(\theta)(1+\beta_j)\mu^{\beta_j} + P\,\theta\,\delta_j\Big)\Big]\mu\,\frac{\Delta\mu}{\mu} < 0$$

indicate $\Delta\mathrm{Gap}^{\mathrm{prod}} < 0$ (the producer-paid pricing gap shrinks) and are shown in green in Fig. 9.

Estimation uses the predictive posteriors (or reported percentiles with a monotone interpolation step) to form field$\times$VPD$\times\theta$ quantiles $Q(p)$, compute $\mu = \int_0^1 Q(p)\,dp$, construct $R_q(p) = Q(p)/\mu$ to recover G, and evaluate $S(\theta; G) = \theta G(\theta) - \int_0^\theta G(r)\,dr$ as in equation (18). Within each field and VPD regime we pair the observed minimum and maximum RCI levels to obtain $(\Delta\mu/\mu, \Delta S)$. The dotted boundary in Fig. 9 uses the Aumann–Shapley first-order approximation evaluated at group-wise means within each VPD$\times\theta$ panel; point colors reflect the *exact* finite-difference change in the producer-paid gap (EI$-$Premium) from equation (14).

First, most points lie in the lower half-plane ($\Delta S < 0$), indicating that RCI systematically reduces downside risk, while mean shifts are modest and typically positive ($\Delta\mu/\mu > 0$). Second, gap reductions dominate across coverage levels, consistent with equation (14): premiums depend only on μ, whereas expected indemnities scale with $S(\theta; G)$, so tail improvements reduce EI relative to premiums. Third, the effect strengthens with coverage and dryness: panels with higher θ and *dry* VPD show steeper downward movements (more negative ΔS) and a larger share of green points, reflecting greater sensitivity of $S(\theta; G)$ to lower-tail mass as θ rises. Because the plane encodes *changes*, points below the boundary signal a reduction in the gap for that transition; they do not by themselves imply that the *level* of the gap is negative. Economically, an APH-only rating attenuates the insurance-based incentive to adopt the practice that reduces the downside tail risk (in our case, the practice is higher RCI) or choose higher coverage, because realized reductions in lower-tail risk are not priced into premiums. These classifications are numerically robust to the shortfall-recovery method within the class of continuous interpolants: Appendix D.2 shows that replacing the baseline linear interpolation with a monotone spline preserves 98.9% of pairwise ΔGap signs and 92.1% of zero-gap-plane classifications.

4.2 Shapley decomposition.

The FCIP's APH-based rating prices premiums off the mean and ignores lower-tail risk. In the pricing identity (14), raising RCI nudges the mean μ slightly upward while compressing the lower tail (reducing S), so expected indemnities fall relative to premiums. The shortfall–mean plane (Fig. 9) makes this mechanism visible: most within-field transitions have $\Delta S < 0$ and lie on the gap-shrinking side of the boundary.

To quantify how much of ΔGap is due to tails versus means, we implement a two–feature finite-difference Shapley allocation (averaging the two orderings $\mu \to S$ and $S \to \mu$), linked to the Aumann–Shapley marginal decomposition in (15). Figure 10 and Table 3 show that the *tail channel dominates* at every coverage and in every VPD regime (every weather scenario).

At $\theta = 0.75$, the mean contribution is small while the tail contribution is large and negative: $(\text{mean}, \text{tail}) = (+0.49, -3.25)$ under *dry*, $(+0.28, -2.67)$ under *somewhat dry*, and $(+0.14, -2.13)$ under *normal*, yielding totals of -2.76, -2.38, and -1.99 \$/acre, respectively.

Measured as $|\text{tail}|/(|\text{mean}|+|\text{tail}|)$, the tail accounts for roughly 87% (*dry*), 91% (*somewhat dry*), and 94% (*normal*) of the magnitude of ΔGap at this coverage.

At higher coverage the pattern strengthens. At $\theta = 0.80$ the contributions are $(+1.56, -5.84)$ (*dry*), $(+1.18, -5.06)$ (*somewhat dry*), and $(+0.89, -4.54)$ (*normal*), with totals -4.28, -3.88, and -3.65. At $\theta = 0.85$ they are $(+3.71, -7.72)$, $(+3.19, -7.33)$, and $(+2.75, -7.08)$, implying totals -4.01, -4.14, and -4.34 \$/acre; the tail's share of the magnitude is about 68–72% at this highest coverage. By contrast, at $\theta = 0.65$ both channels are small (e.g., *dry*: -0.12 from the mean and -0.44 from the tail, total -0.57).

Because producer-paid premiums in the CRF respond only to μ, realized reductions in $S(\theta; G)$ are not rewarded in price. The Shapley results therefore indicate that adopting the downside-risk-reducing agronomic practice (in this case moving from low to high RCI) *reduces the producer-paid pricing gap* (EI − premium) mainly through thinner downside tails, with effects that grow in drier VPD regimes (drier weather scenarios) and at higher coverage (consistent with Fig. 9).

The dominance of the tail channel is also robust to the shortfall-recovery method within continuous interpolants: Appendix D.2 shows that the tail contribution remains dominant in 93.3% of $\theta\times$ VPD cells under both the baseline linear and monotone spline implementations.

4.3 Geography and cross-section of incentives

Figure 11 plots the *level* producer pricing gap at $\theta = 0.75$ under the *normal* VPD regime for the low-RCI benchmark ("1 crop"; colors: red = EI > Premium, blue = EI < Premium). Counties with larger positive gaps are concentrated in a few pockets, while most of the state lies closer to zero; this level picture provides the backdrop against which we interpret within-county transitions.

Turning to incentives, Figure 12 maps the *average* within-county change in the producer pricing gap when adopting the downside risk-reducing practice (moving from the lowest to the highest observed RCI) ($\Delta\text{Gap} = \text{Gap}^{\text{high}} - \text{Gap}^{\text{low}}$), faceted by VPD regime at $\theta =$

0.75. Consistent with the shortfall–mean mechanism in equation (14), most counties exhibit $\Delta\text{Gap} \leq 0$, with larger (more negative) changes in drier panels. Figure 13 complements this by showing, for each county and VPD regime, the *share* of within-field transitions with $\Delta\text{Gap} < 0$; the share is high statewide—especially under *dry*. Panels use a common legend scale within each figure and are ordered left–to–right as *dry*, *somewhat dry*, *normal*.

The cross-section reinforces these patterns. Figure 14 plots county means of μ (horizontal axis) against the county average ΔGap (vertical axis), with point sizes proportional to the number of within-field pairs and panels ordered *dry—somewhat dry—normal*. In *dry* counties the fitted line slopes downward (higher-μ counties tend to experience more negative ΔGap when RCI rises), while in *normal* counties the slope is weakly positive; slopes flatten toward zero as conditions become less stressful. These patterns are consistent with the shortfall term $S(\theta; G)$ bearing more weight where the guarantee bites.

5 Discussion: Incentives under APH-only rating

The main result of the paper is that APH-only rating leaves an important part of practice-induced risk reduction unpriced. In the empirical application, moving to higher rotational complexity primarily reduces downside shortfall, with only modest changes in mean yield. Because producer-paid premiums respond mainly to the mean, these within-field improvements translate into only limited premium reductions, so the producer-paid pricing gap typically falls when RCI rises.

Coverage incentives need not move monotonically with the practice margin. Under our empirical definition,

$$S(\theta; G) = \theta G(\theta) - \int_0^\theta G(r)\, dr,$$

the comparative static with respect to coverage depends on both local downside density near the guarantee and the slope of the premium schedule. As a result, higher coverage need not strengthen the producer-paid wedge. In the data, producer-paid premiums often rise sufficiently quickly with θ that lower or intermediate coverage levels yield a larger (that is, less negative) producer-paid gap. This is a statement about the pricing wedge rather than about total private returns: farmers may still benefit from adopting a risk-reducing practice through mean gains, tail-risk reduction, or other agronomic channels that lie outside the premium schedule.

These results also have implications for policy design. The current CRF recognizes expected yield but not unit-specific downside stabilization. As an illustration, Appendix E considers a one-parameter volatility adjustment that preserves the CRF's multiplicative structure while allowing premiums to respond to changes in predictive dispersion. In our data, that adjustment reduces part of the observed producer-paid wedge but does not fully internalize tail improvements, especially at higher coverage. This suggests that simple dispersion-based refinements may improve alignment at the margin, but they are unlikely to fully price practices whose main effect operates through the extreme lower tail rather than through the mean or overall variance.

6 Conclusion

This paper develops a CRF-consistent pricing framework that separates expected indemnities into a mean component and a shortfall functional of the relative yield distribution, and uses field-level predictive distributions to map practice-induced distributional changes into producer-paid pricing wedges. We study rotational complexity as an empirical application of this broader insurance-pricing question.

First, moving from lower to higher RCI mainly compresses the lower tail of the yield distribution, with only modest changes in mean yield. Second, Shapley decompositions show that the resulting change in the producer-paid pricing gap is driven primarily by the tail channel rather than by the mean channel, with the strongest effects under dry conditions and at higher coverage levels. Under APH-only pricing, the premium schedule therefore provides limited reward for the main risk-management benefit of more complex rotations.

More broadly, the analysis highlights a general limitation of APH-based pricing: when management practices reduce downside risk without proportionally changing expected yield, much of that stabilization is unlikely to be reflected in producer-paid premiums. The empirical application here focuses on rotational complexity, but the same pricing logic applies more generally to practices whose main effect is to reduce downside tail risk rather than to generate large gains in average yield.

Two caveats are important. First, the object of interest in this paper is the pricing gap, not total private payoff, so the results should be interpreted as statements about premium-based incentives rather than about overall profitability or welfare. Second, coverage incentives are not uniformly increasing: the comparative statics depend on downside density near the guarantee and on the slope of the premium schedule, so lower or intermediate coverage can generate a larger producer-paid wedge in many settings.

Appendix E considers a simple volatility-based refinement of the CRF. In our data, that adjustment reduces part of the observed misalignment but does not fully internalize tail improvements, especially where lower-tail risk is most important. Future work could embed these pricing results in a dynamic adoption and coverage framework, explore tail-sensitive proxies implementable with short APH histories, and assess program incidence under broader shifts in risk-reducing management.

Table 1: Summary of Data Sources and Key Variables

Category	**Source**	**Description**
Yields	SCYM (Lobell et al.)	Satellite-based predictive yields at $\sim$30m resolution, validated against NASS and farm-reported data. Converted from biomass to grain yield using a harvest index of 0.5.
Crop rotation	USDA Cropland Data Layer (CDL)	Rotational Complexity Index (RCI) over 6-year window; includes long-run mean and annual deviation.
Weather	PRISM, TerraClimate	Monthly precipitation, temperature, vapor pressure deficit (VPD), Palmer Drought Severity Index (PDSI), and soil moisture deficit (DEF). Emphasis on July max VPD.
Soil productivity	USDA gSSURGO	National Commodity Crop Productivity Index (NCCPI), aggregated to field polygons.

Table 2: Expected indemnities and premiums by rotation complexity (RCI) and VPD regime ($\theta = 0.75$).

RCI (rotation)	$\boldsymbol{\theta}$	**EI ($ / acre)**	**Producer premium ($ / acre)**	**Total premium ($ / acre)**	**Gap (EI−producer) ($ / acre)**
Panel A. Dry					
1 crop	0.75	9.94	5.61	12.47	4.33
2 crop	0.75	8.75	5.71	12.69	3.04
3 crop	0.75	8.33	5.75	12.77	2.59
4 crop	0.75	7.95	5.78	12.85	2.17
5 crop	0.75	7.70	5.80	12.90	1.90
6 crop	0.75	7.40	5.83	12.96	1.57
field average	0.75	8.74	5.71	12.68	3.03
Panel B. Somewhat dry					
1 crop	0.75	7.58	5.81	12.92	1.77
2 crop	0.75	6.59	5.91	13.13	0.68
3 crop	0.75	6.22	5.94	13.21	0.28
4 crop	0.75	5.88	5.98	13.28	−0.10
5 crop	0.75	5.67	6.00	13.33	−0.33
6 crop	0.75	5.41	6.03	13.40	−0.62
field average	0.75	6.57	5.91	13.13	0.66
Panel C. Normal					
1 crop	0.75	5.59	6.01	13.35	−0.42
2 crop	0.75	4.73	6.10	13.55	−1.37
3 crop	0.75	4.43	6.13	13.63	−1.70
4 crop	0.75	4.16	6.17	13.70	−2.01
5 crop	0.75	4.00	6.19	13.75	−2.19
6 crop	0.75	3.81	6.22	13.81	−2.41
field average	0.75	4.69	6.10	13.56	−1.41

Notes: Producer-paid premium = Total premium $\times(1 - \text{subsidy})$; subsidy = 0.55
$\theta = 0.75$ and *basic optional* units
Gap is EI minus producer-paid premium
Price election $P = 3.50$ \$/bu. EI = $P \cdot \mu \cdot S(\theta)$

Table 3: Shapley decomposition of *producer-paid* pricing-gap change by VPD regime and coverage level (θ).

VPD	θ	**Mean (μ) (\$ / acre)**	**Tail (S) (\$ / acre)**	**Total (\$ / acre)**
dry	0.65	−0.12	−0.44	−0.57
	0.70	0.04	−1.53	−1.50
	0.75	0.49	−3.25	−2.76
	0.80	1.56	−5.84	−4.28
	0.85	3.71	−7.72	−4.01
somewhat dry	0.65	−0.14	−0.16	−0.30
	0.70	−0.05	−1.05	−1.10
	0.75	0.28	−2.67	−2.38
	0.80	1.18	−5.06	−3.88
	0.85	3.19	−7.33	−4.14
normal	0.65	−0.14	−0.06	−0.21
	0.70	−0.10	−0.59	−0.69
	0.75	0.14	−2.13	−1.99
	0.80	0.89	−4.54	−3.65
	0.85	2.75	−7.08	−4.34

Notes: Producer-paid gap = EI − producer-paid premium. Values are average Shapley contributions to ΔGap (producer basis). Units: \$/acre. Price election P = 3.50 \$/bu. Subsidy per *basic optional* schedule.

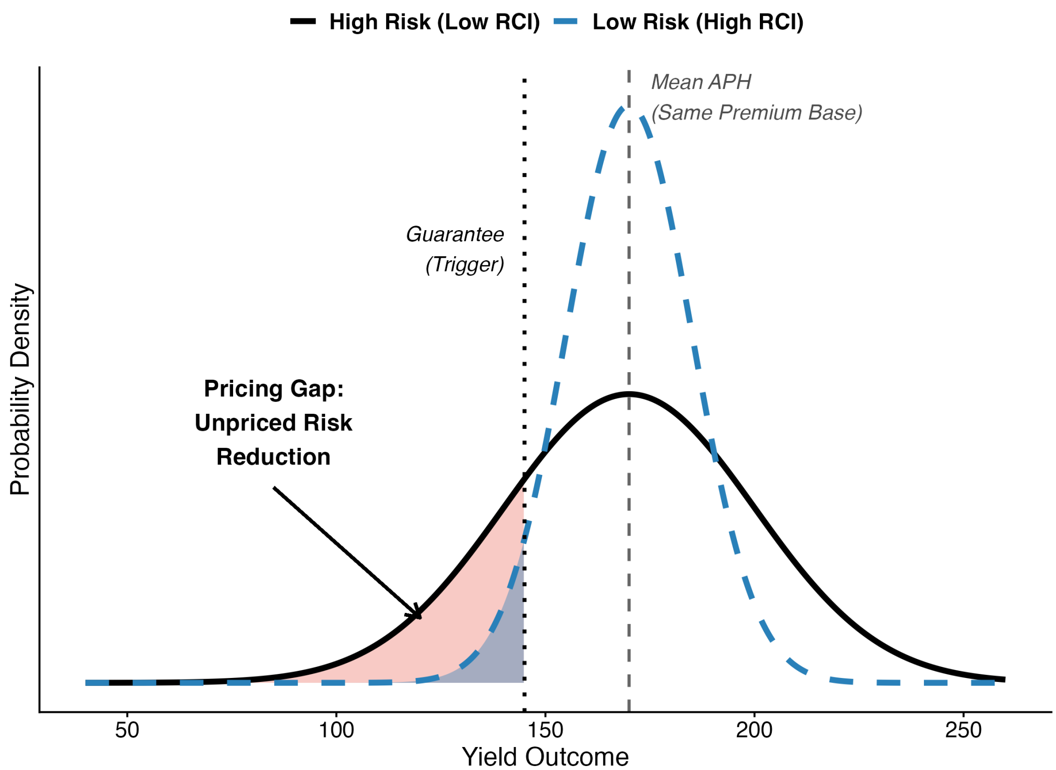

Figure 1: **Conceptual Illustration of the Pricing Gap.** For two units subject to the same county benchmarks, identical mean yields (μ) imply identical premiums. However, the complex rotation (dashed line) exhibits significantly lower risk by compressing the lower tail compared to the simpler rotation (solid line). This divergence between the fixed price and the reduced expected indemnity constitutes the "pricing gap."

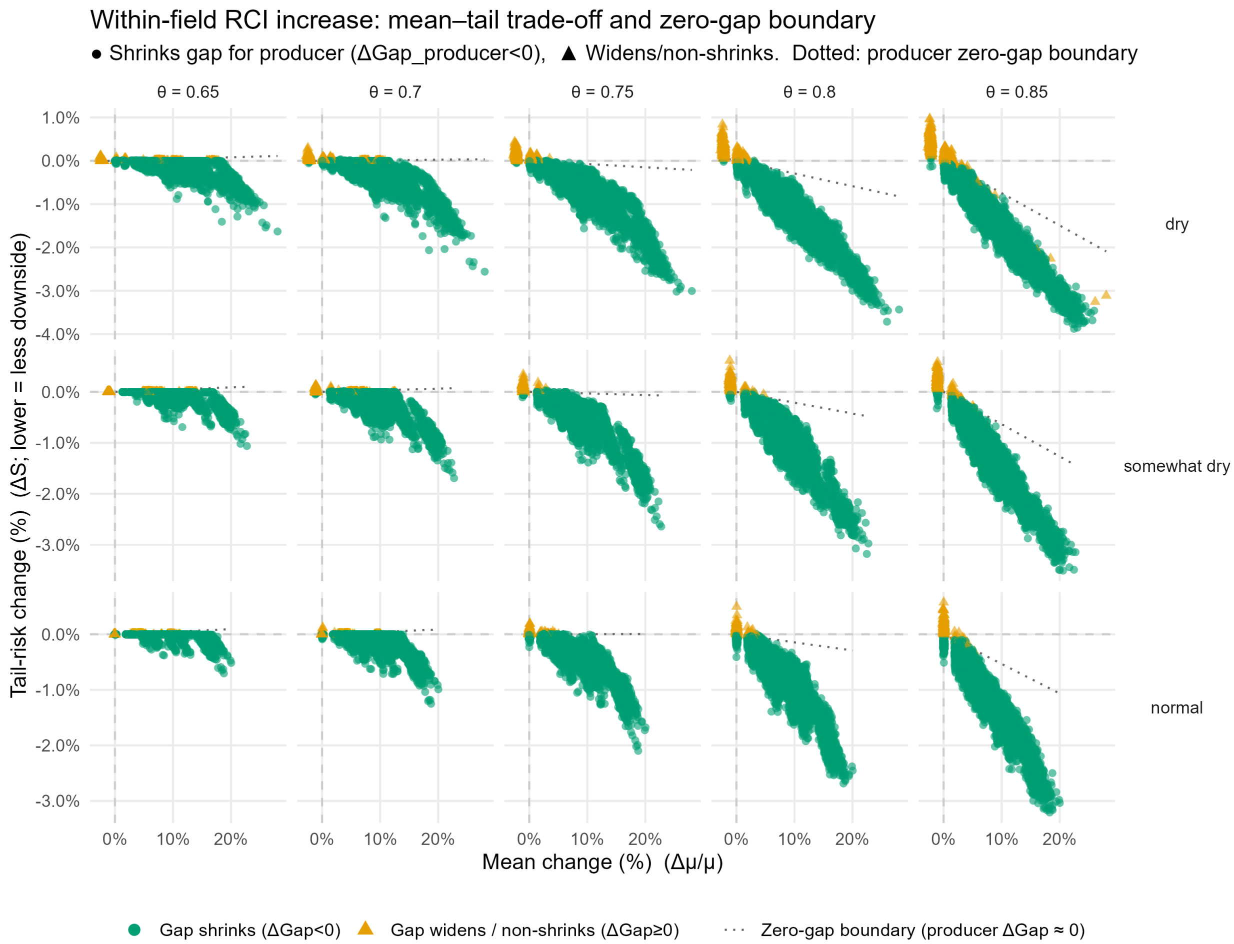

Figure 2: Within-field RCI increase: mean–tail trade-off and zero-gap boundary.
Panels: rows = VPD (dry, somewhat dry, normal); columns = coverage $\theta \in \{0.65, 0.70, 0.75, 0.80, 0.85\}$.
Axes: x = mean change (%, $\Delta\mu/\mu$); y = tail-risk change (%, ΔS; down = less downside).
Markers: green circles = gap shrinks for producer ($\Delta\text{Gap} < 0$); orange triangles = widens / non-shrinks ($\Delta\text{Gap} \geq 0$).
Line: dotted = group-wise zero-gap boundary ($\Delta\text{Gap} \approx 0$); below the line → mostly green, above → mostly orange.
Note: In these data EI is typically below Premium; "gap shrinks" often means the gap becomes *more negative*, not necessarily a welfare gain.

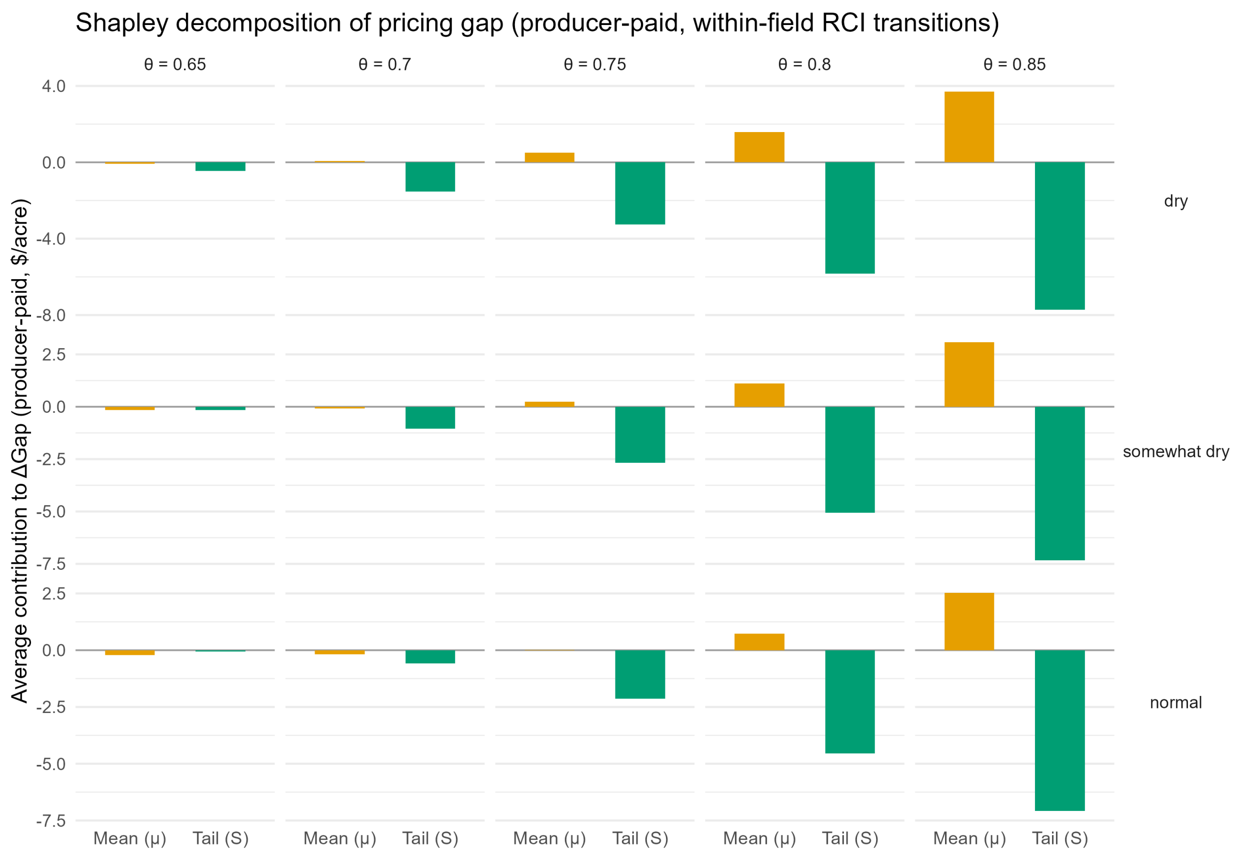

Figure 3: Shapley decomposition of the pricing–gap change for within–field RCI transitions. Bars show mean finite–difference Shapley contributions by channel (rows: VPD regime; columns: coverage θ).

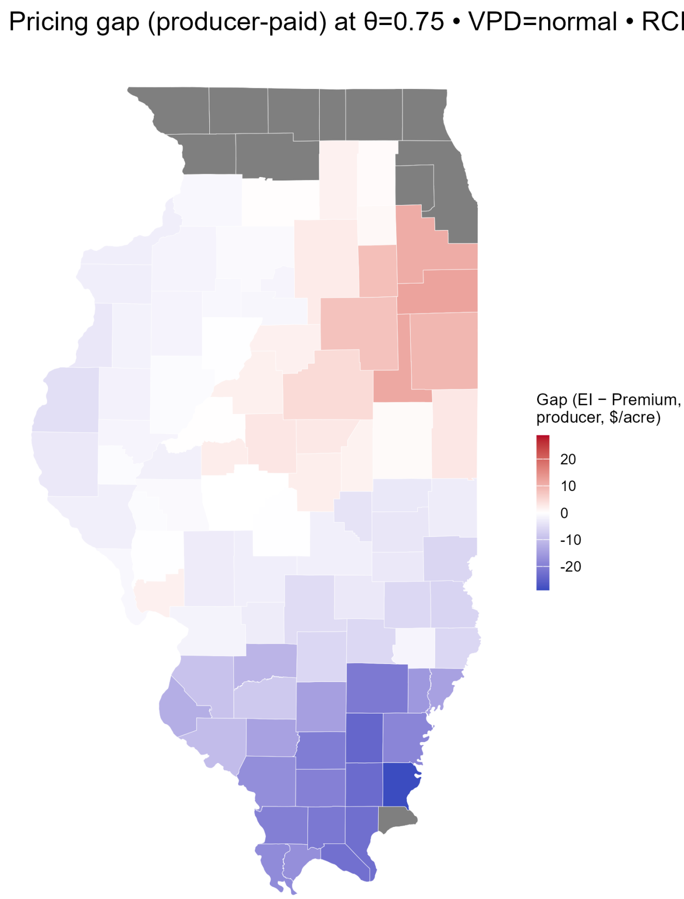

Figure 4: County-level producer-paid pricing gap at coverage $\theta = 0.75$ under the *normal* VPD regime and low rotational complexity (RCI = '1 crop'). Colors use a common blue–white–red scheme (white at zero): red = positive gap (expected indemnities exceed producer-paid premiums), blue = negative gap. This map anchors the geography of the wedge between distribution-sensitive expected indemnities and APH-based CRF premiums before considering within-field practice changes.

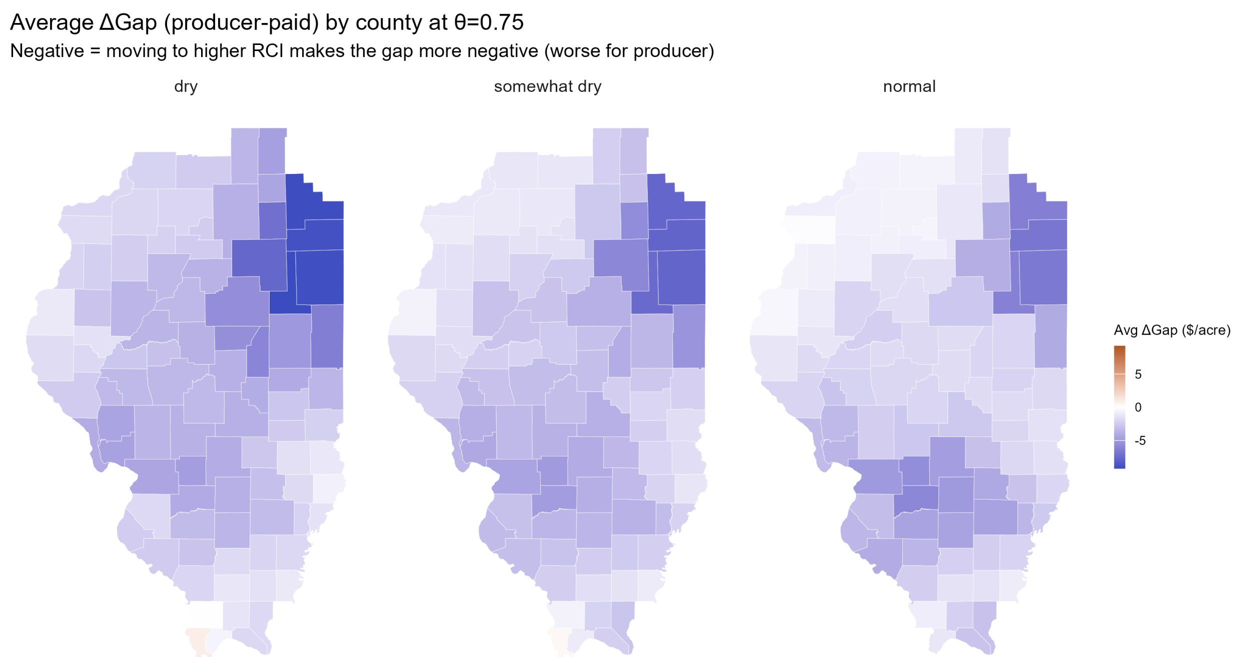

Figure 5: Average change in the producer-paid pricing gap, ΔGap, for within-field transitions from low to high RCI at coverage $\theta = 0.75$. Panels facet VPD regimes in fixed order (dry, somewhat dry, normal). Colors are harmonized across panels (blue = gap shrinks, red = gap widens/non-shrinks). Consistent with the shortfall–mean plane, many counties register negative ΔGap where higher RCI reduces downside shortfalls but premiums remain mean-based.

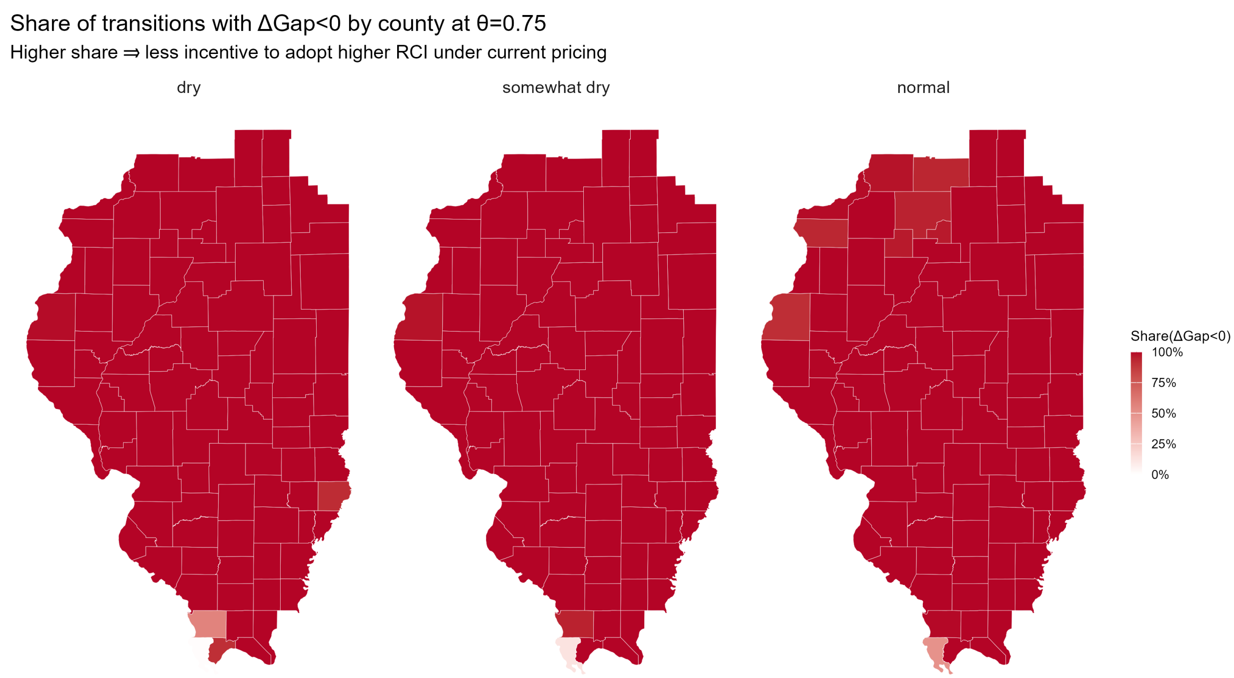

Figure 6: Share of within-field transitions with ΔGap<0 by county at $\theta = 0.75$ (0–1). Higher shares indicate places where moving to a higher RCI more often *reduces* the pricing gap—i.e., weakens the insurance-based incentive to adopt more complex rotations under APH-only pricing. Panels are ordered (dry, somewhat dry, normal) with a common legend.

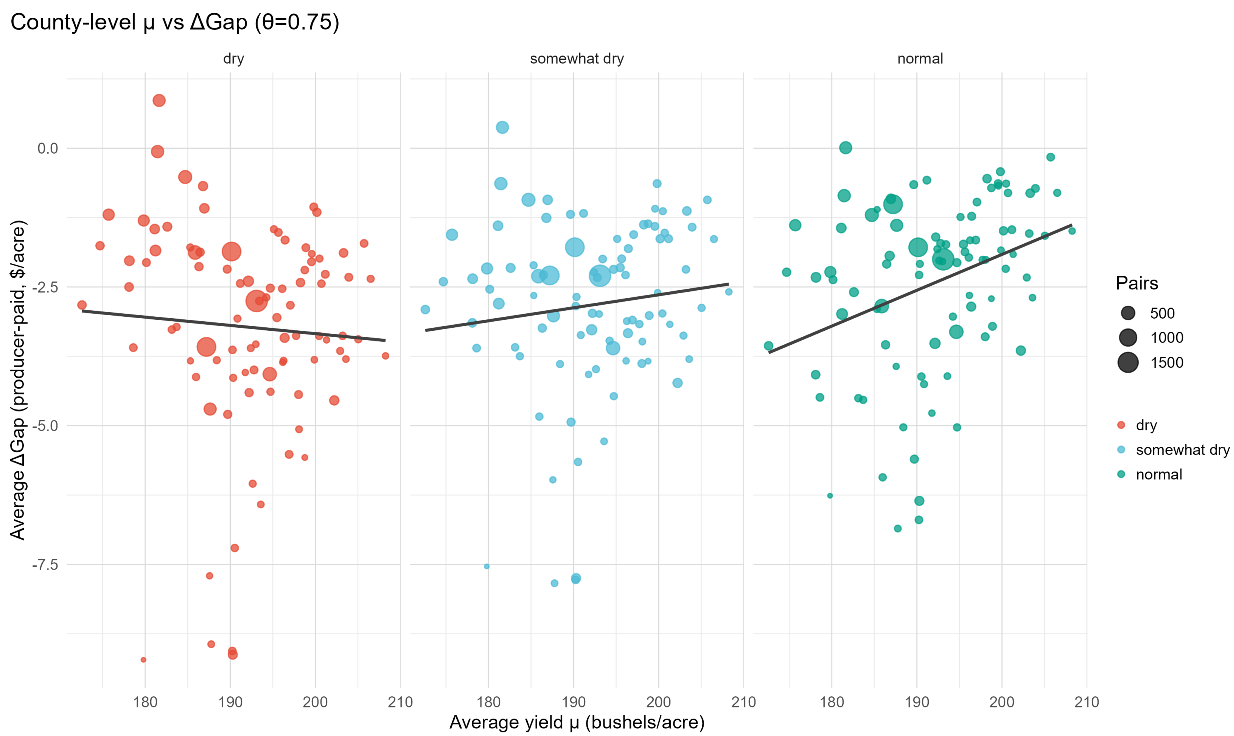

Figure 7: County mean productivity μ (bushels/acre) versus average ΔGap from low to high RCI at $\theta = 0.75$, faceted by VPD regime (dry, somewhat dry, normal). Points are sized by the number of within-field pairs contributing to the county. Fitted lines (OLS) summarize the regime-specific association between productivity and the change in the insurance wedge.

References

Bowles, T. M., Mooshammer, M., Socolar, Y., Calderón, F., Cavigelli, M. A., Culman, S. W., Deen, B., Drury, C. F., Garcia y Garcia, A., Gaudin, A. C. M., Harkcom, W. S., Lehman, R. M., Osborne, S. L., Robertson, G. P., Salerno, J., Schmer, M. R., Strock, J., and Grandy, A. S. (2020). Long-term evidence shows that crop-rotation diversification increases agricultural resilience to adverse growing conditions in north america. *One Earth*, 2(3):284–293.

Burchfield, E. K., Crossley, M. S., and Nelson, K. S. (2024). Rotational complexity across US counties is currently insufficient to observe yield gains in major crops. *Environmental Research Letters*, 19(4):044024.

Hadar, J. and Russell, W. R. (1969). Rules for ordering uncertain prospects. *American Economic Review*, 59(1):25–34.

Lobell, D. B., Thau, D., Seifert, C., Engle, E., and Little, B. (2015). A scalable satellite-based crop yield mapper. *Remote Sensing of Environment*, 164:324–333.

Manski, S., Socolar, Y., Goldstein, B., Pizzo, G., Ahmed, Z., Connor, L., Cross, H., Fettes, K., McLauchlan, A., Pham, L., Viens, F., and Bowles, T. M. (2024). Diversified crop rotations mitigate agricultural losses from dry weather. *agriRxiv*.

Pizzo, G. (2025). *Rotation Complexity and Yield Risk: Field-Level Evidence from the U.S. Corn Belt*. Ph.d. dissertation, Michigan State University. Accessed: ¡YYYY-MM-DD¿.

Rothschild, M. and Stiglitz, J. E. (1970). Increasing risk. i. a definition. *Journal of Economic Theory*, 2(3):225–243.

Socolar, Y., Goldstein, B. R., de Valpine, P., and Bowles, T. M. (2021). Biophysical and policy factors predict simplified crop rotations in the US midwest. *Environmental Research Letters*, 16(5):054045.

Tsiboe, F., Turner, D., and Yu, J. (2025). Utilizing large-scale insurance data sets to calibrate sub-county level crop yields. *Journal of Risk and Insurance*, 92(1):139–165.

Appendix

A Proof of Lemma 2.2

Proof of Lemma 2.2. Let $\phi_\theta(z) := (\theta - z)^+$. Since ϕ_θ is convex (as the maximum of two affine functions), convex-order dominance implies

$$\mathbb{E}\big[\phi_\theta(Z')\big] \;\geq\; \mathbb{E}\big[\phi_\theta(Z)\big]$$

whenever $Z \preceq_{cx} Z'$ and $\mathbb{E}[Z] = \mathbb{E}[Z']$. Substituting $\phi_\theta(z) = (\theta - z)^+$ yields

$$\mathbb{E}\big[(\theta - Z')^+\big] \;\geq\; \mathbb{E}\big[(\theta - Z)^+\big].$$

Because $S(\theta; G) = \mathbb{E}[(\theta - Z)^+]$, it follows that $S(\theta; G)$ is increasing in the convex order. The equivalent SSD statement follows from the standard relationship between convex order and second-order stochastic dominance under equal means. $\square$

Notation remark. With equal means, $Z \preceq_{cx} Z'$ ("Z' more spread") corresponds to $Z \succeq_{SSD} Z'$ in the SSD convention; take care with direction if using SSD notation in the text.

B Derivations and Comparative Statics

Comparative statics. (i) If a (higher RCI) *only* improves variance/tails holding μ_i fixed, then $d\mu_i/da = 0$, hence $d\Pi_i/da = 0$ by (10) (the *no-variance-pricing* property, cf. (7)), while $d\,[I_i]/da = P\,\mu_i\,\frac{dS}{da} < 0$. Insurance therefore attenuates the private value of tail-risk reduction unless premiums co-move with S.

(ii) If, in addition, a raises μ_i, then under the empirical implementation

$$\frac{d\Pi_i}{da} \;=\; (1-\zeta)\Big[A_i(\theta)(1+\beta_j)\mu_i(a)^{\beta_j} + P\,\theta\,\delta_j\Big]\frac{d\mu_i}{da}, \tag{B.1}$$

so the *dollar* premium rises mechanically when the bracketed term is positive. By contrast, the *rate*

$$\text{rate}_i(a) \equiv \frac{\Pi_i(a)}{P\,\mu_i(a)} = (1-\zeta)\left[\frac{A_i(\theta)}{P}\mu_i(a)^{\beta_j} + \theta\,\delta_j\right]$$

contains both a mean-elastic term and the fixed-rate component. Net private incentives in the mean channel are therefore ambiguous (revenue and indemnity-mean effects vs. higher dollar premiums), while pure tail improvements $P\,\mu_i\,\frac{dS}{da}$ remain unpriced under the baseline CRF.

These identities dovetail with the priced objects introduced in the main text:

$$\text{EI}(a) = P\,\mu(a)\,S(\theta; G_a), \qquad \Pi(a) = (1-\zeta)\Big[A_i(\theta)\mu(a)^{1+\beta_j} + P\,\theta\,\delta_j\,\mu(a)\Big],$$

and motivate the wedge analysis

$$\text{Gap}(a) = \text{EI}(a) - \Pi(a)$$

in Section 4.

C Non-monotone practice effects and gap changes

Let $\Delta\mu/\mu \equiv (\mu^{\text{hi}} - \mu^{\text{lo}})/\mu^{\text{lo}}$ and $\Delta S \equiv S^{\text{hi}} - S^{\text{lo}}$. Fix a pivot $a^* a^{\text{lo}}$ and evaluate

$$B_* \;\equiv\; P\,S(\theta; G_{a^*}) \;-\; (1-\zeta)\Big[A_i(\theta)(1+\beta_j)\mu_i(a^*)^{\beta_j} + P\,\theta\,\delta_j\Big].$$

Then the first-order (Aumann–Shapley) zero-gap boundary is

$$B_* \cdot \frac{\Delta\mu}{\mu_i(a^*)} \;+\; P\,\Delta S \;=\; 0. \tag{C.1}$$

Hence the sign of ΔGap_i is governed by $\text{sgn}\big(B_* \cdot \Delta\mu/\mu_i(a^*) + P\,\Delta S\big)$. We summarize the four quadrants:

- **Mean ↑, Tail ↓** ($\Delta\mu/\mu > 0$, $\Delta S < 0$):

$$\Delta \text{Gap}_i < 0 \quad\Longleftrightarrow\quad \begin{cases} \dfrac{\Delta\mu}{\mu_i(a^*)} \;<\; \dfrac{-\,P\,\Delta S}{B_*}, & \text{if } B_* > 0, \\[2ex] \dfrac{\Delta\mu}{\mu_i(a^*)} \;>\; \dfrac{-\,P\,\Delta S}{B_*}, & \text{if } B_* < 0. \end{cases}$$

- **Mean ↑, Tail ↑** ($\Delta\mu/\mu > 0$, $\Delta S > 0$): if $B_* > 0$ then both terms are positive and $\Delta \text{Gap}_i > 0$. If $B_* < 0$, the mean term is negative while the tail term is positive, so

$$\Delta \text{Gap}_i > 0 \iff P\,\Delta S \;>\; |B_*|\,\frac{\Delta\mu}{\mu_i(a^*)} \quad\text{and}\quad \Delta \text{Gap}_i < 0 \iff P\,\Delta S \;<\; |B_*|\,\frac{\Delta\mu}{\mu_i(a^*)}.$$

- **Mean ↓, Tail ↓** ($\Delta\mu/\mu < 0$, $\Delta S < 0$): if $B_* \geq 0$ then both terms are nonpositive and $\Delta \text{Gap}_i < 0$. If $B_* < 0$, the mean term becomes positive while the tail term is negative, so

$$\Delta \text{Gap}_i < 0 \iff P\,|\Delta S| \;>\; |B_*|\,\frac{|\Delta\mu|}{\mu_i(a^*)}.$$

- **Mean ↓, Tail ↑** ($\Delta\mu/\mu < 0$, $\Delta S > 0$): if $B_* < 0$ both terms are positive and

$\Delta\mathrm{Gap}_i > 0$. If $B_* > 0$, the mean term is negative while the tail term is positive, so

$$\Delta\mathrm{Gap}_i > 0 \iff P\,\Delta S \; > \; B_* \frac{|\Delta\mu|}{\mu_i(a^*)}, \quad \text{and otherwise } \Delta\mathrm{Gap}_i < 0.$$

These cases map one-to-one to quadrants in the $(\Delta\mu/\mu, \Delta S)$ plane. Our figures overlay empirical points on the first-order boundary (C.1) for visual diagnosis. For exact finite decompositions one may instead use the two-feature Shapley expressions reported in the Methods, which replace (C.1) with a non-linear (in $\Delta\mu$) equality based on $(S^{\mathrm{lo}}+S^{\mathrm{hi}})/2$ and $\mu^{1+\beta}$.

D Predictive Specification and Recovery Details

D.1 Baseline Predictive Model and Distribution Recovery

Model. We estimate a Bayesian linear mixed–effects model at the field level (Manski et al., 2024),

$$
\begin{aligned}
Y_{it} &= \varpi_0 + \varpi_1\,\mathrm{RCI}_i^{(m)} + \varpi_2\,\mathrm{RCI}_{it}^{(w)} + \varpi_3\,\mathrm{VPD7}_{it} + \varpi_4\,\mathrm{NCCPI}_i + \varpi_5\big(\mathrm{VPD7}_{it}\times\mathrm{RCI}_i^{(m)}\big) \\
&\quad + \varpi_6\big(\mathrm{VPD7}_{it}\times\mathrm{RCI}_{it}^{(w)}\big) + \varpi_7\big(\mathrm{NCCPI}_i\times\mathrm{RCI}_i^{(m)}\big) + \varphi\, t \;+\; \eta_i \;+\; \varepsilon_{it},
\end{aligned}
\tag{D.1}
$$

with field random intercepts $\eta_i \sim \mathcal{N}(0, \sigma_\eta^2)$ and idiosyncratic error $\varepsilon_{it} \sim \mathcal{N}(0, \sigma^2)$, independent across i, t.

Design and scenarios. Let i index fields and define a scenario $c = (i, a, s)$ that sets practice and weather to $a \in \mathcal{A}$ (an RCI level) and $s \in \mathcal{S}$ (a VPD regime). For evaluation we hold $\mathrm{RCI}_i^{(m)}$ and NCCPI_i fixed at field i's values, set

$$
\mathrm{RCI}_{it}^{(w)} = \mathrm{RCI}^{(w)}(a), \qquad \mathrm{VPD7}_{it} = \mathrm{VPD7}(s),
$$

and anchor the time component at a reference year $t^\dagger$ (we use the last sample year so that φt reflects current practice; in year–FE versions we set $\lambda_{t^\dagger} = 0$ under $\sum_t \lambda_t = 0$).

Neighborhood fitting (county–by–county). To capture spatial heterogeneity while limiting edge effects, we fit (D.1) separately for each focal county j. The estimation pool for j includes all fields in j and its adjacent counties $N(j)$ ("neighborhood"); predictions and diagnostics are computed only for $i \in j$. We require a minimum sample size in the pool (e.g., $\geq$500 field–years) to fit the county model; otherwise the county is dropped. Random intercepts η_i are estimated for all fields used in fitting; when predicting within j we use the posterior for the corresponding η_i.

Likelihood and priors. The likelihood is Gaussian as implied by (D.1): $\varepsilon_{it} \sim \mathcal{N}(0, \sigma^2)$ and field intercepts $\eta_i \sim \mathcal{N}(0, \sigma_\eta^2)$. Fixed–effect coefficients $\{\varpi_k\}_{k=0}^{7}$ and the time term φ use weakly–informative Normal priors centered at zero *except* for the key predictors for which we adopt literature–informed priors consistent with the main text: ϖ_1, ϖ_2 (RCI$^{(m)}$, RCI$^{(w)}$) $\sim \mathcal{N}(0.26,\, 0.26)$, ϖ_3 (VPD) $\sim \mathcal{N}(-0.20,\, 0.20)$, ϖ_4 (NCCPI) $\sim \mathcal{N}(1.15,\, 1.15)$, φ (year trend when used) $\sim \mathcal{N}(0.15,\, 0.15)$; interaction terms receive weakly–informative $\mathcal{N}(0,\ s^2)$ priors (diffuse s) or flat priors where noted. Standard deviations σ_η and σ follow half–Student–t priors on the SD scale with ν=3 (brms default shape), providing gentle regularization while allowing the data to dominate.

Computation and diagnostics. We fit the model in `brms` (Stan backend) with 3 chains and at least 2,500 post–warmup draws per chain (default warmup 1,000; increased as needed). Sampler tuning follows standard practice (`adapt_delta` raised up to 0.99 and `max_treedepth` up to 12 if required). Convergence is assessed using rank–normalized $\hat{R} < 1.01$, bulk and tail effective sample sizes $> 1{,}000$ for all reported parameters, and the absence of divergent transitions; we also check energy–BFMI and that Monte Carlo standard errors are small relative to posterior standard deviations. Model adequacy is summarized by the Bayesian R^2 and by empirical coverage of 95% posterior predictive intervals at the county level, complemented with standard posterior predictive checks.

Posterior predictive engine $F_i^{(c)}$. For each scenario $c = (i, a, s)$ and each posterior draw b of $(\varpi, \varphi, \{\lambda_t\}, \eta_i, \sigma)$, we compute the linear predictor at $x_i(a, s)$ and draw

$$Y_i^{(b)}(c) \;=\; \big(\varpi_0^{(b)} + \textstyle\sum_{k=1}^{7} \varpi_k^{(b)} x_{ik}(a,s) + \varphi^{(b)} t^\dagger + \lambda_{t^\dagger}^{(b)}\big) \;+\; \eta_i^{(b)} \;+\; \varepsilon^{(b)}, \quad \varepsilon^{(b)} \sim \mathcal{N}(0, \sigma^{2\,(b)}).$$

Collecting B draws yields the posterior predictive distribution $F_i^{(c)}$. From these we compute the predictive mean $\mu_i^{(c)} = \mathbb{E}[Y_i \mid c]$, percentiles $Q^{(c)}(p)$, and define the relative–yield CDF $G^{(c)}$ via $R = Y/\mu_i^{(c)}$.

D.2 Interpolation and recovery robustness

Because both the predictive mean μ and the shortfall functional $S(\theta; G)$ are recovered from stored percentiles when full posterior draws are not retained, we assess numerical robustness along two margins: recovery of the APH proxy and recovery of the lower-tail shortfall functional.

First, the predictive-mean proxy is extremely stable across alternative feasible implementations, including midpoint weighting, trimmed and winsorized versions of the percentile curve, and a median-based alternative. In premium space and pair space, the resulting objects are nearly unchanged: premium ranks are essentially identical across implementations, county-average ΔGap ranks are virtually unchanged, and pairwise ΔGap signs coincide almost perfectly with the baseline specification. These results indicate that the paper's main comparative statics do not depend on a particular numerical implementation of the APH proxy.

Second, the lower-tail shortfall functional is mechanically more sensitive to interpolation choice than the predictive mean itself, as expected for an object that depends on probability mass near the guarantee threshold. To assess this sensitivity in result space, we recompute cell-level shortfalls, pairwise ΔGap classifications, zero-gap-plane locations, and Shapley decompositions under three implementations: baseline linear interpolation, monotone spline interpolation, and a discontinuous step interpolant. Relative to the baseline linear method, the monotone spline preserves 98.9% of pairwise ΔGap signs and 92.1% of zero-gap-plane classifications, while the tail channel remains dominant in 93.3% of $\theta\times$ VPD cells. By contrast, the discontinuous step interpolant is materially less stable and is therefore best interpreted as a conservative stress test rather than a preferred recovery method.

These checks indicate that the paper's main conclusions are highly robust within the class of continuous recovery methods. The producer-paid pricing wedge, the shortfall–mean plane classifications, and the dominance of the tail channel do not depend on a knife-edge interpolation choice.

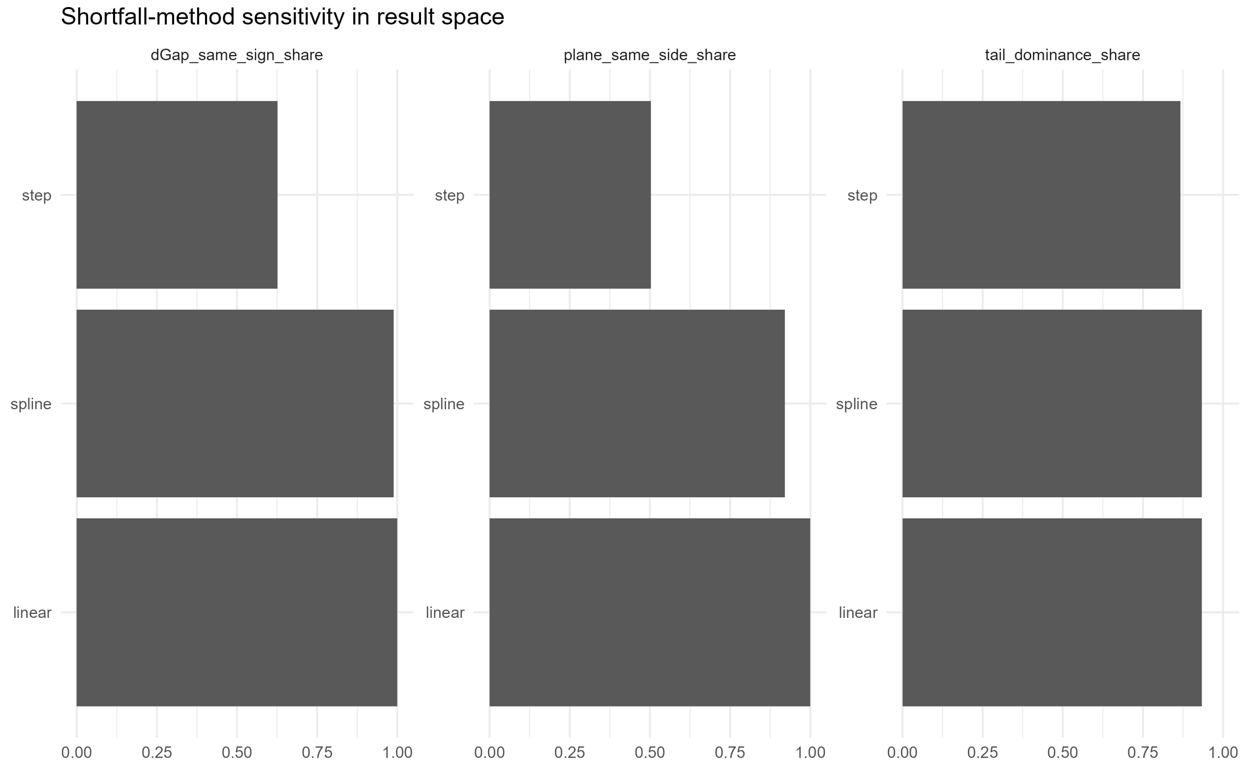

Figure 8: Shortfall-method sensitivity in result space. Bars report robustness of the main result-space objects relative to the baseline linear interpolation. The left panel shows the share of pairwise transitions with the same sign of ΔGap, the middle panel shows the share of transitions that remain on the same side of the zero-gap boundary in the shortfall–mean plane, and the right panel shows the share of $\theta\times$ VPD cells in which the tail Shapley contribution exceeds the mean contribution in absolute value. The monotone spline implementation closely matches the baseline linear method, whereas the discontinuous step interpolant is materially less stable.

Table 4: Robustness of main result-space objects to shortfall interpolation method

Method	Same sign of ΔGap	Same side of zero-gap plane	Tail dominates
Linear	1.000	1.000	0.933
Spline	0.989	0.921	0.933
Step	0.627	0.503	0.867

Entries compare each implementation of the shortfall functional to the baseline linear interpolation. "Tail dominates" denotes the share of $\theta\times$ VPD cells in which the absolute Shapley contribution of the tail channel exceeds that of the mean channel.

D.3 Recovery of (μ, S) from posteriors or percentiles

From $F_i^{(c)}$ we compute $\mu_i^{(c)} = \mathbb{E}[Y_i \mid c]$ and define the relative-yield CDF $G^{(c)}$ via $R \equiv Y/\mu_i^{(c)}$. When only percentiles $\{Q^{(c)}(p_k)\}$ are stored, we enforce monotonicity by rearrangement (cumulative maxima on p), compute $\mu_i^{(c)}$ by trapezoidal integration (with coverage correction if needed), construct the relative quantile $R_q^{(c)}(p) = Q^{(c)}(p)/\mu_i^{(c)}$, invert to recover $G^{(c)}$, and evaluate

$$S\big(\theta; G^{(c)}\big) = \theta\, G^{(c)}(\theta) - \int_0^\theta G^{(c)}(r)\, dr,$$

on a dense grid (we use 600 points). As a cross-check we compute the equivalent lower partial moment $S = \mathbb{E}[(\theta - R)^+]$ on the percentile grid; results are numerically identical up to rounding.

D.4 County-level RMA CRF factors and premium construction

We compute CRF-implied premiums using county-level Continuous Rating Formula factors obtained from RMA administrative baserate data for Illinois corn. For each county, we recover the reference rate α_j, APH exponent β_j, reference yield $\bar{R}_j$, and fixed-rate component δ_j. When multiple administrative records are available within a county, we aggregate them to the county level by taking the mean of each component. Missing county factors, when they occur, are replaced with the statewide median for the corresponding component.

Because unit-level APH is not observed in our data, we proxy APH with the predictive mean $\mu_i^{(c)}$ for each field–RCI–weather scenario $c = (i, a, s)$. Given coverage θ and price election P, we compute

$$A_{ij}(\theta) = \frac{P\, \theta\, \alpha_j}{\bar{R}_j^{\beta_j}},$$

and the CRF-implied premium

$$\Pi_i^{(c)} = A_{ij}(\theta)\, (\mu_i^{(c)})^{1+\beta_j} + P\, \theta\, \delta_j\, \mu_i^{(c)}.$$

This direct administrative implementation replaces the calibration-based approximation used in an earlier draft. In practice, the resulting estimates are quantitatively very similar, and all qualitative conclusions are unchanged.

D.5 Calibration-based approximation

We approximated CRF pricing coefficients using a calibration procedure for K_j (and, where needed, β) within county–plan–coverage–unit strata as a robustness check. Specifically, K_j was anchored to a county benchmark subset through a moment-matching condition based on expected indemnities and predictive means. Because county-level RMA baserate factors later became available, the main text now uses the direct administrative implementation. The calibration-based approximation produces quantitatively similar estimates and does not change the paper's substantive conclusions.

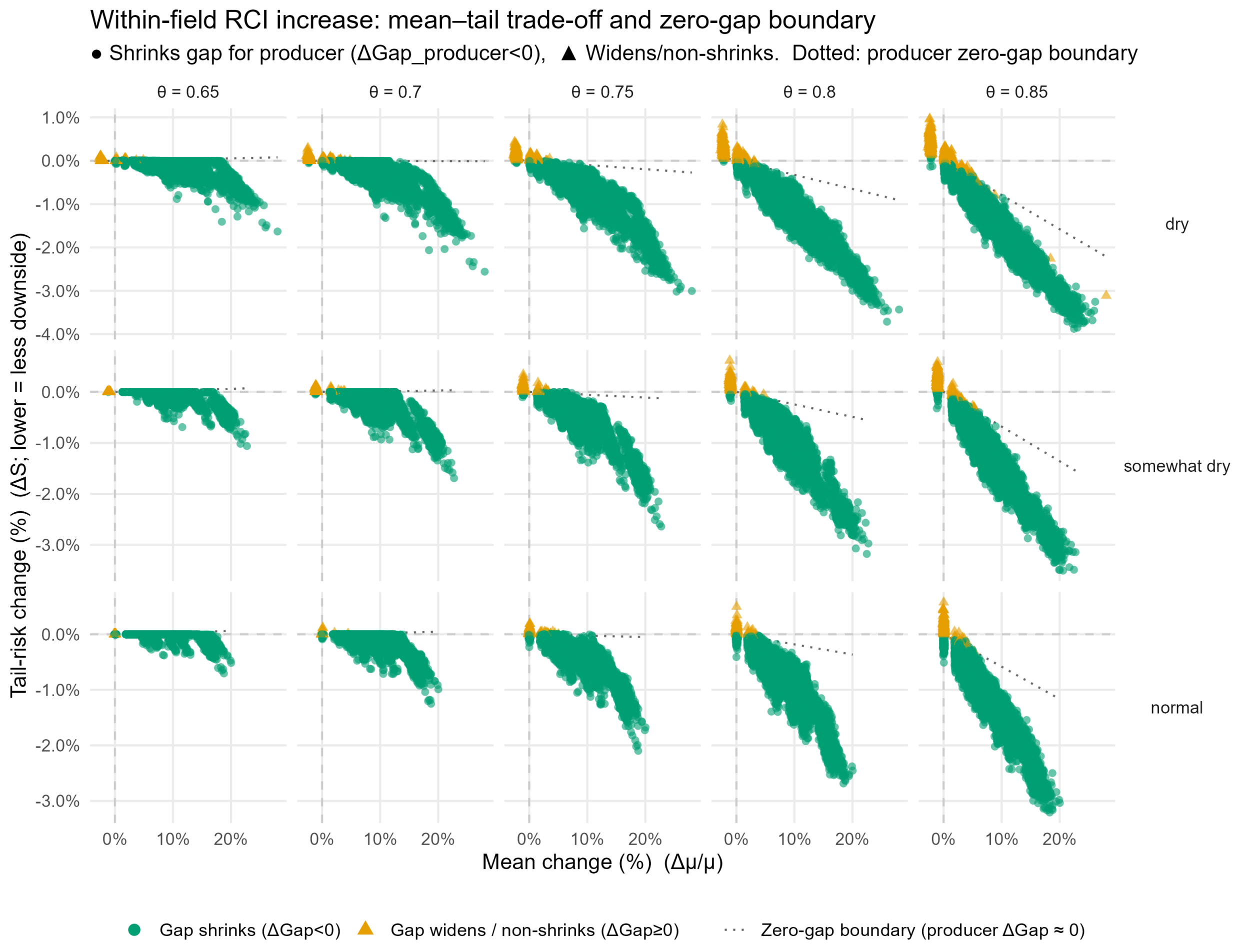

Figure 9: Within-field RCI increase: mean–tail trade-off and zero-gap boundary.
Panels: rows = VPD (dry, somewhat dry, normal); columns = coverage $\theta \in \{0.65, 0.70, 0.75, 0.80, 0.85\}$.
Axes: x = mean change (%, $\Delta\mu/\mu$); y = tail-risk change (%, ΔS; down = less downside).
Markers: green circles = gap shrinks for producer ($\Delta\text{Gap} < 0$); orange triangles = widens / non-shrinks ($\Delta\text{Gap} \geq 0$).
Line: dotted = group-wise zero-gap boundary ($\Delta\text{Gap} \approx 0$); below the line → mostly green, above → mostly orange.
Note: In these data EI is typically below Premium; "gap shrinks" often means the gap becomes *more negative*, not necessarily a welfare gain.

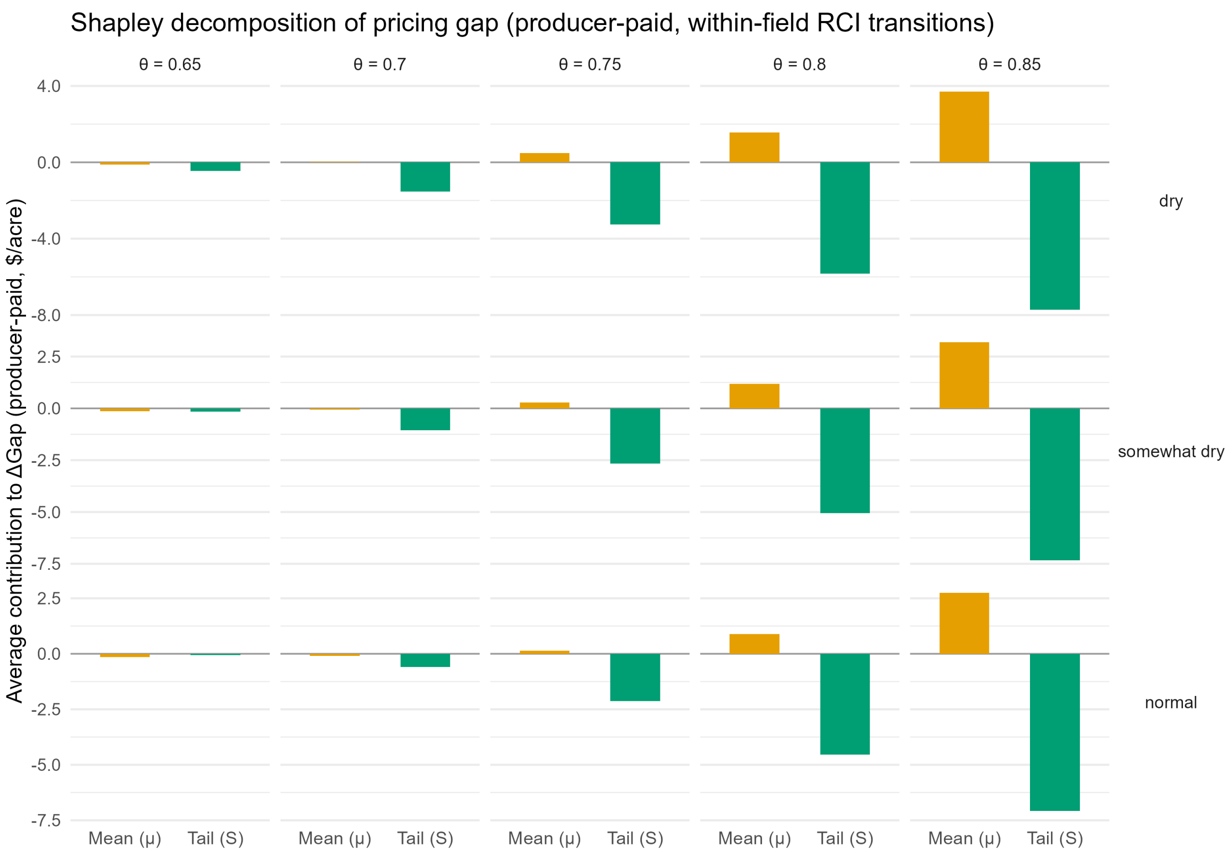

Figure 10: Shapley decomposition of the pricing–gap change for within–field RCI transitions. Bars show mean finite–difference Shapley contributions by channel (rows: VPD regime; columns: coverage θ).

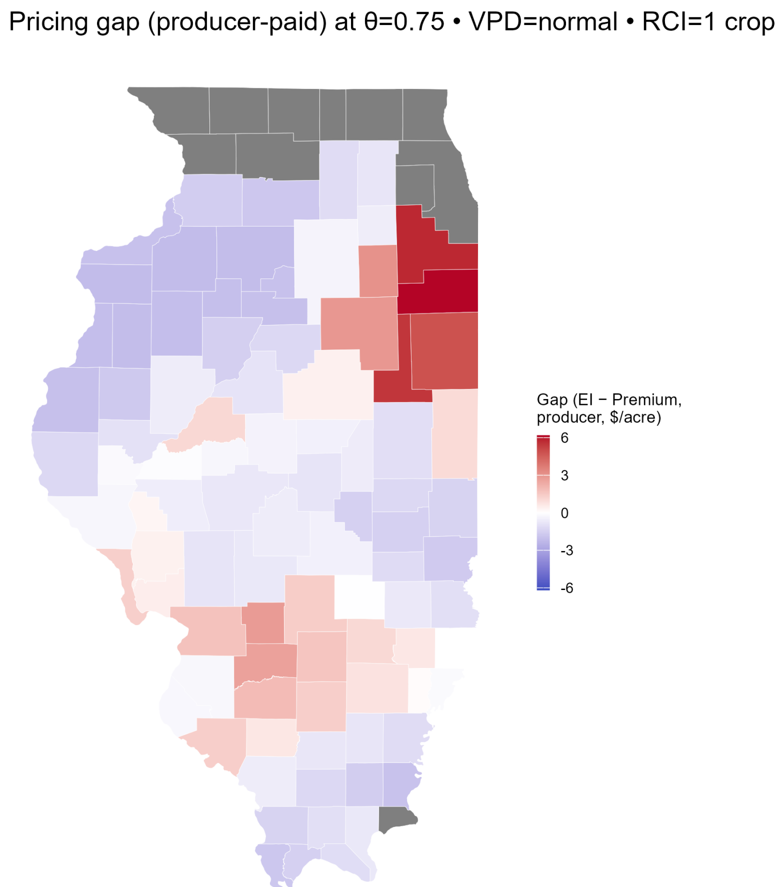

Figure 11: County-level producer-paid pricing gap at coverage $\theta = 0.75$ under the *normal* VPD regime and low rotational complexity (RCI = '1 crop'). Colors use a common blue–white–red scheme (white at zero): red = positive gap (expected indemnities exceed producer-paid premiums), blue = negative gap. This map anchors the geography of the wedge between distribution-sensitive expected indemnities and APH-based CRF premiums before considering within-field practice changes.

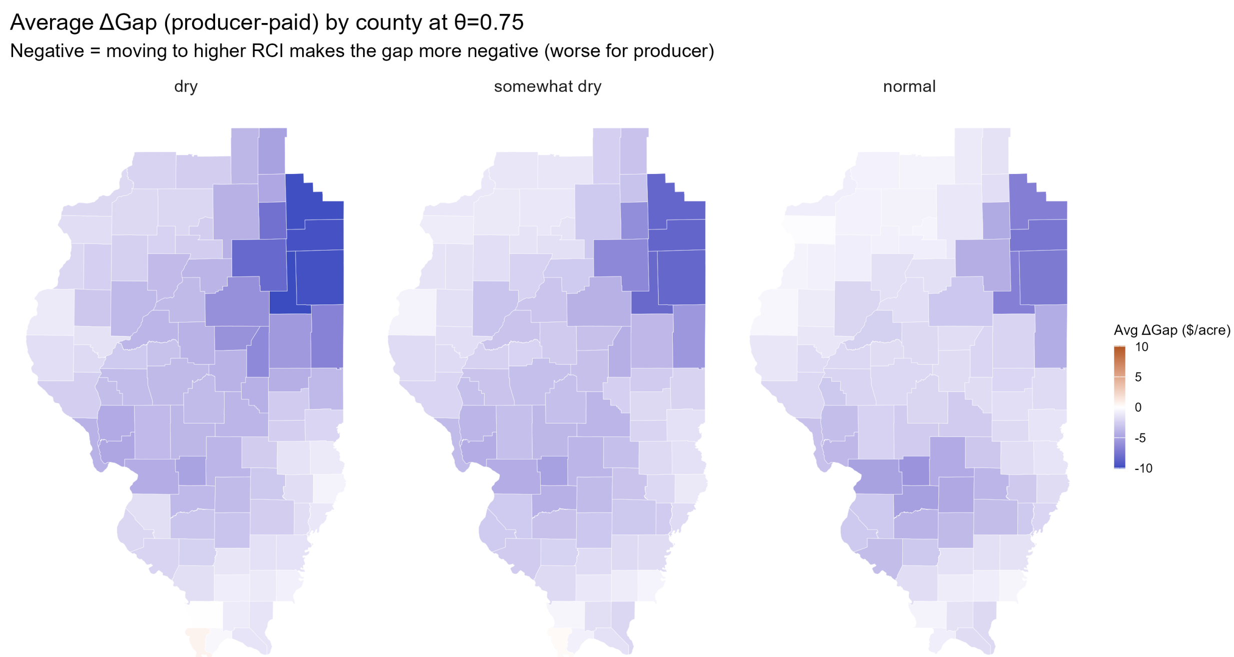

Figure 12: Average change in the producer-paid pricing gap, ΔGap, for within-field transitions from low to high RCI at coverage $\theta = 0.75$. Panels facet VPD regimes in fixed order (dry, somewhat dry, normal). Colors are harmonized across panels (blue = gap shrinks, red = gap widens/non-shrinks). Consistent with the shortfall–mean plane, many counties register negative ΔGap where higher RCI reduces downside shortfalls but premiums remain mean-based.

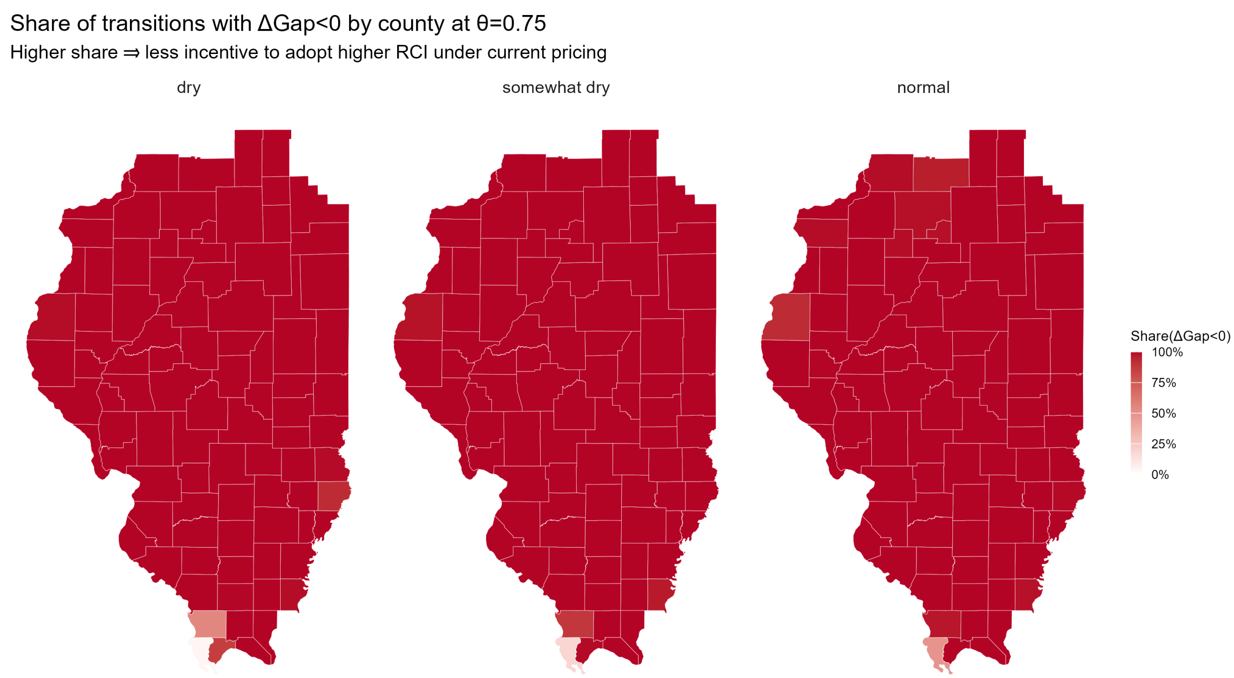

Figure 13: Share of within-field transitions with $\Delta\text{Gap} < 0$ by county at $\theta = 0.75$ (0–1). Higher shares indicate places where moving to a higher RCI more often *reduces* the pricing gap—i.e., weakens the insurance-based incentive to adopt more complex rotations under APH-only pricing. Panels are ordered (dry, somewhat dry, normal) with a common legend.

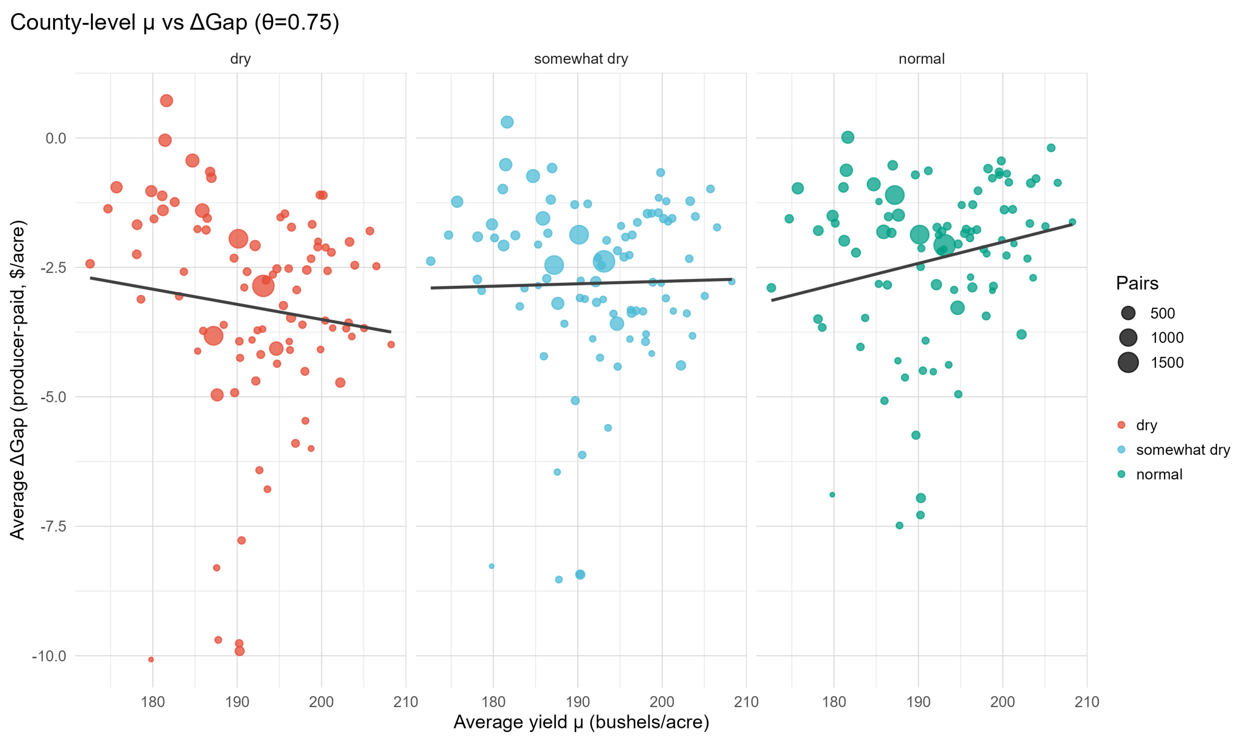

Figure 14: County mean productivity μ (bushels/acre) versus average ΔGap from low to high RCI at $\theta = 0.75$, faceted by VPD regime (dry, somewhat dry, normal). Points are sized by the number of within-field pairs contributing to the county. Fitted lines (OLS) summarize the regime-specific association between productivity and the change in the insurance wedge.

E Volatility–Augmented CRF Pricing

The baseline administrative implementation implies a producer-paid premium of

$$(1-\zeta)\Big[A_i(\theta)\mu_i^{1+\beta_j} + P\,\theta\,\delta_j\,\mu_i\Big].$$

We consider a minimal extension that applies a volatility elasticity only to the mean-elastic component while leaving the fixed-rate component unchanged. Define

$$\tau_{ijtu}^{\gamma} \;=\; \alpha_j\Big(\tfrac{\bar{y}_{ijt}}{R_j}\Big)^{\beta_j}\Big(\tfrac{\sigma_{ijt}}{\bar{\sigma}_j}\Big)^{\gamma} + \delta_j, \tag{E.1}$$

so that the producer-paid dollar premium is

$$\Pi_{ijtu}^{\text{prod}} = (1-\zeta(\theta_{it},u))\Big[A_i(\theta_{it})\,\mu_i^{1+\beta_j}\Big(\tfrac{\sigma_{ijt}}{\bar{\sigma}_j}\Big)^{\gamma} + P\,\theta_{it}\,\delta_j\,\mu_i\Big]. \tag{E.2}$$

A positive γ penalizes unusually volatile units; a negative γ discounts them.[13][14]

Proposition (Zero–gap boundary with volatility adjustment, producer–paid). Let the producer-paid gap be

$$\text{Gap}_i^{\text{prod}} = P\,\mu_i\,S(\theta;G_a) - (1-\zeta)\Big[A_i(\theta)\mu_i^{1+\beta_j}\Big(\tfrac{\sigma_i}{\bar{\sigma}}\Big)^{\gamma} + P\,\theta\,\delta_j\,\mu_i\Big].$$

A first-order approximation for a practice change $a^{\text{low}} \to a^{\text{high}}$ yields the (Aumann–Shapley) zero-gap boundary in the $(\Delta\mu/\mu,\ \Delta S,\ \Delta\sigma/\sigma)$ space:

$$P\,\mu\,\Delta S + \Big[P\,S - (1-\zeta)\Big(A_i(\theta)(1+\beta_j)\mu^{\beta}\Big(\tfrac{\sigma}{\bar{\sigma}}\Big)^{\gamma} + P\,\theta\,\delta_j\Big)\Big]\mu\,\tfrac{\Delta\mu}{\mu} - (1-\zeta)\,A_i(\theta)\,\mu^{1+\beta_j}\Big(\tfrac{\sigma}{\bar{\sigma}}\Big)^{\gamma}\gamma\cdot\tfrac{\Delta\sigma}{\sigma} \;\approx\; 0, \tag{E.3}$$

[13] Sign convention: with $\Pi^{\text{prod}} \propto (\sigma/\bar{\sigma})^{\gamma}$ on the mean-elastic component, a penalty for volatility corresponds to $\gamma > 0$. Within-field contrasts keep θ and u fixed, so $\zeta(\theta, u)$ is constant locally.

[14] Sign convention: with $\Pi^{\text{prod}} \propto (\sigma/\bar{\sigma})^{\gamma}$, a penalty for volatility corresponds to $\gamma > 0$. Within–field contrasts keep θ and u fixed, so $\zeta(\theta, u)$ is constant locally.

where all terms are evaluated at a local reference point. Relative to the baseline boundary without volatility,

$$P\,\mu\,\Delta S+\Big[P\,S-(1-\zeta)\Big(A_i(\theta)(1+\beta_j)\mu^{\beta_j}+P\,\theta\,\delta_j\Big)\Big]\mu\,\tfrac{\Delta\mu}{\mu}=0,$$

the σ-term rotates the boundary: reductions in idiosyncratic volatility ($\Delta\sigma/\sigma<0$) shift observations toward the negative side ($\Delta\text{Gap}^{\text{prod}}<0$) even when ΔS is modest. Equivalently, dividing (E.3) by $P\,\mu$ gives

$$\Delta S=-\Big(S-\tfrac{(1-\zeta)\big[A_i(\theta)(1+\beta_j)\mu^{\beta_j}(\sigma/\bar\sigma)^{\gamma}+P\,\theta\,\delta_j\big]}{P}\Big)\tfrac{\Delta\mu}{\mu}+\underbrace{\tfrac{(1-\zeta)\,A_i(\theta)\,\mu^{\beta_j}(\sigma/\bar\sigma)^{\gamma}\gamma}{P}}_{C(\gamma)}\cdot\tfrac{\Delta\sigma}{\sigma}.$$

Corollary (Sufficient condition under $\mu'\geq 0$, $dS/da\leq 0$, $\sigma'\leq 0$, producer–paid). Totally differentiating gives

$$\begin{aligned}\frac{d\,\text{Gap}_i^{\text{prod}}}{da}\;&=\underbrace{\mu_i'(a)\Big[P\,S-(1-\zeta)\Big(A_i(\theta)(1+\beta_j)\mu_i^{\beta_j}\Big(\frac{\sigma_i}{\bar\sigma}\Big)^{\gamma}+P\,\theta\,\delta_j\Big)\Big]}_{\text{mean channel}}\\&\quad+\;\underbrace{P\,\mu_i\,\frac{dS}{da}}_{\text{tail channel}}\\&\quad-\;\underbrace{(1-\zeta)\,A_i(\theta)\,\mu_i^{1+\beta_j}\Big(\frac{\sigma_i}{\bar\sigma}\Big)^{\gamma}\gamma\,\frac{d\sigma_i/da}{\sigma_i}}_{\text{volatility channel}}.\end{aligned}$$

When $\mu_i'(a)\geq 0$, $dS/da\leq 0$, and $d\sigma_i/da\leq 0$, a sufficient condition for $d\,\text{Gap}_i^{\text{prod}}/da<0$ is

$$\frac{\mu_i'(a)}{\mu_i(a)}\;<\;\frac{-\,P\,\frac{dS}{da}\;+\;(1-\zeta)\,A_i(\theta)\,\mu_i^{1+\beta_j}\,\big(\frac{\sigma_i}{\bar\sigma}\big)^{\gamma}\,\gamma\,\frac{d\sigma_i/da}{\sigma_i}}{P\,S\;-\;(1-\zeta)\Big[A_i(\theta)(1+\beta_j)\mu_i^{\beta_j}\big(\frac{\sigma_i}{\bar\sigma}\big)^{\gamma}+P\,\theta\,\delta_j\Big]}. \tag{E.4}$$

Thus a small volatility elasticity $\gamma>0$ partially internalizes tail improvements by pricing units that become idiosyncratically more stable on a producer-paid basis. In Section 3, we estimate σ_{ijt} from predictive posteriors (e.g., posterior sd or a quantile–based dispersion for robustness) and implement a counterfactual rating exercise by choosing γ to (i) minimize the median $\Delta\text{Gap}^{\text{prod}}$ or (ii) match the cross–sectional share with $\Delta\text{Gap}^{\text{prod}} \approx 0$ at each θ. This quantifies how much of the observed negative producer–paid gap can be closed by a one–parameter volatility adjustment.

E.1 Implementation and summary

For each within-field transition $a^{\text{low}} \to a^{\text{high}}$ inside a VPD regime s, let σ_{low} and σ_{high} denote the predictive dispersion of Y for the two practice states.[15] Our default change metric is the log–difference $\Delta \log \sigma = \log \sigma_{\text{high}} - \log \sigma_{\text{low}}$; when logs are not used we report the first–order equivalent ratio $\Delta\sigma/\sigma = (\sigma_{\text{high}} - \sigma_{\text{low}})/\sigma_{\text{low}}$. To make σ dimensionless, we divide by a county benchmark

$$\bar{\sigma}_j \;=\; \text{median}_{i\in j,\, c}\{\sigma_i^{(c)}\},$$

with a `FIPS` default (pooled over practice–regime scenarios) and a `FIPS_normal` variant (restricted to *normal* VPD when available). Thus $\sigma^{\text{rel}}_{\text{low}} = \sigma_{\text{low}}/\bar{\sigma}_j$ and $\sigma^{\text{rel}}_{\text{high}} = \sigma_{\text{high}}/\bar{\sigma}_j$. Missing $\bar{\sigma}$ are imputed by the statewide median.

Given (θ, u) and subsidy $\zeta(\theta, u)$, the producer-paid premium under the volatility-augmented CRF is given by equation (E.2); in implementation we use

$$\Pi_i^{\text{prod}}(\gamma) = (1-\zeta)\Big[A_i(\theta)\, \mu_i^{1+\beta_j} \Big(\tfrac{\sigma_i}{\bar{\sigma}_j}\Big)^{\gamma} + P\,\theta\,\delta_j\,\mu_i\Big], \tag{E.5}$$

using the same county-level RMA baserate factors as in the baseline implementation. Producer-paid expected indemnity is $EI = P\,\mu\,S(\theta; G)$, and the within-field change is

$$\Delta\text{Gap}^{\text{prod}}(\gamma) = \Big[P\,\mu\,S - (1-\zeta)\Big(A_i(\theta)\mu^{1+\beta_j}(\tfrac{\sigma}{\bar{\sigma}})^{\gamma} + P\,\theta\,\delta_j\,\mu\Big)\Big]^{\text{high}} - \Big[\,\cdot\,\Big]^{\text{low}}.$$

The Aumann–Shapley first-order boundary with volatility (equation (E.3)) can be writ-

[15] Baseline: posterior sd of Y; robustness: a quantile–based dispersion $(Q_{0.84} - Q_{0.16})/2$.

ten, after dividing by $P\,\mu$, as

$$\Delta S = -\Big(S - \tfrac{(1-\zeta)\left[A_i(\theta)(1+\beta_j)\mu^{\beta_j}(\sigma/\bar{\sigma})^{\gamma}+P\,\theta\,\delta_j\right]}{P}\Big)\,\tfrac{\Delta\mu}{\mu}\;+\;C(\gamma)\cdot\tfrac{\Delta\sigma}{\sigma},\quad C(\gamma)\equiv\tfrac{(1-\zeta)\,A_i(\theta)\,\mu^{\beta_j}(\sigma/\bar{\sigma})^{\gamma}\,\gamma}{P}.$$

Equivalently, define a *tilted* tail metric $\Delta\tilde{S}\equiv\Delta S - C(\gamma)\,\Delta\sigma/\sigma$. The dotted line in Fig. 15 plots $\Delta\tilde{S} = m_\gamma\cdot(\Delta\mu/\mu)$, with slope

$$m_\gamma = -\Big(S - \tfrac{(1-\zeta)\left[A_i(\theta)(1+\beta_j)\mu^{\beta_j}(\sigma/\bar{\sigma})^{\gamma}+P\,\theta\,\delta_j\right]}{P}\Big)$$

evaluated at a local reference point. Reductions in idiosyncratic volatility ($\Delta\sigma/\sigma<0$) then shift observations toward smaller producer wedges even when ΔS is modest.

For each θ we scan $\gamma\in\{-0.60,-0.58,\ldots,0.80\}$ and compute $\Delta\mathrm{Gap}^{\mathrm{prod}}(\gamma)$ across all within–field pairs in that slice. We record

$$\mathrm{med}(\gamma)=\mathrm{median}\{\Delta\mathrm{Gap}^{\mathrm{prod}}(\gamma)\},\qquad \mathrm{share}_{\approx 0}(\gamma)=\Pr\left(|\Delta\mathrm{Gap}^{\mathrm{prod}}(\gamma)|\le\varepsilon\right),$$

with $\varepsilon=\$0.10/\mathrm{acre}$. Our main choice sets $\widehat{\gamma}(\theta)=\arg\min_\gamma|\mathrm{med}(\gamma)|$ (ties $\to$ smaller $|\gamma|$); a complementary choice targets $\mathrm{share}_{\approx 0}(\gamma)\approx 0.5$.

For each coverage level θ, we compare the baseline

$$\Delta\mathrm{Gap}^{\mathrm{prod}}(0)\quad\text{to}\quad\Delta\mathrm{Gap}^{\mathrm{prod}}(\widehat{\gamma}(\theta)),$$

with results visualized in the histogram series, the rotated plane in Fig. 15, and the summary table in Table 5. Two facts emerge. (i) A small $\gamma>0$ substantially *closes* the producer–paid gap at low to intermediate coverage, pushing mass toward $\Delta\mathrm{Gap}^{\mathrm{prod}}\!\approx\!0$—consistent with the σ–tilt observed in the plane. (ii) At $\theta\!\ge\!0.80$, tail sensitivity dominates: dispersion narrows but a pronounced negative mass remains, indicating that modest volatility pricing cannot fully internalize lower–tail improvements where $S(\theta;G)$ matters most.

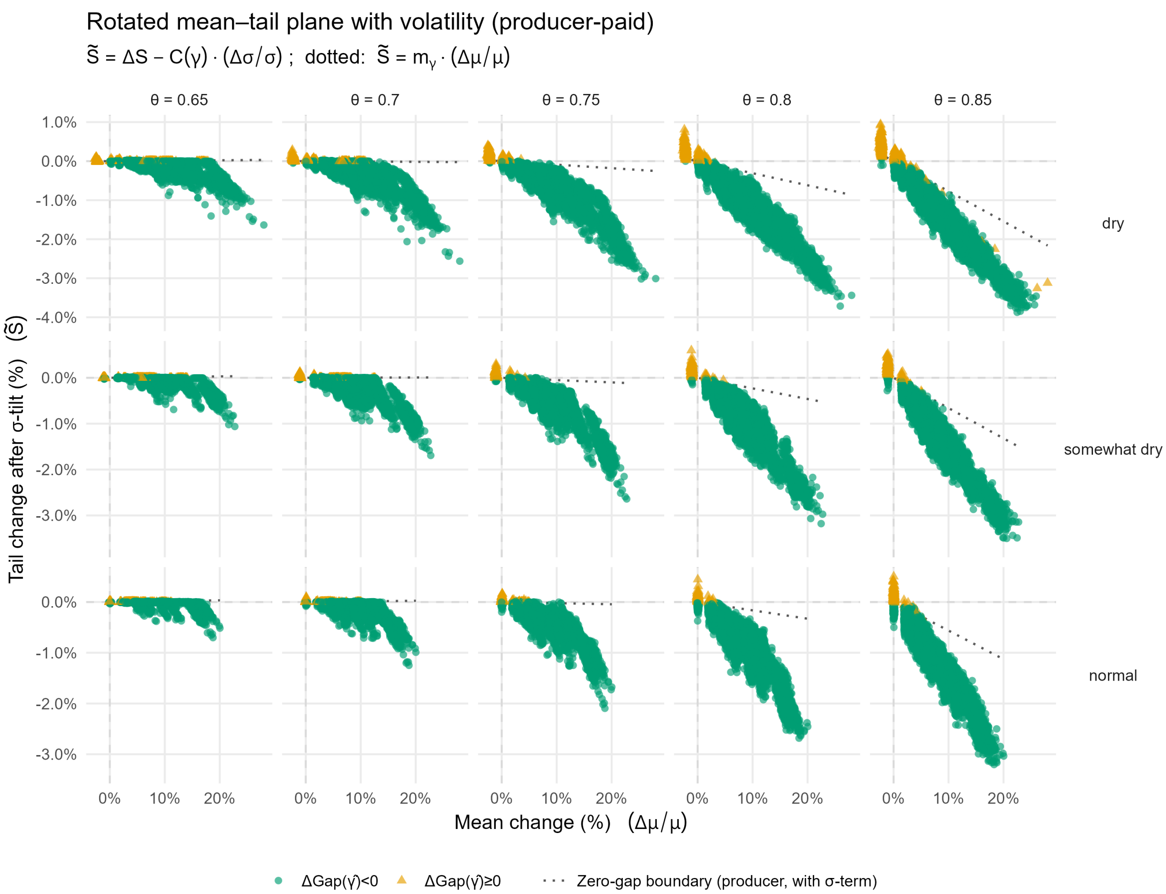

Figure 15: Rotated mean–tail plane with volatility (producer–paid). Vertical axis uses $\Delta\tilde{S} = \Delta S - C(\gamma)\Delta\sigma/\sigma$; dotted line is $\Delta\tilde{S} = m_\gamma \cdot (\Delta\mu/\mu)$. Green points have $\Delta\text{Gap}^{\text{prod}}(\gamma) < 0$.

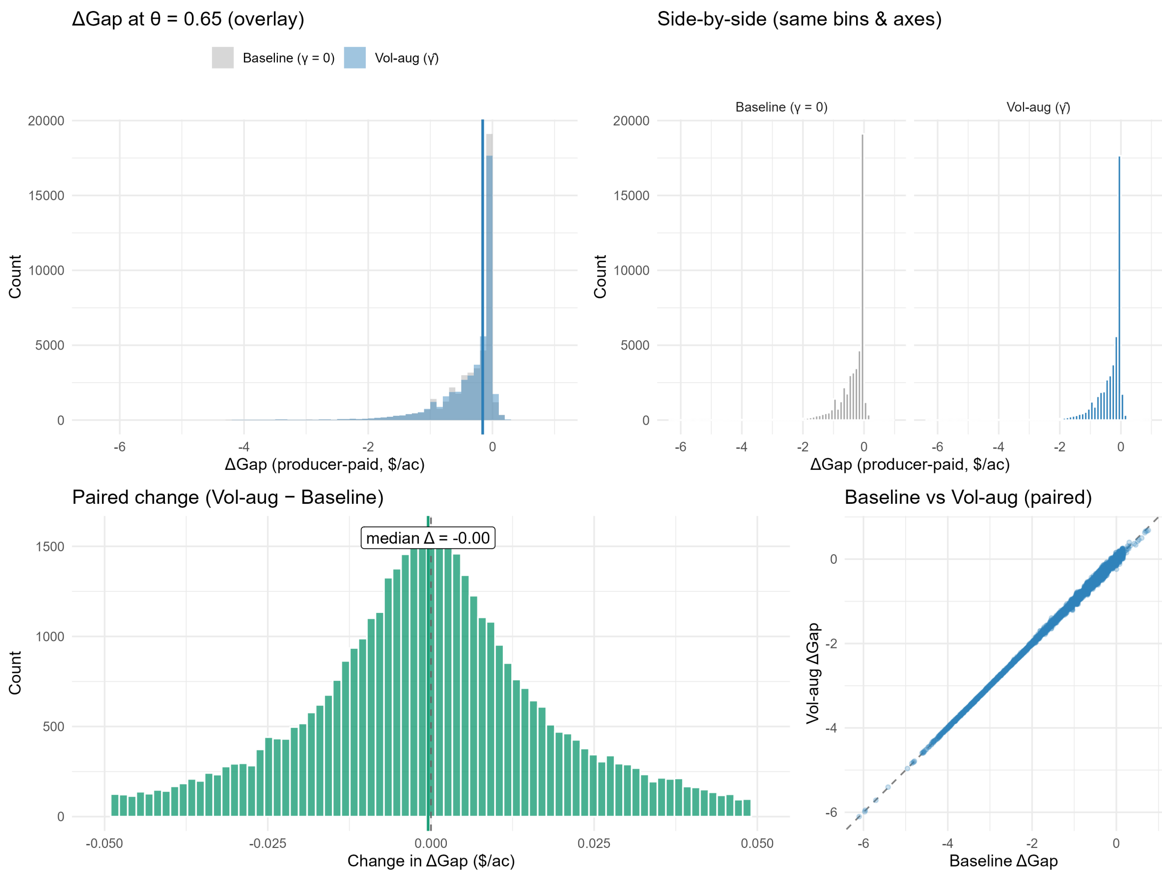

Figure 16: $\Delta\text{Gap}^{\text{prod}}$ before (gray) and after (blue) volatility augmentation at $\theta = 0.65$. The mass near zero indicates that a small positive γ largely neutralizes mispricing at low coverage.

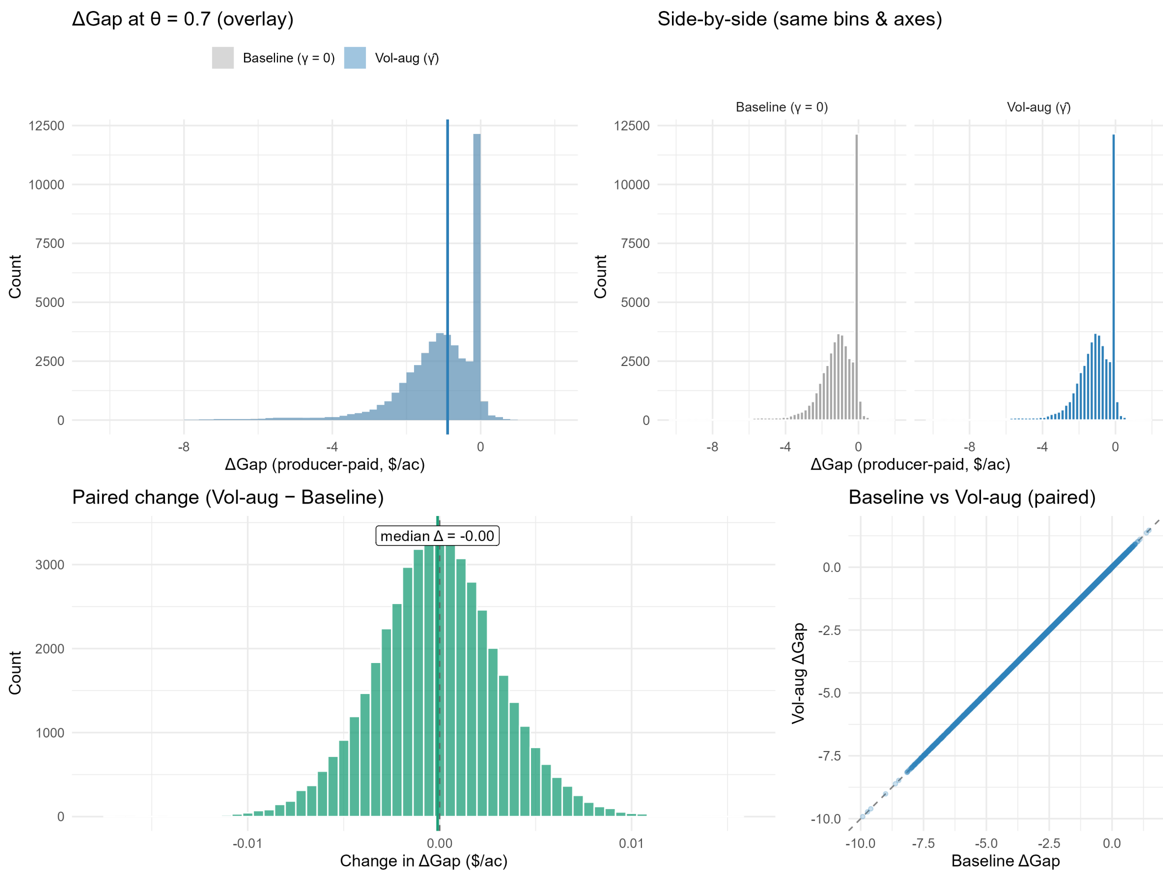

Figure 17: $\theta = 0.70$. A prominent spike around zero and a left tail remain; the adjustment reduces dispersion but does not eliminate negative gaps for all pairs.

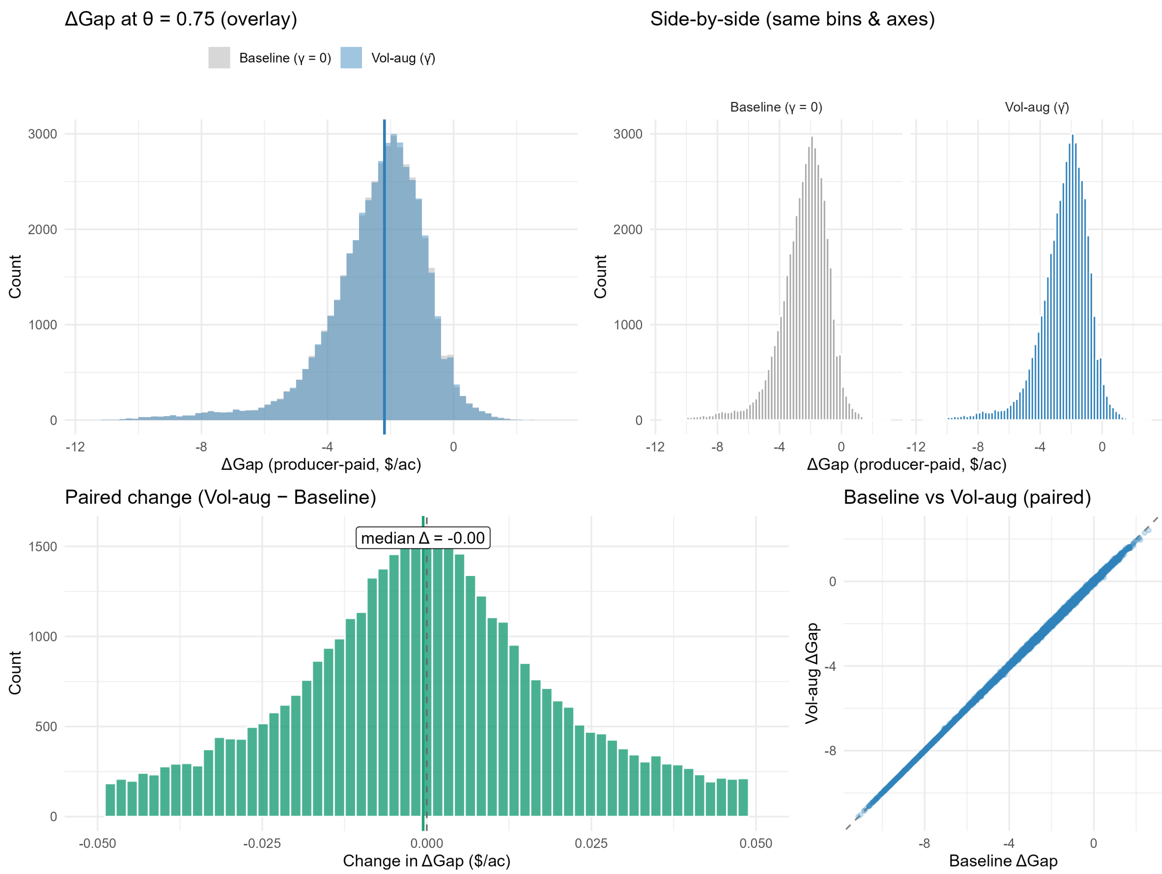

Figure 18: $\theta = 0.75$. The distribution shifts left by $\sim$ \$1–\$2 on average in our sample, consistent with the rotated plane in Fig. 15; tail improvements dominate.

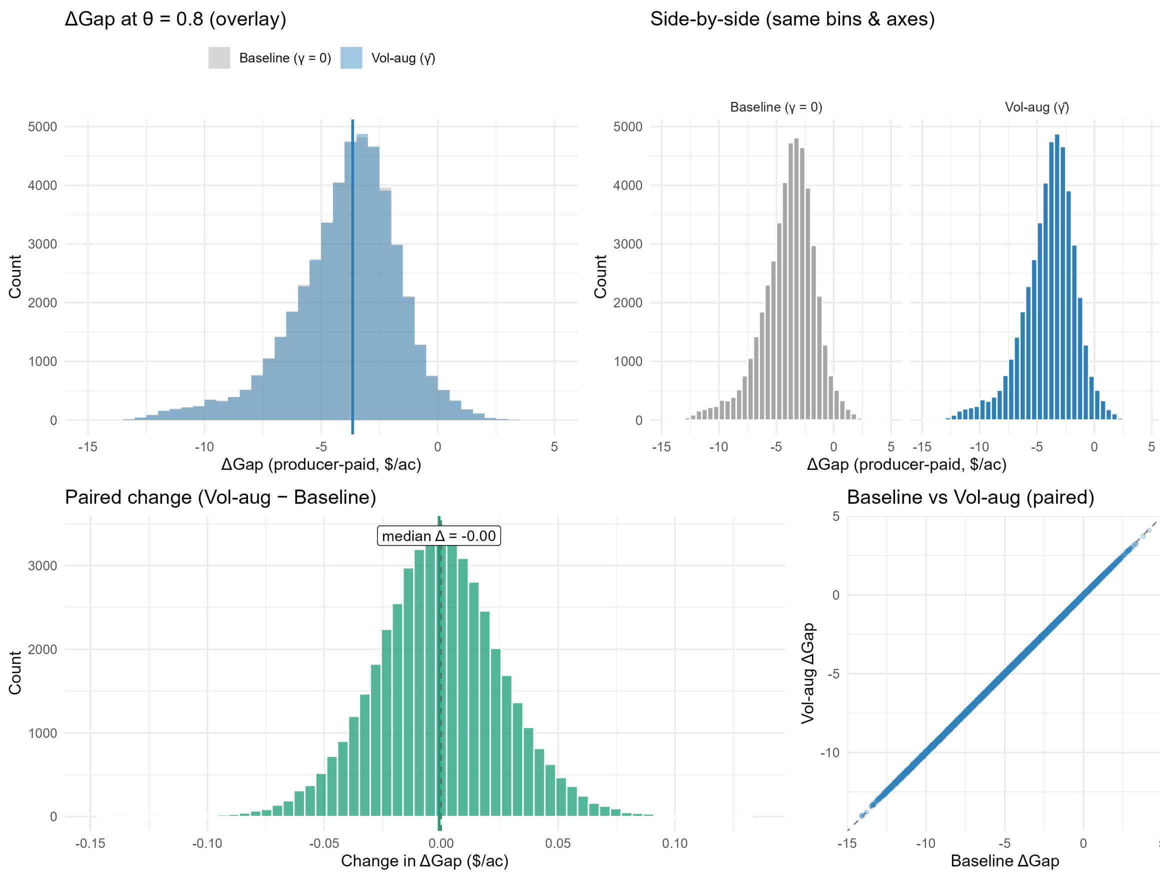

Figure 19: $\theta = 0.80$. Mass concentrates between \$-6 and \$-2 even after the volatility tilt; high coverage remains sensitive to tail risk.

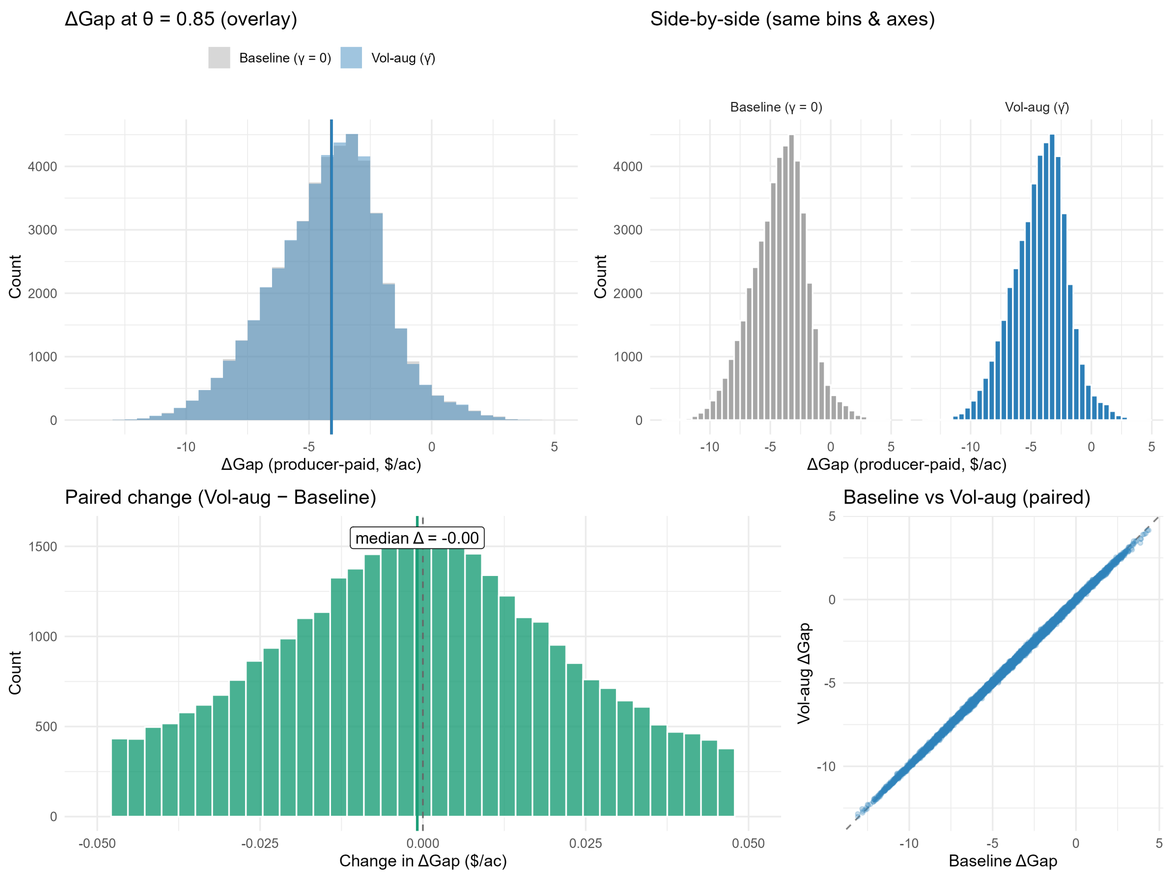

Figure 20: $\theta = 0.85$. Most pairs remain negative, reflecting strong tail effects at high coverage; the σ-tilt reduces dispersion but cannot fully offset tail improvements.

Table 5: Before/after summary of $\Delta\text{Gap}^{\text{prod}}$ by (θ, VPD).

θ	VPD	N	Median [p25, p75] $\Delta\text{Gap}^{\text{prod}}$ (\$/ac), baseline (γ=0)	Median [p25, p75] $\Delta\text{Gap}^{\text{prod}}$ (\$/ac), vol–aug ($\widehat{\gamma}$)	$\Pr(\Delta\text{Gap}^{\text{prod}} < 0)$ base (%)	$\Pr(\Delta\text{Gap}^{\text{prod}} < 0)$ vol (%)	$\Pr(\lvert\Delta\text{Gap}^{\text{prod}}\rvert \le \varepsilon)$ base (%)	$\Pr(\lvert\Delta\text{Gap}^{\text{prod}}\rvert \le \varepsilon)$ vol (%)
0.65	dry	15505	-0.23 [-0.82, -0.10]	-0.22 [-0.82, -0.10]	96.9	96.4	24.0	24.5
0.65	somewhat dry	15505	-0.16 [-0.22, -0.11]	-0.16 [-0.22, -0.11]	97.6	97.8	20.4	21.2
0.65	normal	15505	-0.14 [-0.18, -0.10]	-0.14 [-0.19, -0.10]	98.9	98.8	27.4	26.2
0.70	dry	15505	-1.28 [-2.05, -0.62]	-1.27 [-2.05, -0.62]	96.0	95.8	6.0	6.4
0.70	somewhat dry	15505	-1.01 [-1.52, -0.22]	-1.01 [-1.52, -0.22]	97.5	97.7	14.0	14.4
0.70	normal	15505	-0.22 [-1.07, -0.10]	-0.22 [-1.07, -0.10]	98.8	98.8	24.9	24.0
0.75	dry	15505	-2.52 [-3.75, -1.44]	-2.51 [-3.75, -1.44]	96.1	96.0	1.3	1.3
0.75	somewhat dry	15505	-2.19 [-3.03, -1.44]	-2.19 [-3.03, -1.43]	97.7	97.7	1.0	1.0
0.75	normal	15505	-1.74 [-2.57, -1.02]	-1.74 [-2.57, -1.02]	98.6	98.5	2.3	2.3
0.80	dry	15505	-4.16 [-5.94, -2.47]	-4.17 [-5.94, -2.48]	96.0	96.0	0.7	0.7
0.80	somewhat dry	15505	-3.64 [-4.91, -2.55]	-3.64 [-4.91, -2.55]	97.6	97.6	0.4	0.4
0.80	normal	15505	-3.35 [-4.59, -2.27]	-3.35 [-4.58, -2.28]	98.6	98.6	0.4	0.4
0.85	dry	15505	-3.87 [-5.79, -2.29]	-3.86 [-5.79, -2.29]	95.8	95.8	0.8	0.7
0.85	somewhat dry	15505	-4.06 [-5.43, -2.75]	-4.06 [-5.42, -2.76]	97.5	97.5	0.4	0.3
0.85	normal	15505	-3.96 [-5.65, -2.79]	-3.95 [-5.64, -2.80]	98.6	98.6	0.2	0.2

Notes: $\varepsilon = 0.10$ (\$/ac). Medians and quartiles are for $\Delta\text{Gap}^{\text{prod}}$; shares are percentages of pairs in each (θ, VPD) cell.