

参赛学生姓名： 吴博洋

中学： 北京师范大学附属实验中学

省份： 北京市

国家/地区： 中华人民共和国

指导老师姓名： 兰茹佳 郁露

指导老师单位： 北京师范大学经济与工商管理学院

北京师范大学附属实验中学

论文题目： 保障失灵，粮食贱卖：极端气候下小农户双重困局的
宏观经济学解释

Weak Protection, Cheap Grain

A Macroeconomic Explanation of the Dual Predicament of Chinese Smallholder Farmers under Extreme Climate Conditions

Boyang Wu¹

The Experimental High School Attached to Beijing Normal University

Abstract: This paper studies why Chinese smallholder farmers may face both payouts that do not match actual losses and low-price sales of wet grain after a climate shock, and how these household choices affect a regional grain price. I develop a two-stage model to answer these questions. Households choose coverage before the disaster and choose the distress sale after they see the payout. Aggregating household sales in a local market determines a common price. The theoretical analysis and numerical calibration produce three results. First, a higher probability of a missed payout widens the post-disaster cash gap, reduces optimal coverage, and increases immediate grain sales. Second, when the climate distribution is uncertain, households put more weight on more pessimistic climate models, so the marginal effect of premium subsidies declines. Third, synchronized distress sales lower the common price sellers face and may reduce next-season inputs. These results imply that simply raising premium subsidies cannot fully address the problem. More direct measures include improving loss assessment and payout credibility and providing drying, storage, and short-term liquidity during the disaster window. This paper identifies a macroeconomic transmission channel through which household cash constraints can transmit a local weather shock into regional price pressure, with implications for agricultural protection and food security.

Keywords: macroeconomic consequences; climate extremes; distress selling; regional grain prices; food security

JEL codes: E21; E44; Q18; Q54

¹ Paper submitted to the S.-T. Yau High School Science Award (Economics and Finance). Affiliation: The Experimental High School Attached to Beijing Normal University. Email: wuboyang0112@outlook.com. I thank Professor Rujia Lan of Beijing Normal University and Ms. Lu Yu of The Experimental High School Attached to Beijing Normal University for guidance on the model setup and the proofs. All remaining errors are mine.

Contents

1	Introduction.....	4
2	Background.....	6
2.1	Policy-based agricultural insurance and payout mismatch.....	6
2.2	Quality loss, farm types, and an aging workforce	7
2.3	Post-disaster distress sales and quality-specific prices	9
3	Theory and comparative statics	11
3.1	Framework and the ideal-insurance benchmark	11
3.2	Payout mismatch and state-contingent wealth.....	12
3.3	Two-stage choice, the cash constraint, and endogenous distress sales.....	13
3.4	Prudence, smooth ambiguity, and climate severity	15
3.5	Optimal coverage and the main comparative statics.....	17
4	Numerical results, regional aggregation, and policy implications.....	21
4.1	Parameters, the solution method, and coverage demand	21
4.2	Regional market equilibrium, climate severity, and the older-household share	24
4.3	Macroeconomic consequences and policy implications	26
5	Conclusion	29
	Appendix A Mathematical derivations	31
A.1	Notation, regularity, and optimality.....	31
A.2	The ideal-insurance benchmark	31
A.3	Missed payouts and coverage demand: Proposition 1	31
A.4	Existence and monotonicity of the distress-sale solution	32
A.5	Amplification of the distress-sale discount: Proposition 2	33
A.6	Belief distortion under smooth ambiguity: Proposition 3.....	33
A.7	Concavity of the objective and the local link: Proposition 4	33
A.8	Local comparative statics of climate severity: Proposition 5	34
A.9	Heterogeneous households, aggregation, and the regional price: Proposition 6	35
	Appendix B Numerical algorithm and parameters	37
	Appendix C Price anchors and sources.....	39
	References.....	40
	Acknowledgement	41

1 Introduction

Agriculture matters for both farm income and national food security. Extreme weather has become more frequent and severe (Dell, Jones and Olken, 2014; Lesk, Rowhani and Ramankutty, 2016). In 2025, spring drought, midsummer heat, and autumn rain affected the main producing regions in succession.²³⁴ National grain output remained resilient, but local households could still face quality loss and a cash gap at the same time. This paper starts from that autumn rain and asks a macroeconomic question: if expected insurance payouts fail to arrive when they are most needed, will many smallholder farmers sell grain at once? How would the regional price and next-season inputs respond?

During the 2025 National Day holiday, I went back to our hometown in Hebei with my family. The North China Plain was in the midst of a prolonged autumn rain.⁵ Two facts stood out during the visit, and both affected the same households. Most premiums are fiscally subsidized, yet the households I spoke with still expressed distrust. Output and quality had deteriorated, and wet grain had to be sold at a steep discount. I call the first a protection failure and the second a distress sale. When both affect the same household, they form the dual predicament in the title. Together they alter the household cash position and the regional clearing price of grain.

Consider the protection failure first. In 2025, the central budget paid 51.7 billion yuan in agricultural insurance premium subsidies and supported total premiums of more than 155 billion yuan.⁶ Households pay about one-fifth of the premium. High subsidies did not automatically create confidence in payout performance. Yi et al. (2022) show that a lump-sum premium-subsidy rule lets insurers split funds between administration and pure premium, so payouts can drift away from actual loss. Yi et al. (2025) report the 2013-2023 industry loss ratio at roughly 60 to 70 percent, below the 80 percent target. Tuo and Li (2003) have already

² Wei Fenghua, deputy director of the Rural Department of the National Bureau of Statistics, on 2025 summer-grain output: Henan, Shaanxi and other main producing areas had serious drought. Henan summer grain fell by 355,000 tonnes and Shaanxi by 145,000 tonnes. See the NBS note of 10 July 2025.

³ Office of the National Commission for Disaster Reduction and the Ministry of Emergency Management, national natural disasters in the first three quarters of 2025: national mean high-temperature days were 15.9, the second highest for that period since 1961. July mean temperature was the highest for that month since 1961.

⁴ Wei Fenghua, director of the Rural Department of the National Bureau of Statistics, on 2025 grain output: harvest-time rain in the Huang-Huai-Hai region left some crops moldy. See the NBS note of 12 December 2025.

⁵ Monitoring by the National Climate Centre shows that this autumn rain was the heaviest for the period since 1961. Much of the Huang-Huai-Hai land was too wet for machines to enter. Corn that should have been harvested around National Day was delayed. Corn still on the stalk began to mold, and some kernels had already sprouted. Disaster and price details are compiled from public reports in October 2025 in Caixin, China Newsweek, the People's Daily app, and Jiemian News. Hebei and Henan were both inside this rain. The ten-day quotes used later as a check are concentrated in the Henan producing area, where the public quality-specific series is more complete. Quality-specific price anchors are in Figure 1 and Appendix C. The main text does not claim an econometric identification. The numbers are qualitative checks.

⁶ The 51.7 billion yuan, the premium book above 155 billion yuan, 125 million household-times, and risk cover above 5 trillion yuan are from the Ministry of Finance, Report on the Implementation of China's Fiscal Policy in 2025. The 125 million household-times and risk cover above 5.2 trillion yuan also appear in remarks by the Finance Department of the Ministry of Finance at a State Council Information Office briefing.

identified a divergence in objectives: the government wants broad coverage, while household demand remains limited. Turn next to distress sales. In a normal year, corn spot prices are typically around 1.1 to 1.4 yuan per jin. In some hard-hit areas in autumn 2025, standard buyers refused wet grain, so the grain moved into the feed or industrial market, and the purchase price fell to about 0.6 yuan per jin. In early October around Zhoukou, grain with 30 to 35 percent moisture (30 to 35 "water points" in the trade) sold for about 1,260 yuan per tonne, or about 0.63 yuan per jin, while good dried grain nearby was still near 1.07 yuan per jin.

In the textbook, if supply only shifts left from its pre-disaster position, the equilibrium price should rise. The field setting also includes a lower-valued secondary market. After standard buyers leave, wet grain can only enter the feed or industrial market. Sales volume is also constrained by the cash gap, so price adjustment occurs before the physical supply response is complete. The existing sales volume is then absorbed by a weaker demand segment, and prices fall. The mechanism is straightforward. If insurance payouts fail when they are most needed, a missed payout opens a cash gap. The gap forces a distress sale, which depresses the price. The lower price makes protection appear less valuable, leading households to buy less coverage in the next round and to enter the next disaster with weaker protection.

Existing work has already shown that basis risk can reduce insurance demand (Clarke, 2016; Bryan, 2019; Elabed and Carter, 2015; Cole et al., 2013), and that a binding cash need can push asset prices below fundamental value (Shleifer and Vishny, 1992, 2011; Lorenzoni, 2008). Related literature asks how a borrower's balance sheet can amplify and propagate a shock: net worth falls, borrowing capacity falls, and spending and prices fall again (Kiyotaki and Moore, 1997; Bernanke et al., 1999). Grain is not collateral. This paper follows the amplification structure. When the probability of climate risk itself is no longer stable, Ellsberg (1961) and Klibanoff et al. (2005) give a way to separate risk aversion from ambiguity aversion. Zhao et al. (2025) also show that payouts computed from meteorological data at different spatial scales can differ by a wide margin. Most of these papers focus on one link. What is missing is a chain that runs from payout mismatch to grain sales and then to the regional price. This paper connects these strands by combining analytical derivations with numerical calibration.

The paper does three things. First, it places insurance and determines the coverage before the disaster and sales after it, so the payout outcome affects distress sales through the cash gap. Second, it embeds the distress sale in a downward-sloping spot price, so household choices can aggregate into common price pressure. Third, it keeps risk aversion and ambiguity aversion in the same model, then computes regional coverage, regional distress sales, and the common clearing price when many households choose at once, and it also introduces climate severity and the share of older households at the aggregate layer. When the miss probability rises from

0.10 to 0.50, optimal coverage falls from 0.966 to 0.356, and expected distress sales rise from 0.070 to 0.408. Subsidies still help, but the same subsidy buys less coverage when the miss rate is high. When climate risk is more severe and the share of older households is higher, aggregate coverage falls, regional distress sales rise, the clearing price falls, and the pressure can pass through to next-season inputs.

The object of analysis is staple grain, with corn as the leading case. The analysis considers quality loss and focuses on smallholder farmers with limited cash buffers. The visit motivates the problem, and the model explains the proposed mechanism. The paper is a theory-and-calibration exercise. It does not claim a causal estimate. The later price points are quality-specific plausibility checks rather than estimates. Section 2 gives the institutional background and the facts. Section 3 builds the model and states six propositions. Section 4 reports the numbers, the regional aggregation, and the policy discussion. Section 5 concludes. Proofs are in Appendix A. The algorithm and the baseline parameters are in Appendix B. The price anchors used in Figure 1 are in Appendix C.

2 Background

2.1 Policy-based agricultural insurance and payout mismatch

China is a large agricultural country dominated by small farmers. The Third National Agricultural Census shows that large-scale farm operators are only 2.0 percent of households; the other 98.0 percent are smallholder farmers. The mean cultivated area per household is 7.8 mu.⁷ The demographic structure matters alongside farm size. At the census, agricultural workers aged 55 and above already made up 33.6 percent.⁸ Later this paper treats this share as the older-household share ω . Smallholder constraints extend beyond land. Many also lack short-term storage and the cash needed to wait. A model that does not adequately capture smallholder farmers' cash constraints, storage limits, and demographic structure is unlikely to reveal the specific problems of China's agricultural insurance. The model must begin with these households' actual cash inflows and outflows.

Agricultural insurance has grown quickly since the 2007 central premium-subsidy pilots. Premium income rose from 5.333 billion yuan in 2007 to 148.373 billion yuan in 2024, a compound annual rate of about 21.61 percent.⁹ In 2025 total agricultural-insurance premiums exceeded 155 billion yuan. The central budget paid 51.7 billion yuan in premium subsidies and

⁷ Han Jun used the Third National Agricultural Census in this way at a State Council Information Office briefing on 1 March 2019 when he summarized "a large country of small farms." See the Xinhua report.

⁸ See Leading Group Office of the Third National Agricultural Census and the National Bureau of Statistics, Main Data Bulletin of the Third National Agricultural Census (No. 5). The definition is agricultural workers, not the whole rural population. The later 33.6 percent share of older households is only a reference. The 2016 cross-section is not treated as a 2025 estimate.

⁹ The 2007 figure of 5.333 billion yuan is the usual industry series. The 2024 figure of 148.373 billion yuan matches the NFRA agricultural-insurance premium total (148,373.02 million yuan). 2025 premiums, subsidies, and cover are in the previous note.

covered more than 5.2 trillion yuan of risk for 125 million household-times. This expansion has been driven mainly by fiscal subsidies. From 2024, full-cost insurance and planting-income insurance for the three staple grains (rice, wheat and corn) have been available nationwide. The central budget pays 45 percent of the premium in the central and western regions and the northeast, provincial budgets pay at least 25 percent, and, with city and county funds, the public share is close to 80 percent. Households pay about one-fifth. Full-cost insurance can cover material input costs, land costs, and labor costs. The sum insured is in principle no higher than 80 percent of the crop's output value. With the subsidy, households pay a small share of the premium, so nominal enrollment is broad.

On loss assessment, the three staples use field inspection and plot sampling, not weather-index insurance. Smallholder farmers are many and scattered, so group claims often rest on a village-average loss, and assessment often happens before harvest. Sampling, triggers, and a blind spot on quality can all pull the assessed loss away from the true loss. Some households who should be paid are not paid; some who should not be paid are paid.¹⁰ In the visit, households summarized the mismatch plainly: a large disaster gets a small payout, a small disaster gets none, and a year with no disaster can still get a check. On paper, full-cost insurance can pay about one thousand yuan per mu for a total loss at maturity. Quality loss is often averaged first and then further reduced by grading disputes, so the cash a household actually receives can fall far below the season's sunk material and labor costs. With fiscal subsidy, the effective source of demand has largely shifted from households to the government. Tuo and Li (2003) observed that under a fully voluntary rule, the smallholder farmers who farm most of the sown area have "very little demand; subsidy or no subsidy, they are not interested." High subsidies raise nominal coverage. They do not automatically raise the credibility of the payout. Basis risk in this paper is the broad gap between the payout and true loss.

2.2 Quality loss, farm types, and an aging workforce

The 2025 autumn rain in the Huang-Huai-Hai region damaged crop quality in a way that differed from many earlier disasters. Prolonged rain caused corn to sprout and mold on the stalk before harvest. Standard satellite-based biomass measures or NDVI showed no anomaly at the time. The plants still looked whole, and some fields were still green. In downstream markets, moldy corn had already lost its place in the food chain, so its value as standard commercial grain fell sharply. Two kinds of loss need to be kept apart. Physical quantity need not fall much. Economic value can still fall a lot. When station data or a satellite index still sits in the "normal" range, the grain in the household's hands may already be badly damaged. That is the physical source of "hit by the disaster, not paid." In its note on 2025 grain output, the

¹⁰ A missed payout that should have been paid is often called a Type I error. A payout that should not have been paid is a Type II error. The visits kept returning to the first kind. This paper focuses on the first kind.

National Bureau of Statistics wrote a separate sentence: harvest-time rain in the Huang-Huai-Hai region left some crops moldy; after deducting the mold loss, national autumn grain was 536.62 million tonnes, up 1.5 percent from the year before.¹¹ National output rose, yet quality loss was singled out, meaning the official record now treats it as a separate fact. Stable aggregate output is not the same as the absence of household loss. Zhao et al. (2025) compare indemnities computed from meteorological data at different spatial scales and find large gaps, so a station index can miss what is happening in the field.

The same rainfall affects different farm types differently. Using National Bureau of Statistics figures cited by the Ministry of Agriculture and Rural Affairs, rural per capita disposable income in 2023 was 21,691 yuan, with wages at 42.2 percent and net operating income at 34.3 percent.¹² Since wages first passed operating income in 2015, part-time work has been the rural norm. The national average is annual household income. This paper focuses on cash on hand during the harvest-disaster window. Wages, transfers, and migrant-work earnings can provide a cash buffer and lift the worst-state floor, so insurance demand is less likely to collapse. The mechanism is most relevant for smallholder farmers whose buffer is still thin and who face quality loss and cash outlays at the same time.

Media reports make the equipment gap clear. Scaled operators with crawler harvesters and drying towers can harvest during brief dry intervals and use mechanical drying to limit mold. Smallholder farmers who rely on hired labor and have no dryer see wet grain start to ferment after two days in a pile, and then they have to sell at a steep discount. Table 1 splits this across four dimensions and maps each to a later model variable.

Table 1 How different farms behave in a severe autumn rain

Feature	Scaled operators	Small traditional farms	In the model
Capital and equipment	Crawler harvesters, drying towers, standard storage	Hired machines, open-air drying	Capital stock sets the shock boundary
Harvest and storage	Rush harvest and high-heat drying	Machines cannot enter; wet grain ferments again	Quality decay is uneven across farms
Market access	Standard market or reserve channels	Refused by the standard market; moved to the secondary market	The product is segmented
Cash flow	Reserves; grain can be held	Debt and living costs are rigid; forced to sell	The cash constraint triggers a minimum distress sale

Age also changes the cash constraint within the same group of smallholder farmers. After younger workers leave, the people still in the field are older. Late-life farming, outward wage

¹¹ See Wei Fenghua, director of the Rural Department of the National Bureau of Statistics, on grain output, 12 December 2025. Full-year grain output was 1,429.75 billion jin, up 16.75 billion jin or 1.2 percent from the year before.

¹² See the Ministry of Agriculture and Rural Affairs note on rural household income in 2023. The original series is the National Bureau of Statistics household survey.

flows, and transfers from children overlap. A half-work, half-farm pattern puts the younger labor outside earning wages and leaves the land to older people at home. The census figure of 33.6 percent is about who is still working the land, not who still lives in the village.¹³ A large wage flow leaves the village, so the post-disaster cash buffer is thinner. Money sent by children can help, but it often arrives late and may not fill the gap in that harvest week. In the later model, older households have higher risk aversion and less disaster cash. Proposition 6 aggregates those two differences into the regional price.

2.3 Post-disaster distress sales and quality-specific prices

Quality standards alter the set of eligible buyers. After reserve depots and large processors leave, sales concentrate in the secondary market, and the clearing price deviates from its usual range.¹⁴ Figure 1 plots quality-specific corn purchase prices in autumn 2025. In early August, national grade-two spot corn was about 2,403 yuan per tonne, or 1.202 yuan per jin. At the end of September, Jiaozuo deep-processing was about 2,290 yuan per tonne, or 1.145 yuan per jin. In early October, Henan deep-processing was about 2,155 yuan per tonne, or 1.078 yuan per jin, while Zhoukou wet grain with 30 to 35 percent moisture sold for about 1,260 yuan per tonne, or 0.630 yuan per jin. The gap opened clearly in October.




Figure 1 Quality-specific corn purchase prices, autumn 2025

Note: Only points with a direct source are kept. The line only joins those points. Missing ten-day slots are left blank. Sources are in Appendix C.

¹³ The questionnaire survey for the “National Rural Operational Patterns” project in 2024 found that 35.8% of the people surveyed were aged 60 or above, and only 7.53 percent were under 40. The main text relies solely on the more authoritative census figure of 33.6 percent.

¹⁴ The finance literature calls a short window of forced sales that push prices down together a fire sale (Shleifer and Vishny, 2011). The term comes from Shleifer and Vishny (1992): when an asset's natural buyers are short of cash, the liquidation price can fall below the asset's value in use. Grain spoils, and quality standards also rewrite the set of buyers, so this paper does not treat grain as the same object as a financial asset. It borrows one structure: a rigid cash need sends sellers into a short-term market together, and total sales push down the common price.

Once quality crosses the reserve or food-use line, price adjustment may be discontinuous. Figure 2 draws this cliff. To the left of the dashed line, price eases down as quality worsens. Past the line, standard buyers leave and the grain shifts into the feed or industrial market. Figure 2 is a sketch, not a fitted curve. It only gives an empirical direction for the later distress-sale discount.




Figure 2 The price drop after quality crosses the line (schematic)

Note: The figure shows the split between the standard market and the secondary market.

Together, these facts imply the next step. When insurance fails to pay because loss assessment does not match the actual loss, a smallholder farmer with few other options has to sell the wet grain for cash before quality worsens. That forced sale lowers cashable income in the disaster state, makes a missed payout more costly, and makes protection less worth buying. Two facts that appeared separate now form a single mechanism. Figure 3 draws this feedback as the bridge into the model.




Figure 3 Feedback from household choice to the regional price

Note: Climate extremes worsen payout mismatch. After protection demand falls, the disaster cash gap and concentrated distress sales both rise. The regional grain price and worst-state wealth fall, which weakens the appeal of protection.

3 Theory and comparative statics

3.1 Framework and the ideal-insurance benchmark

This section builds the model from the simple case to the fuller one. It first restates, in the cleanest setting, why expected utility predicts high coverage. It then adds payout mismatch, post-disaster distress sales, and climate severity. Later, household distress sales are added up to a common price, which is the regional grain price.

On ambiguity, the max-min expected-utility rule of Gilboa and Schmeidler (1989) takes the worst case in a set of priors, but the objective is often not differentiable at the boundary and it locks ambiguity aversion at an extreme. This paper uses smooth ambiguity utility instead, because it directly captures the idea that the decision maker no longer knows which probability model is in use. Its inner layer measures risk aversion, while its outer layer measures ambiguity aversion. The outer layer is smooth and differentiable, and when it is linear, the model reduces to expected utility. Table 2 compares the three setups.

Table 2 Three utility setups for uncertainty

Dimension	Expected utility	Max-min expected utility	Smooth ambiguity
Object	Risk with known probabilities	A set of priors; take the worst	The probability distribution itself is unknown
Fit to climate	Weak; the distribution is treated as known	Medium, but too pessimistic	Strong; the distribution can drift
Risk and ambiguity separable	No ambiguity dimension	Hard to separate	Inner and outer layers are separate
Tractability	Smooth	Often not differentiable at the boundary	Smooth and differentiable
Contains expected utility	Itself	No	Falls back when the outer layer is linear
Strength of ambiguity aversion	Not adjustable	Fixed at an extreme	Adjustable through the outer function

Frictionless expected utility remains a useful benchmark. Under ideal insurance, its prediction is clear: risk aversion plus a subsidized premium should produce high coverage. The low trust and limited perceived benefit in the visits do not match that prediction. The mismatch itself points the way. Households can be rational and risk averse. The classic prediction can fail if real insurance fails to pay when it is most needed, if the household has to meet cash bills before the indemnity arrives, and if the climate distribution can no longer be read off a stable historical frequency. The next three layers match those three facts.

The ambiguity here has a clear experimental source. Ellsberg (1961) compared two urns. In one urn, the red and black balls are known to be half and half. In the other urn, only the total of one hundred is known, not the mix. Most people would rather bet on the first urn, even though the objective chance of winning can be the same. People dislike more than the ups and downs of winning and losing. They also dislike not knowing the odds. Climate change pushes households into the second urn. Past rainfall frequencies no longer tell them, in a stable way, what next year will look like. What follows is how that dislike stacks with missed payouts and the cash constraint, pulling insurance demand down layer by layer.

The paper uses constant relative risk aversion. For $\gamma \neq 1$,

$$u(W) = \frac{W^{1-\gamma}}{1-\gamma}, u'(W) = W^{-\gamma}, R(W) = \gamma \quad (1)$$

So relative risk aversion is constant at γ . A larger γ means the household is less willing to let wealth swing widely across states. Absolute and relative risk aversion are defined in Appendix A.1.

Consider a risk-averse representative household with initial wealth W_0 and a sunk fixed production cost C . At harvest the state of nature is $s \in \{G, B\}$: G is normal, with baseline output value Y_G ; B is the disaster, and baseline output value falls to $V_B < Y_G$. The true prior probability of disaster is π . First imagine an ideal insurance contract. It pays the loss $L = Y_G - V_B$ if and only if the disaster occurs. The premium per unit of cover is P , and the fiscal subsidy rate is $\sigma \in (0,1)$. The household chooses a coverage share $\alpha \in [0,1]$ and pays $(1 - \sigma)P\alpha$ out of pocket. Expected utility is

$$EU(\alpha) = (1 - \pi)u(W_0 - C + Y_G - (1 - \sigma)P\alpha) + \pi u(W_0 - C + V_B + \alpha L - (1 - \sigma)P\alpha) \quad (2)$$

So, if the subsidized premium is below the fair actuarial value, that is $(1 - \sigma)P < \pi L$, optimal coverage under ideal insurance is $\alpha^* = 1$, and wealth is fully smoothed across the two states.

The derivative of expected utility has the sign of a weighted gap in marginal utilities across the two states. When insurance is a good deal for the household, the derivative is positive on the whole of $[0,1]$, so the optimum sits at the right endpoint. The full steps are in Appendix A.2. This benchmark pushes the tension onto the assumptions. If rationality, risk aversion, and the subsidy all hold, the problem must lie in the ideal-insurance assumption.

3.2 Payout mismatch and state-contingent wealth

In real agricultural insurance, a payout must first pass judgment on true loss. That judgment is not easy. On cash crops, a weather or remote-sensing index can miss the field loss. For the

three staples, field inspection and plot sampling are limited by triggers, averaging, and a blind spot on quality, so they also struggle to match each household's true loss. A signal $\theta \in \{0,1\}$ writes this gate in one symbol; $\theta = 1$ means a payout is granted. The mismatch between the signal and the true state is basis risk in the broad sense. Write

$$p_{B0} = \text{Prob}(\theta = 0 \mid s = B), p_{G1} = \text{Prob}(\theta = 1 \mid s = G) \quad (3)$$

p_{B0} is the friction that runs through the whole paper: the miss probability. The harder it is for current loss assessment to see a quality loss, the higher this probability. p_{G1} is the probability of a payout when there is no loss. End-of-period wealth then has four states: harvest and no payout W_{G0} , harvest and a payout W_{G1} , disaster and a payout W_{B1} , and disaster and no payout W_{B0} . Before distress sales are added, they can be written as

$$W_{G0}(\alpha) = W_0 - C + Y_G - (1 - \sigma)P\alpha \quad (4)$$

$$W_{G1}(\alpha) = W_{G0}(\alpha) + \alpha L \quad (5)$$

$$W_{B1}(\alpha) = W_0 - C + Y_B + \alpha L - (1 - \sigma)P\alpha \quad (6)$$

$$W_{B0}(\alpha) = W_0 - C + Y_B - (1 - \sigma)P\alpha \quad (7)$$

Two symbols need to be kept apart. Y_B is the baseline output value in the disaster state. In equations (6) and (7), it is first written as a given number. The next subsection drops that and lets distress sales and the spot price determine it. W_{B0} is the core case: the disaster has already happened, the premium has already been paid, and the payout has not arrived. It will determine the result later. The payout in the model is written as αL . If a "large disaster, small payout" is written as a discount $\delta \in (0,1]$ on αL , that discount enters the first-order condition in the same way as a higher miss rate or a smaller nominal loss L .¹⁵ Equations (4) to (7) therefore do not need a separate discount.

3.3 Two-stage choice, the cash constraint, and endogenous distress sales

The model has two decision dates. Before the disaster, the household chooses coverage α and pays the premium. After the disaster, the state s and the payout signal θ are already realized, and the household then chooses the immediate sale $q_{s\theta}$. Immediate sales cover debt, machine hire, and living costs. Grain that is not sold at once, after drying and storage, earns a delayed value R_s per unit.

¹⁵ The discount multiplies the indemnity and is equivalent to a smaller nominal loss. A "large disaster, small payout" is already inside the same parameter. A missed payout is the corner where the discount is zero.

In state s , cash on hand before the sale is A_s , the harvest quantity is y_s , and the near-term rigid outlay is D . A payout adds income αL . The cash gap that has to be filled by an immediate sale is

$$x_{s\theta}(\alpha) = D - A_s - \theta\alpha L + (1 - \sigma)P\alpha \quad (8)$$

A larger rigid outlay widens the gap. A cash buffer and a payout shrink it. The premium already paid widens it. The same ledger can be read as a post-disaster cash-flow statement. Inflows are cash before the sale and any payout that arrives. Outflows are the rigid outlay and the premium already paid. After netting the two sides, a positive gap means the household must sell some grain at the spot price. Insurance shows up twice on this statement: the pre-disaster premium is an outflow, and a payout is an inflow. A missed payout leaves only the outflow, so the gap is larger. If $x_{s\theta} \leq 0$, the household does not need a distress sale for current bills. The mass of households is normalized to one. The spot price in state s is locally linear:

$$p_s(q) = a_s - \kappa_s q, \kappa_s > 0 \quad (9)$$

a_s is the intercept for that quality grade when sales are not bunched. κ_s measures short-run absorptive capacity. Fewer buyers or tighter funds raise κ_s . The cash constraint requires spot sales revenue to cover the gap:

$$(a_s - \kappa_s q_{s\theta})q_{s\theta} \geq x_{s\theta}(\alpha) \quad (10)$$

In the parameter range used here, the unit value of delayed sales is at least the marginal return on a current distress sale, so the household sells only the smallest quantity that meets the cash constraint. When the constraint binds and the solution sits on the stable branch,

$$q_{s\theta}(\alpha) = \frac{a_s - \sqrt{a_s^2 - 4\kappa_s x_{s\theta}(\alpha)}}{2\kappa_s} \quad (11)$$

If the gap is not positive, then $q_{s\theta} = 0$. The distress-sale solution must also raise enough cash in the spot market and cannot exceed the harvest. Putting spot income and delayed income back into wealth gives

$$W_{s\theta}(\alpha) = A_s - (1 - \sigma)P\alpha + \theta\alpha L - D + p_s(q_{s\theta})q_{s\theta} + R_s(y_s - q_{s\theta}) \quad (12)$$

The disaster-state baseline output value is no longer an exogenous Y_B . The distress sale and the spot price jointly set it. Define the discount relative to "no distress sale, all grain sold at R_s " as

$$\Delta_{\text{fire}}(q_{s\theta}) = (R_s - p_s(q_{s\theta}))q_{s\theta} \quad (13)$$

A missed payout raises x_{B0} , which raises q_{B0} , lowers p_B , raises Δ_{fire} , and pushes worst-state wealth W_{B0} down further. The fire-sale discount here results from the cash gap, not an extra term pasted on. More capital or better working-capital terms show up as a higher A_s or a lower κ_s . Both cut the minimum distress sale and shrink the discount. R_s is the delayed value of grain that is not sold at once. Under the minimum-sale rule, raising R_s alone improves the return on retained grain and does not directly change the $q_{s\theta}$ pinned by the cash gap. Later I split the role of public drying into three channels: preserving value, widening absorptive capacity, and supplying working capital.

3.4 Prudence, smooth ambiguity, and climate severity

For $\gamma \neq 1$, marginal utility is $u'(W) = W^{-\gamma}$, the second derivative is $u''(W) = -\gamma W^{-\gamma-1} < 0$, and the third is $u'''(W) = \gamma(\gamma + 1)W^{-\gamma-2} > 0$. A positive third derivative is prudent: as wealth falls, marginal utility rises and becomes more sensitive to a further loss. Figure 4 draws marginal utility at $\gamma = 3$. When a distress sale pushes W_{B0} still lower, the curve enters its steepest stretch, and one extra unit of premium paid out carries a large utility penalty. That is the mathematical source of Proposition 2.




Figure 4 CRRA marginal utility

Note: When wealth is pushed toward the worst state, the curve steepens quickly, so a further small loss carries a large utility penalty.

If probability model k is known, including the disaster probability and the miss probability that belong to that model, then

$$EU_k(\alpha) = \pi_k p_{B0}^k u(W_{B0}) + \pi_k (1 - p_{B0}^k) u(W_{B1}) + (1 - \pi_k) (1 - p_{G1}^k) u(W_{G0}) + (1 - \pi_k) p_{G1}^k u(W_{G1}) \quad (14)$$

Under climate extremes, the household cannot tell which probability model governs the real world. Let the set of possible models be $\mathcal{K} = \{1, \dots, K\}$, and let the household place a second-order belief $\mu(k) \geq 0$ on them, with $\sum_k \mu(k) = 1$. Smooth ambiguity preferences are

$$V(\alpha) = \sum_{k=1}^K \mu(k) \phi(EU_k(\alpha)) \quad (15)$$

The outer function satisfies $\phi' > 0$. If $\phi'' < 0$, the household dislikes ambiguity. If ϕ is linear, the objective falls back to ordinary expected utility over models. Risk aversion (concavity of the inner u) and ambiguity aversion (concavity of the outer ϕ) are thus kept in two layers. The paper assumes the following: u and ϕ are twice continuously differentiable, $u' > 0$, $u'' < 0$, $u''' > 0$, $\phi' > 0$, $\phi'' \leq 0$, and wealth is positive on the feasible set. The numerical work uses the normalized exponential $\phi_\rho(x) = (1 - e^{-\rho x})/\rho$ ($\rho > 0$) and sets $\phi_0(x) = x$. Then $\phi'_\rho(x) = e^{-\rho x} > 0$ and $\phi''_\rho(x) = -\rho e^{-\rho x} \leq 0$, and the criterion returns to expected utility as $\rho \rightarrow 0$. For any $\rho > 0$, this normalization is only a positive affine transform of the original exponential objective, so it does not change optimal coverage.

Next, write extreme climate into the baseline disaster-state value. That value is already endogenous. It is set by the price intercept a_B , short-run absorptive capacity κ_B , delayed unit value R_B , and the harvest y_B . A climate-severity index $\xi \in [0,1]$ writes these parameters as functions of one index, and the baseline disaster value then has an explicit form:

$$Y_B(\xi) \equiv R_B(\xi)y_B(\xi), \frac{dY_B}{d\xi} < 0 \quad (16)$$

On a quality path, the harvest quantity can stay almost still, while market and payout conditions worsen with ξ :

$$\begin{aligned} a_B(\xi) &= a_B^0 - \lambda_a \xi, \kappa_B(\xi) = \kappa_B^0 + \lambda_\kappa \xi, R_B(\xi) \\ &= R_B^0 - \lambda_R \xi, p_{B0}(\xi) = p_{B0}^0 + \lambda_p \xi \end{aligned} \quad (17)$$

Where $\lambda_a, \lambda_\kappa, \lambda_R, \lambda_p > 0$. A quantity path mainly raises the disaster probability π and cuts y_B ; the price intercept, delayed value, and the miss rate can stay roughly unchanged. Both paths can lower Y_B , yet insurance behaves quite differently. Ordinary drought or flood cuts quantity. Measuring the crop is relatively easy, loss assessment is clearer, and missed payouts are less likely to rise systematically. Long rain can leave the plant green and the kernel already ruined. Standard buyers leave, a_B and R_B shift down together, and κ_B and p_{B0} rise together. Quantity barely moves, yet loss assessment struggles to match the field. That is why "hit by the disaster, not necessarily paid" sounded so sharp in the 2025 autumn rain, and why this channel is included in the model.

Numerically, I set $\lambda_a = 0.06$, $\lambda_b = 0.06$, $\lambda_k = 0.06$, $\lambda_R = 0.12$, $\lambda_p = 0.20$, and the baseline case is set by $a_B^0 = 1$, $\kappa_B^0 = 0.35$, $R_B^0 = 1.10$, $p_{B0}^0 = 0.30$. At $\xi = 0$ the model returns to the baseline calibration. At $\xi = 1$, Y_B falls from 1.100 to 0.980. The parameters keep the unpaid cash gap fundable even in the worst case $\xi = 1$. The model solves for optimal coverage via a grid search.

3.5 Optimal coverage and the main comparative statics

Write $\varpi_{k,s\theta}$ for the joint probability of state (s, θ) under model k , and write $m_{s\theta}(\alpha) \equiv \partial W_{s\theta} / \partial \alpha$. Under a single known model, the full first derivative is

$$\begin{aligned} \frac{\partial EU_k}{\partial \alpha} &= \sum_{(s,\theta) \in \{G,B\} \times \{0,1\}} [\varpi_{k,s\theta} u'(W_{s\theta}) m_{s\theta}(\alpha)], m_{s\theta}(\alpha) \\ &\equiv \frac{\partial W_{s\theta}}{\partial \alpha} \end{aligned} \quad (18)$$

Let $c = (1 - \sigma)P$. When the cash constraint is slack, $m_{s\theta} = \theta L - c$. When it binds, $p_s(q_{s\theta})q_{s\theta} = x_{s\theta}$, equation (12) becomes $W_{s\theta} = R_s(y_s - q_{s\theta})$, and therefore $m_{s\theta} = -R_s(\partial q_{s\theta} / \partial x_{s\theta})(c - \theta L)$. The missed-payout state's marginal effect is then always negative. The paid state's marginal effect is positive when $L > c$. The distress-sale channel widens the gap through $\partial q / \partial x$. Equation (18) therefore holds both the direct premium effect and the indirect effect of the cash constraint.

Proposition 1 (Missed payouts and coverage demand) Under expected utility, let $p_{B0}^k > 0$.

- (i) If the marginal-utility term in the missed-payout state is negative and has a common lower bound, and if the positive marginal contributions from the other states have a common upper bound, then there is a wealth threshold $\bar{W} > 0$. When $W_{B0}(\alpha) \leq \bar{W}$ holds for every $\alpha \in [0,1]$, optimal coverage is 0.
- (ii) If the missed-payout state is the worst state and $W_{B0}(\alpha) < W_{B0}(0)$ holds for every $\alpha > 0$, then under CRRA any sequence of optima satisfies $\alpha^*(\gamma) \rightarrow 0$ ($\gamma \rightarrow \infty$). Coverage demand therefore need not rise with risk aversion. The calibration in this paper sits in the range where coverage falls with γ .

Part (i) of Proposition 1 comes from $u'(W_{B0}) = W_{B0}^{-\gamma}$. When the worst state nears zero, the negative marginal effect in the missed-payout state can dominate the remaining finite terms. Part (ii) cannot be read from the first derivative at $\alpha = 0$ alone, because wealth in the paid and missed states then coincides. A strict proof uses the CRRA certainty equivalent. As γ goes to infinity, the certainty equivalent goes to minimum wealth. Any positive coverage pays a

premium in the missed-payout state and lowers that minimum, so optimal coverage goes to zero. The full proof is in Appendix A.3.

Proposition 2 (Marginal-utility amplification of the distress-sale discount) When wealth is positive and $u''' > 0$, a distress-sale discount raises the marginal-utility penalty of the direct premium cost in the missed-payout state. Write the absolute value of that direct cost as $MC_{B0} = c\pi_k p_{B0}^k u'(W_{B0})$. Then

$$\frac{\partial MC_{B0}}{\partial \Delta_{\text{fire}}} = (1 - \sigma)P\pi_k p_{B0}^k \gamma W_{B0}^{-\gamma-1} > 0 \quad (19)$$

and this sensitivity rises as W_{B0} falls.

The proof uses the chain rule. The derivative of W_{B0} with respect to the discount is -1 . $u'' < 0$ gives the positive sign, and $u''' > 0$ makes the sensitivity larger as wealth falls. If the discount is written as an endogenous distress sale, one also needs $\Delta_{\text{fire}}/\partial q = R_s - a_s + 2\kappa_s q \geq 0$, which is the condition that delayed value is at least the current marginal return on a distress sale. Economically, a spot-market distress sale moves wealth onto a steeper portion of marginal utility and thereby amplifies the burden of a missed payout. The insurance market and the spot market therefore sit on the same balance sheet. The feedback in Figure 3 has a mathematical counterpart here.

Proposition 3 (Pessimistic belief weights) Under smooth ambiguity, if $\phi' > 0$ and $\phi'' < 0$, the first-order condition can be written as an expected-marginal-utility condition under a distorted probability $\hat{\mu}$:

$$\sum_{k=1}^K [\hat{\mu}(k) \frac{\partial EU_k}{\partial \alpha}] = 0, \hat{\mu}(k) = \frac{\mu(k)\phi'(EU_k)}{\sum_j [\mu(j)\phi'(EU_j)]} \quad (20)$$

The proof is obtained by differentiating the objective function and applying a normalization. From $\phi'' < 0$, ϕ' falls with expected utility, so $\hat{\mu}(k)/\mu(k)$ falls with EU_k . More precisely, a lower-expected-utility model gets a larger weight multiple than a higher-expected-utility model. After normalization, it is not true that every lower-utility model must have an absolute weight above its prior. Figure 5 reports the weights at the baseline optimum. The heavier-disaster model receives a larger decision weight than its prior.




Figure 5 Belief distortion induced by ambiguity aversion

Note: Light bars are prior beliefs. Dark bars are the weights used in the decision. The values are computed at optimal coverage under the baseline parameters.

Proposition 3 says that ambiguity aversion tilts decision weight toward lower-expected-utility models. If those models also have a higher miss probability, the decision places less weight on the payout gain. That direction needs the pessimistic models to carry a higher miss rate. It does not follow from $\phi'' < 0$ alone. Section 4 reports the result under this paper's calibration: ambiguity aversion modestly reduces coverage, but the effect is clearly smaller than the miss probability and the cash constraint.

Proposition 4 (Local link between coverage and distress sales) Let $F(\alpha; \beta) = V_\alpha(\alpha; \beta)$. Suppose optimal coverage is interior and locally unique, and that $L > c$ and $V_{\alpha\alpha} < 0$. If the worst-state effect dominates, so that $F_{p_{B0}} < 0$ and $F_{\Delta_{fire}} < 0$, then both the miss probability and the distress-sale discount lower optimal coverage. In a neighborhood of the baseline parameters, $d\alpha^*/d\sigma > 0$, but that derivative shrinks as the miss probability and the strength of ambiguity aversion rise. If both payout outcomes sit on a differentiable stable branch and $(1 - p_{B0})(L - c)q_x(x_{B1}) > p_{B0}cq_x(x_{B0})$, then a fall in coverage raises expected distress sales in the disaster state, added up across payout outcomes. Under a linear-gap benchmark the condition simplifies to $(1 - p_{B0})L > c$.

Proposition 4 is written as a local, conditional comparative statics result. The implicit function theorem gives $d\alpha^*/d\beta = -F_\beta/V_{\alpha\alpha}$, so optimal coverage moves in the same sign as F_β . The negative signs on the miss rate and the discount come from the worst-state marginal-utility term. Whether the force of the subsidy fades monotonically with the miss rate or with ambiguity still depends on the outer weights and wealth in each state. It cannot be stated as a global theorem without extra conditions. The analysis is kept inside the baseline calibration and a neighborhood of it. Section 4 checks the key results with the same algorithm.

The role of initial wealth also needs to be stated, so that the paper is not read as saying that insurance cannot sell whenever a missed payout exists. If the worst state is still far from a survival floor, absolute risk aversion γ/W is low, and the premium is held down by the subsidy, full coverage can still be reasonable. Demand falls sharply only in the stretch where the worst state is pushed toward that floor and the steep rise in marginal utility makes the premium hard to bear. Once wage income, transfer income, and cash from migrant work are incorporated into initial wealth, they raise the worst-state wealth level, making such a decline harder to trigger. This paper focuses on smallholder farmers whose safety buffers are already thin and who are simultaneously squeezed by autumn-harvest quality losses and cash expenditures, not on all rural households.

Proposition 5 (Local effects of a quality-type climate shock) Let optimal coverage be interior, and let the cash constraint bind in the missed-payout state on the stable branch. If $V_{\alpha\xi} < 0$, and if the direct increase in distress sales from a fall in a_B and a rise in κ_B outweighs the premium saving from lower coverage, then $d\alpha^*/d\xi < 0$, $dq_{B0}/d\xi > 0$, and $dp_B/d\xi < 0$. The baseline parameterization of equation (17) meets these conditions for $\xi \in [0, 0.8]$.

A dominance condition is needed here. A rise in $p_{B0}(\xi)$ moves probability from the paid state to the missed state. A fall in a_B and a rise in κ_B raise distress sales at a given gap. A fall in R_B lowers the value of grain that is not sold at once. All three forces move toward lower coverage and a lower price. Lower coverage also cuts the premium already paid in the missed-payout state, which eases the cash gap a little. Proposition 5 requires the first three direct forces to dominate. Appendix A.8 gives the exact inequalities. The numerical work confirms that the condition holds on the interval reported in the paper.

In this calibration, a quantity path only raises π and cuts y_B . Coverage still falls, but the miss probability does not rise with climate severity, and cash after a payout can still cover the gap. For the same fall in Y_B , the quality path hurts coverage demand more.

Proposition 1(ii) already says that when γ is large enough, households are less willing to buy insurance that can miss a payout. Extend that result to an aggregation of heterogeneous households. Split households into two types. Let the share of older households be $\omega \in [0, 1]$. They have higher risk aversion than younger households, $\gamma_H > \gamma_L$, and a thinner disaster-window cash buffer, $A_B^H < A_B^L$. Both types face the same local market and each chooses its own coverage. Aggregate coverage is

$$\bar{\alpha}(\omega) = \omega\alpha^*(\gamma_H) + (1 - \omega)\alpha^*(\gamma_L) \quad (21)$$

Aggregate distress sales Q_B are solved from the fixed point of the common price. Missed and paid households, older and younger, all take the price as given. After the minimum sales are added up, the price returns to $p = a_B - \kappa_B Q_B$.

Proposition 6 (Heterogeneous households, aggregation, and the regional price) Let $\alpha_H^* < \alpha_L^*$ and $A_B^H < A_B^L$, and let older households have higher expected distress sales at a given common price. Let the common-price fixed point be locally unique and stable, $1 + \kappa_B Q_p > 0$, where $Q_p < 0$ is the derivative of total distress sales with respect to price at a given price. Then $d\bar{\alpha}/d\omega < 0$, $dQ_B/d\omega > 0$, and $dp_B/d\omega < 0$. If the direct widening of the sales gap between the two types outweighs the feedback from the clearing price as total sales change, the local dominance condition in Appendix A.9 holds and $\partial^2 Q_B / (\partial \omega \partial \xi) > 0$.

The first set of results comes from type replacement: as the share of older households rises, total distress sales move with it. Lower coverage and a thinner cash buffer together imply higher expected distress sales for older households at a given price. At a stable fixed point, total sales rise and the common price falls. The interaction with climate severity also requires the direct widening of the type gap to dominate the feedback from the slope of the fixed point. The numerical results in this paper meet that local condition.

The six propositions take up missed payouts, amplification through distress sales, distorted beliefs, the local link, climate severity, and aggregation across heterogeneous households. Together, they form the causal chain behind the two field facts: protection failure sets off the distress sale, and the distress sale then further weakens the value of protection in the household's eyes.

4 Numerical results, regional aggregation, and policy implications

4.1 Parameters, the solution method, and coverage demand

The numerical work normalizes household resources by a normal-year net resource unit. Baseline parameters are a unit premium of 0.40, a subsidy rate of 0.80, a unit indemnity of 0.80, a rigid cash outlay of 0.70, relative risk aversion of 2.0, and ambiguity aversion of 0.80. In the normal state, cash before the sale, the harvest, the price intercept, the price slope, and delayed unit value are 1.00, 1.20, 1.10, 0.20, and 1.12. In the disaster state, they are 0.25, 1.00, 1.00, 0.35, and 1.10. The three candidate climate models have disaster, miss, and false-payout probabilities (0.15, 0.10, 0.02), (0.25, 0.30, 0.03), and (0.35, 0.55, 0.05), with second-order beliefs 0.30, 0.40, and 0.30. Coverage is searched on a 500-point grid on $[0,1]$. The algorithm, grid robustness, and a per-mu reading are in Appendix B. The curves below all come from the same grid search.

In Figure 6, the horizontal axis is the miss probability, and the vertical axis is coverage or distress sales. The solid line is optimal coverage under smooth ambiguity. The dashed line is the expected-utility check with a linear outer function. The dotted line is expected distress sales in the disaster state. When the miss probability rises from 0.10 to 0.50, optimal coverage falls from 0.966 to 0.356 and expected distress sales rise from 0.070 to 0.408. The expected-utility curve lies a little above the ambiguity curve and moves in the same direction. The missed payout itself pulls demand down.




Figure 6 Optimal coverage and expected disaster-state distress sales as the miss probability changes

Note: The curves come from a grid search. The dashed vertical lines mark 0.10 and 0.50.

Figure 7 puts the subsidy rate and the miss rate on the same figure. To isolate the interaction, Figure 7 sets the miss probability in all three candidate models to 0.10 or to 0.50.¹⁶ When the miss rate is 0.10, raising the subsidy from 0 to 0.80 can lift coverage from 0.224 to 1. When the miss rate is 0.50, the same subsidy increase raises coverage only to 0.426. The subsidy still helps, but a high miss rate limits the demand a given subsidy can buy. This aligns with Cole et al. (2013), who find that cutting the price still leaves the adoption problem in place.

¹⁶ When Figure 6 moves the center it keeps the probability gaps among the candidate models, so coverage at 0.50 need not match across the two figures.




Figure 7 How the effect of the subsidy on optimal coverage changes with the miss probability

Note: Both curves are computed using the same parameters and algorithm.

The other comparative statics are summarized as follows; details are in Table 3. A thicker cash buffer in the disaster window raises coverage and cuts distress sales. When pre-sale disaster cash rises from 0.15 to 0.45, optimal coverage rises from 0.276 to 1.000 and expected distress sales fall from 0.525 to 0.114. A steeper spot-price slope, that is, a thinner market, makes the price fall from a distress sale deeper, and insurance demand weakens with it. When the slope rises from 0.15 to 0.47, coverage falls from 0.908 to 0.282 and the spot price falls from 0.974 to 0.806. As relative risk aversion rises from 0.5 to 5.0, coverage in this calibration falls from 1.000 to 0.198 and distress sales rise from 0.211 to 0.421. As ambiguity aversion rises from 0 to 2.4, coverage only falls from 0.582 to 0.526. The sign is negative and the size is small. Ambiguity is therefore not the main reason demand falls sharply. It changes decision weights, but under this calibration the quantity effect is smaller than that of the miss rate and the cash constraint.

When delayed sale value changes alone, coverage and distress sales change little, yet public drying still matters. They barely move because the minimum-sale rule already ties the sale to the cash gap. The household will not sell less just because delayed value is higher, unless the gap itself shrinks. For drying and storage to work, they have to raise cash that can turn over in the disaster window, keep more buyers so the price slope falls, or protect quality so the household does not have to clear the bin before the quality line arrives.

Table 3 Numerical comparative statics and the matching propositions

Parameter change	Coverage	Distress sales	Proposition
Miss probability rises	Falls, and can fall very low	Rises	Propositions 1 and 4
Risk aversion rises	Falls monotonically in this calibration	Rises	Proposition 1

Parameter change	Coverage	Distress sales	Proposition
Distress-sale discount rises	Falls	It is the distress sale	Proposition 2
Ambiguity aversion rises	Falls a little in this calibration	Rises indirectly	Proposition 3
Subsidy rate rises	Rises; weaker when the miss rate is high	Falls indirectly	Proposition 4
Cash buffer rises	Rises	Falls	Propositions 2 and 4
Climate severity rises	Falls	Rises; the price falls	Proposition 5
Share of older households rises	Aggregate coverage falls	Regional distress sales rise	Proposition 6

Splitting the disaster state into missed and paid cases makes the mechanism more concrete. Table 4 reports wealth, distress sales, and the spot price as coverage moves from 0 to 1 under the baseline. With no insurance, missed and paid are the same. The household must sell 0.560 to cover the rigid outlay of 0.70, the spot price falls to 0.804, and terminal wealth is 0.484. At coverage 0.50, the paid household almost need not sell in distress and wealth rises to 0.998; the missed household has to sell 0.628 and wealth falls to 0.409. At full cover, the paid household sells nothing and wealth is 1.370; the missed household sells 0.703 and is left with wealth at 0.327. Insurance is cash when it pays. When it misses, it is a cost already paid that is still using up liquidity. That is why a higher miss probability makes households buy less insurance and sell more grain at the same time.

Table 4 How coverage rewrites missed and paid outcomes in the disaster state (baseline)

Coverage	Missed wealth	Paid wealth	Missed sales	Paid sales	Missed price	Paid price
0.00	0.484	0.484	0.560	0.560	0.804	0.804
0.25	0.448	0.768	0.593	0.302	0.792	0.894
0.50	0.409	0.998	0.628	0.093	0.780	0.967
0.75	0.369	1.190	0.665	0.000	0.767	1.000
1.00	0.327	1.370	0.703	0.000	0.754	1.000

4.2 Regional market equilibrium, climate severity, and the older-household share

When many disaster-affected farm households enter the same local market, both missed and paid sellers face the same spot price. Missed households have a larger cash gap and sell more. Total sales push the common price down, so paid households have to accept a lower price even if they are less in a hurry. Solving this fixed point at the baseline, a miss probability of 0.10 gives regional distress sales of 0.054 and a common price of 0.981; missed households sell 0.537 and paid households need not sell in distress. At a miss probability of 0.50, regional distress sales rise to 0.389 and the common price falls to 0.864; missed households sell 0.554 and paid households are also forced to sell 0.224. The extra sale by paid households is dragged down by the common price. Their own gap has not widened in step. Household decisions thus become regional price pressure. In the macro numbers, national autumn grain in 2025 was

536.615 million tonnes, up 1.5 percent, and full-year grain was 714.875 million tonnes, up 1.2 percent. The National Bureau of Statistics also said harvest-time rain in the Huang-Huai-Hai region left some crops moldy, and that autumn grain is reported after deducting mold losses.¹⁷ National output can rise while local disaster-hit balance sheets are hurt. Once scattered distress sales line up in time, they show up as a lower regional price and lower farm income, and they can pass through to next-season inputs, local public finance, and grassroots financial institutions. Timely payouts and liquidity arranged in advance can therefore bridge a single household's cash gap and support regional price stability.

Figure 8 puts climate severity ξ on the horizontal axis. At $\xi = 0$, the model returns to the original calibration: coverage 0.562, expected disaster-state distress sales 0.223, and a spot price of 0.922. When ξ rises to 0.40, coverage falls to 0.372, distress sales rise to 0.375, and the price falls to 0.836. When ξ rises further to 0.80, coverage is 0.214, distress sales are 0.515, and the price is 0.747. Three forces move together: p_{B0} rises from 0.30 to 0.46, Y_B falls from 1.100 to 1.004, and the market's absorptive capacity weakens.




Figure 8 Climate severity, optimal coverage, expected disaster-state distress sales, and the spot price
Note: The left panel is coverage or distress sales. The right panel is the price.

Hold Y_B at 1.004 and compare the two paths. The difference is clearer. On the quality path, coverage is only 0.214, expected distress sales are 0.515, and the miss rate rises to 0.46. On the quantity path, the harvest is cut to 0.913, and the baseline model's disaster probability is lifted from 0.25 to 0.35, while the miss rate stays at 0.30. Coverage can still reach 0.494 and distress sales are only 0.257. The same drop in Y_B hurts insurance more on the quality path. That is the trouble with the autumn rain. Output need not fall, loss assessment is harder, households are less willing to buy, and they have less choice but to sell.

¹⁷ The definition is from Wei Fenghua, director of the Rural Department of the National Bureau of Statistics, on grain output, 12 December 2025. Section 2 already cites the same source.

Figure 9 puts the older-household share ω on the horizontal axis. Younger households have $\gamma_L = 1.5$ and disaster cash 0.28. Older households have $\gamma_H = 3.0$ and disaster cash 0.18. At $\xi = 0$, the two optimal coverage rates are 0.814 and 0.224. As ω rises from 0 to the census share 0.336, aggregate coverage falls from 0.814 to 0.616, regional distress sales rise from 0.154 to 0.259, and the common price falls from 0.946 to 0.910. At $\xi = 0.40$, the same census share gives aggregate coverage of only 0.409, regional distress sales of 0.333, and a common price of 0.852. As ω moves from 0 to 0.60, the increment in regional distress sales is 0.194 at $\xi = 0$ and 0.235 at $\xi = 0.40$. The more uncertain the climate, the more aging pushes the price down.




Figure 9 The older-household share, regional distress sales, and the common price

Note: The solid line is $\xi = 0$. The dashed line is $\xi = 0.40$. The dashed vertical line marks the 33.6 percent share of agricultural workers aged 55 and above in the Third National Agricultural Census.

33.6 percent is only the 2016 age mix of agricultural workers. It is a checkable reference for ω . Once household choices are added up, they become regional distress sales and a common price. The next subsection joins this aggregate layer to the policy discussion and gives three suggestions.

4.3 Macroeconomic consequences and policy implications

The 2024 Central Document No. 1 called for "more accurate agricultural insurance enrollment and claims, with payouts made whenever they are due." On 14 November of the same year, the financial regulator, the Ministry of Finance, the Ministry of Agriculture and Rural Affairs, and the National Forestry and Grassland Administration issued a notice on more accurate enrollment and claims (Jin Fa [2024] No. 35). The notice calls for more granular payout rules by crop, disaster, and growth stage, updated technical norms for loss identification, local loss committees, and room for third-party inspectors. It also supports advance payouts, a higher advance share, claim-cycle thresholds by line, a ban on delayed or withheld payouts, and direct payment of indemnities to the insured. The first set of rules maps to a lower miss

rate. The second maps to holding the disaster-window cash gap. In the three suggestions below, loss-assessment review and disaster-window working capital align with the current notice. Public drying is not on that list. It is an additional arrangement highlighted by the model.

After the model is aggregated, three aggregate implications follow. First, climate risk is an aggregate shock that is hard to spread locally. Standard insurance relies on the law of large numbers and pools independent individual risks. When a prolonged rain event hits the Huang-Huai-Hai region at once, local mutual aid and commercial reinsurance both tighten, and the risk has to move up to reinsurance, the fiscal budget, and the national food-security system. Whether a household buys insurance then becomes a question of how aggregate risk is allocated.

Second, price spillovers from bunched distress sales are a pecuniary externality. If one household sells a little more, everyone faces a lower price. In plainer words, a household that sells extra to cover a cash need leaves its neighbors with a lower price. In a complete market, this spillover need not mean inefficiency. When markets are incomplete and liquidity constraints bind, no household puts "my extra sale will push down everyone's price" into its own cost, so the group sells too much and the price falls too far (Lorenzoni, 2008; Shleifer and Vishny, 1992, 2011). That is the amplification and propagation the financial accelerator literature writes about, where the balance sheet is firm net worth and the price is a collateral price. Here the balance sheet is a household cash ledger and the price is the local grain price (Kiyotaki and Moore, 1997; Bernanke et al., 1999). A weather shock one household might absorb can become a regional shock, and it can last into the next season. Policy intervention therefore has an efficiency rationale.

Third, household fragility can aggregate into a food-security risk. Stable aggregate output is not the same as the absence of household loss. When many smallholder farmers face cash and quality constraints at the same time, next-season seed, fertilizer, and land rent are all affected. Households cut spending after the disaster; in aggregate, this can reduce next-season production capacity. At the base of food security is whether these micro agents can still plant after a bad year. When the constraint binds, the spot price sits below the delayed value of the same grain, and the discount deepens as distress sales grow. The lower the price, the worse the worst state, and the less protection looks worth buying. The less they buy, the heavier the distress-sale pressure in the next disaster, and the more they sell. Each household's choice is rational. Together, these choices can sustain vulnerability across disasters. $R_B - p_B(q) > 0$

The numbers do not support writing policy as keep raising the subsidy in a straight line and that will do. The subsidy helps, especially when the miss rate is low. When the miss rate is high, the same subsidy buys clearly less coverage. What can raise insurance demand and cut distress sales at the same time is a lower miss rate, cash held through the disaster window, and

a way for wet grain to wait instead of being sold at once. The three suggestions below each map to one parameter in the model.

First, improve loss assessment to detect quality loss, and lower the miss probability. If current assessment leans on high-altitude remote sensing or a regional average yield, cases that "still look green while the kernel is already ruined" are easy to miss. Putting low-altitude remote sensing, field sampling, and a reviewable quality test into one assessment process is a way to turn quality loss into evidence that can be checked and appealed. After the miss probability falls, the limit penalty in Proposition 1 is weaker and the pessimistic weights in Proposition 3 are lighter, so the marginal effect of the subsidy can recover.

Second, build public drying and ventilated storage that can operate effectively during prolonged rain windows. Public drying works through at least three channels: it protects the delayed value R_s of grain that is not sold in distress, it keeps more buyers and thereby lowers κ_s , and it can raise disaster-window working cash A_s through warehouse receipts or a purchase arrangement. Under the minimum-sale rule used here, raising R_s alone only improves the return on retained grain. It does not directly change the sale pinned by the cash gap. Lowering κ_s or raising A_s is what cuts distress sales directly. Once that split is made, the policy language and the model's comparative statics line up.

Third, in the days after a weather warning and before the indemnity arrives, put working capital in place that can arrive quickly and close the cash gap. Warehouse receipts or grain pledges, pre-signed emergency lines, and treating days-to-payment as a performance measure all belong here. In the model, this is a higher post-disaster cash balance A_s or a smaller share of the rigid outlay D that must be paid at once. After the gap shrinks, the gap in distress sales between paid and missed cases also narrows, and regional price pressure eases. The most direct score for this suggestion is how many days pass from the claim to the actual payment.

Table 5 Policies, parameters, and mechanisms

Policy	Matching parameter	Role
Loss-assessment review and quality identification	Miss probability falls	Eases the worst-state penalty and restores the marginal value of insurance
Public drying and storage	Delayed value rises, the price slope falls, disaster cash rises	Preserves value; capacity and working capital then cut distress sales directly
Disaster-window working capital	The cash buffer rises; rigid outlays are postponed	Holds the gap until the indemnity arrives
Comparison: raise the subsidy only	The subsidy rate rises	Helps, but much weaker when the miss rate is high

If agricultural insurance is placed in a wider climate-risk financing frame, the three measures should be split by how often the risk arrives and how deep the loss is. High-frequency, shallower quality swings lean more on drying, storage, and short-term working capital. Low-

frequency, very deep, spatially correlated climate extremes need the central budget and reinsurance to set triggers and funding channels before the disaster. Grassroots loss assessment determines whether insurance can work in the few days when households need it most. Without that layer, the subsidy scale above is hard to turn into a sense that the policy works in the field, and it is hard to hold the regional price.

5 Conclusion

The paper starts with quality loss in the 2025 autumn rain, combines payout mismatch, the cash constraint, and distress sales into a single household decision, and aggregates those decisions to the regional price. The model has two stages. Households choose coverage before the disaster and choose the distress sale after they see the payout. A payout shrinks the cash gap and cuts distress sales. A missed payout leaves the premium already paid and the indemnity still unpaid, exactly when cash is tightest, and raises the pressure to sell. The analysis shows that a higher miss probability raises expected disaster-state distress sales and lowers the marginal value of insurance. The numbers then show that the cash buffer and the market's absorptive capacity matter a great deal. The subsidy always helps, but a high miss rate means a given subsidy buys less demand. The roles of risk aversion and ambiguity depend on the parameters. A quality path hurts insurance more than a quantity path. A higher share of older households pushes aggregate distress sales up and the regional price down. A lower price cuts income and can pass through to next-season inputs, local public finance, and grassroots financial institutions. Timely payouts and liquidity arranged in advance therefore act on both the household cash gap and regional price clearing.

The six propositions gather the two field facts that looked odd on their own onto one line of logic. Under ideal insurance, a subsidized risk-averse household should buy full cover. Once a missed payout exists, the worst-state marginal-utility penalty can outweigh the incentive to buy, and the most risk-averse households are the least willing to buy. The distress-sale discount amplifies that penalty through prudence. Ambiguity aversion is like deciding under a more pessimistic measure, which makes it harder for the subsidy alone to restore demand. Climate severity moves a_B , R_B , κ_B , and p_{B0} together, and coverage and the price move down together. High risk aversion and thin cash buffers among older households become regional price pressure when combined. Protection failure and the distress sale are therefore two sides of the same mechanism.

On policy, climate risk financing cannot rely on a single instrument and it requires a policy mix. Premium subsidies are one tier of the climate risk financing system. Grassroots loss assessment and verification, public drying services, disaster-period working capital, central government fiscal support, and reinsurance need differentiated roles by risk size, with trigger

conditions and transmission channels specified ex ante. One of the clearest performance measures is the number of days between a claim and payment.

The paper has several limits. The paper simplifies the price function and the storage return. The parameters are not estimated on household data. The aggregation does not solve a full general equilibrium, and insurers' assessment costs and strategic behavior are outside the model. Later work needs matched data on weather, loss, claims, sales, and cash flow to test the predictions here: a higher miss rate should come with both lower effective enrollment and higher post-disaster distress sales; if the indemnity arrives well after the cash outlay, distress sales should be more bunched; and if public drying or working capital is truly available inside the window, the quality-specific price gap should narrow.

Appendix A Mathematical derivations

A.1 Notation, regularity, and optimality

Baseline utility is CRRA. When $\gamma \neq 1$, $u(W) = W^{1-\gamma}/(1-\gamma)$. When $\gamma = 1$, $u(W) = \ln W$. The domain is $W > 0$. Then $u'(W) = W^{-\gamma} > 0$, $u''(W) = -\gamma W^{-\gamma-1} < 0$, and $u'''(W) = \gamma(\gamma+1)W^{-\gamma-2} > 0$. As $W \rightarrow 0^+$, $u'(W) \rightarrow +\infty$. The outer function uses $\phi_\rho(x) = (1 - e^{-\rho x})/\rho$ ($\rho > 0$) and $\phi_0(x) = x$. Then $\phi'_\rho > 0$ and $\phi''_\rho \leq 0$, and the criterion returns continuously to expected utility as $\rho \rightarrow 0$. Absolute and relative risk aversion are $A(W) = -u''(W)/u'(W) = \gamma/W$ and $R(W) = -Wu''(W)/u'(W) = \gamma$.

Write the objective as $V(\alpha)$, $\alpha \in [0,1]$. An interior optimum satisfies $V'(\alpha^*) = 0$ and $V''(\alpha^*) < 0$. If the derivative is always negative or always positive on the interval, the optimum is at an endpoint. Appendix A.7 proves that V is concave when wealth is positive, the distress-sale solution sits on the stable branch, and $R_s \geq a_s$. If those global conditions fail, later implicit comparative statics are used only near a local optimum that satisfies $V''(\alpha^*) < 0$. If $V'(\alpha^*; \beta) = 0$ defines $\alpha^*(\beta)$ implicitly, then

$$\frac{d\alpha^*}{d\beta} = -\frac{\partial^2 V / \partial \alpha \partial \beta}{\partial^2 V / \partial \alpha^2} \quad (22)$$

The sign matches the cross partial. Later comparative statics are all derived from this.

A.2 The ideal-insurance benchmark

Under ideal insurance there are only two wealth states: $W_G = W_0 - C + Y_G - (1-\sigma)P\alpha$ and $W_B^+ = W_0 - C + V_B + \alpha L - (1-\sigma)P\alpha$. The derivative is

$$EU'(\alpha) = -(1-\sigma)P(1-\pi)u'(W_G) + [L - (1-\sigma)P]\pi u'(W_B^+) \quad (23)$$

At $\alpha = 1$ the two wealth levels are equal. Substitute $W_G = W_B^+ = W^*$ into equation (23) to get $EU'(1) = u'(W^*)[\pi L - (1-\sigma)P]$. If $(1-\sigma)P < \pi L$, then $EU'(1) > 0$. Because EU is concave, its derivative does not rise with α , so $EU'(\alpha) \geq EU'(1) > 0$ holds for every $\alpha \in [0,1]$ and the optimum is $\alpha^* = 1$. This is more direct than checking the derivative only at $\alpha = 0$, and it proves that the right endpoint is optimal.

A.3 Missed payouts and coverage demand: Proposition 1

Write $m_{B0}(\alpha) = \partial W_{B0} / \partial \alpha < 0$. Suppose there exists $\underline{m} > 0$ such that $-m_{B0}(\alpha) \geq \underline{m}$. Write the first-derivative contribution of the remaining states as $B(\alpha)$, and suppose $B(\alpha) \leq \bar{B} < \infty$. Then $EU'_k(\alpha) \leq -Au'(W_{B0}) + \bar{B}$, where

$$A \equiv \pi_k p_{B0}^k \underline{m} > 0, -m_{B0}(\alpha) \geq \underline{m} \quad (24)$$

Choose $\bar{W}$ so that $A\bar{W}^{-\gamma} > \bar{B}$. If $W_{B0}(\alpha) \leq \bar{W}$ holds for every $\alpha \in [0,1]$, then $EU'_k(\alpha) < 0$ on the whole interval, so optimal coverage is 0. This proves Proposition 1(i). The condition is written as an upper bound on the whole interval so that a low wealth at one coverage rate is not mistaken for a global result.

For Proposition 1(ii), let the probabilities of a finite set of states be $\varpi_j > 0$ and wealth be $W_j(\alpha) > 0$. The certainty equivalent of CRRA expected utility is

$$CE_\gamma(\alpha) = \left\{ \sum_j [\varpi_j W_j(\alpha)^{1-\gamma}] \right\}^{\frac{1}{1-\gamma}}, \lim_{\gamma \rightarrow \infty} CE_\gamma(\alpha) = \min_j W_j(\alpha) \quad (25)$$

The generalized-mean theorem gives $CE_\gamma(\alpha) \rightarrow \min_j W_j(\alpha)$ ($\gamma \rightarrow \infty$). If $W_{B0}(0) = \min_j W_j(0)$ and $W_{B0}(\alpha) < W_{B0}(0)$ holds for every $\alpha > 0$, then for any $\varepsilon > 0$ the largest minimum wealth on the compact set $[\varepsilon, 1]$ is strictly below the value at $\alpha = 0$. For γ large enough the optimum cannot sit in $[\varepsilon, 1]$. Because ε is arbitrary, $\alpha^*(\gamma) \rightarrow 0$. The argument does not require the first derivative at $\alpha = 0$ to be negative.

Once the cash constraint is written in, the missed-payout state has $W_{B0} = R_B(y_B - q_{B0})$ and $\partial q_{B0}/\partial \alpha = (\partial q_{B0}/\partial x_{B0})c > 0$, so $W_{B0}(\alpha)$ falls strictly. The distress-sale channel therefore delivers the minimum-wealth condition needed for Proposition 1(ii).

A.4 Existence and monotonicity of the distress-sale solution

Lemma Let $\bar{q}_s = \min\{y_s, a_s/(2\kappa_s)\}$ and $\bar{X}_s = (a_s - \kappa_s \bar{q}_s)\bar{q}_s$. If $0 < x_{s\theta} < \bar{X}_s$, the cash constraint has a unique solution on the feasible stable branch $0 < q < \bar{q}_s$, given by the smaller quadratic root in the main text.

Proof. The function $r(q) = (a_s - \kappa_s q)q$ is strictly increasing on $[0, \bar{q}_s]$ and rises continuously from 0 to $\bar{X}_s$. So any $x_{s\theta} \in (0, \bar{X}_s)$ maps to a unique feasible q . The discriminant is positive and the smaller root sits on the stable branch. The larger root sits on the downward branch of the revenue curve, or it exceeds the harvest cap y_s , so it is not "the smallest distress sale that meets the cash constraint."

Differentiating the constrained solution gives

$$\frac{\partial q_{s\theta}}{\partial x_{s\theta}} = \frac{1}{\sqrt{a_s^2 - 4\kappa_s x_{s\theta}}} > 0 \quad (26)$$

A larger gap means a larger minimum distress sale. Relative to a payout, a missed payout raises x by αL , so $q_{B0} > q_{B1}$ (on the interval where both are positive). The spot price $p_s = a_s - \kappa_s q$ then falls and the discount $\Delta_{\text{fire}} = (R_s - p_s)q$ rises. This turns the exogenous discount in Proposition 2 into a comparative static of the cash gap.

A.5 Amplification of the distress-sale discount: Proposition 2

Let $MC_{B0}(\Delta_{\text{fire}}) = (1 - \sigma)P\pi_k p_{B0}^k u'(W_{B0}(\Delta_{\text{fire}}))$, and let the derivative of W_{B0} with respect to the discount be -1 . The chain rule gives

$$\begin{aligned} \frac{dMC_{B0}}{d\Delta_{\text{fire}}} &= (1 - \sigma)P\pi_k p_{B0}^k u''(W_{B0}) \cdot (-1) \\ &= (1 - \sigma)P\pi_k p_{B0}^k \gamma W_{B0}^{-\gamma-1} > 0 \end{aligned} \quad (27)$$

Differentiating again with respect to W_{B0} gives a negative value. The lower the wealth, the larger the marginal-utility penalty of one more unit of discount. If the discount is replaced by an endogenous q_{B0} , the expression is also multiplied by $\partial\Delta_{\text{fire}}/\partial q_{B0} = R_s - a_s + 2\kappa_s q_{B0}$. The sufficient condition $R_s \geq a_s$ keeps that term non-negative. A weaker local condition is $R_s \geq a_s - 2\kappa_s q_{B0}$.

A.6 Belief distortion under smooth ambiguity: Proposition 3

The objective is $V(\alpha) = \sum_k \mu(k) \phi(EU_k(\alpha))$. The chain rule gives

$$V'(\alpha) = \sum_k \mu(k) \phi'(EU_k) EU_k'(\alpha) \quad (28)$$

Let $Z = \sum_j \mu(j) \phi'_\rho(EU_j) > 0$ and define $\hat{\mu}(k) = \mu(k) \phi'_\rho(EU_k)/Z$. Then $\hat{\mu}(k) \geq 0$ and $\sum_k \hat{\mu}(k) = 1$, and $V'(\alpha) = 0$ is equivalent to $\sum_k \hat{\mu}(k) EU_k'(\alpha) = 0$. For the normalized exponential, $\phi'_\rho(x) = e^{-\rho x}$, so the distorted weights are still

$$\hat{\mu}(k) = \frac{\mu(k) \exp(-\rho EU_k)}{\sum_j \mu(j) \exp(-\rho EU_j)} \quad (29)$$

Figure 5 is computed from this formula at the baseline optimum.

A.7 Concavity of the objective and the local link: Proposition 4

For a given model k , write $EU_k(\alpha) = \sum_{s,\theta} \varpi_{k,s\theta} u(W_{s\theta}(\alpha))$. When the cash constraint is slack, $W_{s\theta}$ is affine in α . When it binds, $W_{s\theta} = R_s(y_s - q_{s\theta})$ and $q_{xx} = 2\kappa_s/(a_s^2 - 4\kappa_s x)^{3/2} > 0$, so $W_{s\theta}' = -R_s q_{xx}(c - \theta L)^2 \leq 0$. The condition $R_s \geq a_s$ also keeps the slope from jumping up at the splice $x = 0$. Wealth in each state is therefore concave, and $EU_k' \leq 0$. Differentiating V again:

$$V''(\alpha) = \sum_k \{\mu(k) [\phi''(EU_k)(EU_k')^2 + \phi'(EU_k) EU_k'']\} \leq 0 \quad (30)$$

Both terms in equation (30) are non-positive, so V is concave. If at least one term is strictly negative, the interior solution is unique. The result uses positive wealth, the stable branch, and $R_s \geq a_s$.

Let $F(\alpha^*, \beta) = V_\alpha(\alpha^*; \beta) = 0$. The implicit function theorem gives $d\alpha^*/d\beta = -F_\beta/V_{\alpha\alpha}$. Because the denominator is negative, the derivative has the same sign as F_β . When $L > c$ and the worst-state term dominates, raising p_{B0} or Δ_{fire} makes $F_\beta < 0$ and coverage falls.

In general, write $m_{s\theta} = \partial W_{s\theta}/\partial \alpha$. The cross partial with respect to the subsidy rate is

$$\frac{\partial EU'_k}{\partial \sigma} = \sum_{s,\theta} \{\bar{\omega}_{k,s\theta} [u''(W_{s\theta}) W_{s\theta,\sigma} m_{s\theta} + u'(W_{s\theta}) m_{s\theta,\sigma}]\} \quad (31)$$

In the affine benchmark with a slack cash constraint, $W_{s\theta,\sigma} = P\alpha$ and $m_{s\theta,\sigma} = P$. After the constraint binds, both terms also contain q_x and q_{xx} . The subsidy cross partial therefore has no unconditional global sign. The baseline grid and a neighborhood of it give $d\alpha^*/d\sigma > 0$, and the derivative shrinks when the miss rate is high and ambiguity is stronger.

A note on the regional fixed point. Disaster-hit households of unit mass share the price $p = a_B - \kappa_B Q$. The missed household's gap is $x_0 = D - A_B + (1 - \sigma)P\alpha$ and the paid household's gap is $x_1 = x_0 - \alpha L$. A price-taking minimum distress sale is $q_i = \min\{y_B, \max\{0, x_i/p\}\}$. Total sales are $Q = p_{B0}q_0 + (1 - p_{B0})q_1$. Iterating on the interval where $p > 0$ yields a stable point. Individual decisions ignore their own effect on p , so an externality is present.

A.8 Local comparative statics of climate severity: Proposition 5

Let $G(\alpha, \xi) = V_\alpha(\alpha; \xi) = 0$ describe the interior optimum. From Appendix A.7, $V_{\alpha\alpha} < 0$, so $d\alpha^*/d\xi$ has the same sign as $G_\xi = V_{\alpha\xi}$. Proposition 5 assumes $G_\xi < 0$ and further requires that the direct worsening of market conditions outweigh the premium saving from lower coverage.

The miss channel is $p_{B0}(\xi) = p_{B0}^0 + \lambda_p \xi$, $\lambda_p > 0$. Raising p_{B0} moves probability from the paid state to the missed state. When $L > c$ and the worst-state term dominates, this lowers G_ξ . The condition is checked with numerical derivatives on the calibration interval used here.

The gap channel. On the stable branch, $q = (a - \sqrt{a^2 - 4\kappa x})/(2\kappa)$ and $\partial q/\partial x = 1/\sqrt{a^2 - 4\kappa x} > 0$. Differentiating with respect to a and κ :

$$\frac{\partial q}{\partial a} = \frac{1}{2\kappa} \left(1 - \frac{a}{\sqrt{a^2 - 4\kappa x}}\right) < 0, \quad \frac{\partial q}{\partial \kappa} = \frac{-q\sqrt{a^2 - 4\kappa x} + x}{\kappa\sqrt{a^2 - 4\kappa x}} > 0 \quad (32)$$

The second derivative is positive because, at the solution $x = aq - \kappa q^2$, the numerator becomes $\kappa q^2 > 0$. Equivalently, $\partial q / \partial \kappa = q^2 / (a - 2\kappa q) > 0$. So a fall in a_B and a rise in κ_B both raise q_{B0} at a given x . After the feedback through optimal coverage is included, a sufficient condition is $q_a a'_B + q_\kappa \kappa'_B + q_x c \alpha^{*'}(\xi) > 0$. The last term is negative. Proposition 5 requires the first two terms to dominate.

The delayed-value channel. $W_{B0} = R_B(y_B - q_{B0})$, so $\partial W_{B0} / \partial R_B = y_B - q_{B0} \geq 0$. A fall in R_B lowers wealth and raises the worst-state marginal-utility cost, but under the minimum-sale rule it does not change q directly. The quality path holds y_B fixed, so $dY_B / d\xi = y_B dR_B / d\xi < 0$.

Taken together, $G_\xi < 0$ gives $d\alpha^* / d\xi < 0$, and the dominance condition $q_a a'_B + q_\kappa \kappa'_B + q_x c \alpha^{*'} > 0$ gives $dq_{B0} / d\xi > 0$. From $p_B = a_B - \kappa_B q_{B0}$ one then has $dp_B / d\xi = a'_B - \kappa'_B q_{B0} - \kappa_B q'_{B0} < 0$. The baseline parameters of equation (17) satisfy these inequalities on $\xi \in [0, 0.8]$.

$$\frac{dp_B}{d\xi} = \frac{da_B}{d\xi} - q_{B0} \frac{d\kappa_B}{d\xi} - \kappa_B \left(\frac{\partial q_{B0}}{\partial \xi} + \frac{\partial q_{B0}}{\partial \alpha} \frac{d\alpha^*}{d\xi} \right) \quad (33)$$

Equation (33) is only the expansion of that total price differential. Without the dominance condition, the premium saving from lower coverage can offset part of the direct distress-sale pressure, so $q'_{B0} > 0$ cannot be read off $a'_B < 0$ and $\kappa'_B > 0$ alone. The numerical work confirms that q_{B0} and expected disaster-state distress sales both rise monotonically, and the price falls monotonically, on the interval reported in the paper.

A.9 Heterogeneous households, aggregation, and the regional price: Proposition 6

Let $\alpha_H^* < \alpha_L^*$. Aggregate coverage is $\bar{\alpha} = \omega \alpha_H^* + (1 - \omega) \alpha_L^*$, so $d\bar{\alpha} / d\omega = \alpha_H^* - \alpha_L^* < 0$.

Let $Q(p, \omega)$ be total minimum distress sales at a given common price. The fixed point satisfies $Q = Q(a_B - \kappa_B Q, \omega)$. Write $Q_\omega > 0$ for the fact that older households have higher expected distress sales at a given price. In this parameterization that inequality comes from the cash buffer and coverage together, not from either condition alone. A lower price raises the required sale, so $Q_p < 0$. Implicit differentiation gives $dQ / d\omega = Q_\omega / (1 + \kappa_B Q_p)$. The stability condition $1 + \kappa_B Q_p > 0$ then yields $dQ / d\omega > 0$, and $p = a_B - \kappa_B Q$ yields $dp / d\omega < 0$.

For the interaction with climate severity, let $H(Q, \omega, \xi) = Q - Q(a_B(\xi) - \kappa_B(\xi)Q, \omega, \xi)$. Along the stable fixed point, $H_Q > 0$ and $Q_{\omega\xi} = -[H_{\omega\xi} + H_{\omega Q} Q_\xi + (H_{Q\xi} + H_{QQ} Q_\xi) Q_\omega] / H_Q$. A negative term in brackets is therefore a local sufficient condition for $Q_{\omega\xi} > 0$. A widening type gap in sales at a given price, as ξ rises, is the pushing term, but stability by itself does not rule out a reverse effect from a change in the slope of the fixed point. The parameterization

used here meets the dominance condition after a numerical check: as ω rises from 0 to 0.60, the increment in regional distress sales is 0.194 at $\xi = 0$ and 0.235 at $\xi = 0.40$.

2026 S.-T. Yau High School Science Award
仅用于2026丘成桐中学科学奖论文公示

Appendix B Numerical algorithm and parameters

The baseline comparative statics, climate severity, aging aggregation, and the interaction all use the same grid search. When climate severity is zero the model returns to the original calibration. Household resources are normalized by a normal-year net resource unit. Coverage is enumerated on $\{0, 1/500, \dots, 1\}$, and the point that maximizes the smooth-ambiguity objective is kept. On the same baseline, refining the grid to 100, 250, 500, and 1,000 points gives optimal coverage 0.560, 0.564, 0.562, and 0.562. The difference is no larger than one grid step, and the objective changes by no more than order 10^{-6} . Assign coverage rates that produce an infeasible cash gap a value of minus infinity, and they do not enter the optimum.

Table B1 Baseline parameters (normalized units)

Parameter	Value	Meaning
Unit premium	0.40	Actuarial premium for full cover (normalized)
Subsidy rate	0.80	Close to the current all-level fiscal share
Unit indemnity	0.80	Amount that enters the cash book when paid
Rigid outlay	0.70	The part that must be paid at once after the disaster
Risk aversion	2.00	A usual mid-range CRRA value
Ambiguity aversion	0.80	Concavity of the outer exponential
Normal / disaster cash	1.00 / 0.25	Cash buffer before the sale
Normal / disaster harvest	1.20 / 1.00	Saleable quantity
Disaster price intercept / slope	1.00 / 0.35	Weaker secondary-market capacity
Delayed unit value	1.12 / 1.10	Net return after drying and storage
Climate slopes $\lambda_a, \lambda_k, \lambda_R, \lambda_p$	0.06, 0.06, 0.12, 0.20	Quality path; $\xi=0$ returns to the baseline
Older / younger risk aversion	3.0 / 1.5	The two types in Proposition 6
Older / younger disaster cash	0.18 / 0.28	The weighted mean is near the baseline 0.25

For an approximate per-mu interpretation, one normalized unit can be read as about one thousand yuan. In a normal year corn yields about 500 kg per mu at around 2.6 yuan per kg, so output value is about 1,300 to 1,400 yuan. Full-cost insurance costs about 40 yuan per mu, of which the household pays about 8 yuan. Rigid outlays in the disaster window include machine hire, interest, and short-term living costs. This interpretation is only a reading aid. It is not used in estimation. The three climate models stand for lighter, baseline, and heavier distributions that a household can keep in mind. They are not a statistical fit to historical frequencies.

When Figure 6 moves the center of the horizontal axis it keeps the miss-rate gaps among the candidate models. At a center of 0.10, optimal coverage is 0.966 and expected disaster-state

distress sales in the baseline model are 0.0698. At a center of 0.50, coverage is 0.356 and distress sales are 0.4084. Figure 7 sets the three models' miss probabilities to a common 0.10 or 0.50 in order to isolate the subsidy interaction, so coverage at the same horizontal value need not match Figure 6. In the climate analysis, coverage is 0.562, distress sales 0.223, and the price 0.922 at $\xi = 0$; coverage is 0.214, distress sales 0.515, and the price 0.747 at $\xi = 0.80$. At the same $Y_B = 1.004$, quality-path coverage is 0.214 and quantity-path coverage is 0.494. In the aging analysis, at $\omega = 0.336$ and $\xi = 0$ aggregate coverage is 0.616 and regional distress sales are 0.259; at $\xi = 0.40$ coverage is 0.409 and distress sales are 0.333. The other scans are recomputed under the same grid rule.

Appendix C Price anchors and sources

Figure 1 keeps only anchors that can be pointed to a public source. Quotes in yuan per tonne are divided by 2,000 to convert to yuan per jin.

Table C1 Price anchors used in Figure 1

Date	Quality	Yuan/jin	Source note
early Aug	Grade-two spot	1.202	National grade-two corn about 2,403 yuan/tonne; Beijing Grain Reserve Bureau, 2025 Q3 corn market note
late Aug	High-quality / deep-processing	1.195	North China deep-processing mean about 2,390 yuan/tonne; same source
late Sep	High-quality / deep-processing	1.145	Jiaozuo deep-processing about 2,290 yuan/tonne; Zhengzhou Grain Reserve Bureau September note
late Sep	Rough / damaged	1.075	Southern Henan rough grain about 1.05 to 1.10 yuan/jin; Feed Industry Information Net
early Oct	High-quality / deep-processing	1.078	Henan deep-processing 2,140 to 2,170 yuan/tonne after National Day, midpoint 2,155; Beijing Grain Reserve Bureau Q3 report
early Oct	Wet grain	0.630	Zhoukou wet grain with 30 to 35 percent moisture about 1,260 yuan/tonne; same source
mid Oct	Damaged grain	0.630	Damaged grain still near 0.6 yuan/jin; same source
late Oct	High-quality / wet grain	1.102 / 0.78	A local rebound; Huang-Huai wet grain below 0.8 yuan/jin; industry monitoring

References

1. Bernanke, B. S., Gertler, M., and Gilchrist, S. (1999). The financial accelerator in a quantitative business cycle framework. In J. B. Taylor and M. Woodford (eds.), *Handbook of Macroeconomics*, Vol. 1C, ch. 21, pp. 1341-1393. Elsevier.
2. Bryan, G. (2019). Ambiguity aversion decreases the impact of partial insurance: Evidence from African farmers. *Journal of the European Economic Association*, 17(5), 1428-1469.
3. Clarke, D. J. (2016). A theory of rational demand for index insurance. *American Economic Journal: Microeconomics*, 8(1), 283-306.
4. Cole, S., Giné, X., Tobacman, J., Topalova, P., Townsend, R., and Vickery, J. (2013). Barriers to household risk management: Evidence from India. *American Economic Journal: Applied Economics*, 5(1), 104-135.
5. Dell, M., Jones, B. F., and Olken, B. A. (2014). What do we learn from the weather? The new climate-economy literature. *Journal of Economic Literature*, 52(3), 740-798.
6. Elabed, G., and Carter, M. R. (2015). Compound-risk aversion, ambiguity and the willingness to pay for microinsurance. *Journal of Economic Behavior and Organization*, 118, 150-166.
7. Ellsberg, D. (1961). Risk, ambiguity, and the Savage axioms. *Quarterly Journal of Economics*, 75(4), 643-669.
8. Gilboa, I., and Schmeidler, D. (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2), 141-153.
9. Kiyotaki, N., and Moore, J. (1997). Credit cycles. *Journal of Political Economy*, 105(2), 211-248.
10. Klibanoff, P., Marinacci, M., and Mukerji, S. (2005). A smooth model of decision making under ambiguity. *Econometrica*, 73(6), 1849-1892.
11. Lesk, C., Rowhani, P., and Ramankutty, N. (2016). Influence of extreme weather disasters on global crop production. *Nature*, 529, 84-87.
12. Lorenzoni, G. (2008). Inefficient credit booms. *Review of Economic Studies*, 75(3), 809-833.
13. Shleifer, A., and Vishny, R. W. (1992). Liquidation values and debt capacity: A market equilibrium approach. *Journal of Finance*, 47(4), 1343-1366.
14. Shleifer, A., and Vishny, R. W. (2011). Fire sales in finance and macroeconomics. *Journal of Economic Perspectives*, 25(1), 29-48.
15. Tuo, G., and Li, J. (2003). Agricultural insurance experiments in China: Achievements, tensions, and a way forward. *Journal of Financial Research* (《金融研究》), (9), 88-98.
16. Yi, F., Lu, Y., and Wang, K. (2022). Small payouts after large losses, large payouts after small losses: The paradox of agricultural production risk and indemnity under a lump-sum premium subsidy, with evidence from policy corn insurance. *Chinese Rural Economy* (《中国农村经济》), (3), 128-144.
17. Yi, F., Wang, S., and Ding, Y. (2025). Premium rebates in agricultural insurance as compensation for an ambiguity premium: A theory and experimental evidence. *China Economic Quarterlyly* (《经济学（季刊）》), 25(6), 1660-1678.
18. Zhao, Y., Chen, S., Li, X., Lou, W., and Zhang, Y. (2025). Basis risk in weather index insurance products under alternative meteorological data. *Insurance Studies* (《保险研究》), (1), 57-67.

Acknowledgement

The question first took shape in 2023, when I went to Dubai for a youth program on sustainable development and climate action and joined a Model United Nations debate on climate change. That is when I started to pay attention to climate change and extreme weather. I later read William Nordhaus's *The Climate Casino*. The book treats greenhouse-gas emissions as a bet placed inside the climate system. It explains where climate change comes from, what it does, and how it might be slowed. What stayed with me were the externality, uncertainty, and irreversibility behind climate change, and I began to see climate risk as an economics problem. With that lens, I became more alert to climate effects around me. During the 2025 National Day holiday, I went back with my parents to our hometown in Hengshui, Hebei. It is a major grain-producing area, and harvest time coincided with days of rain. The fields were standing in deep water, the roads were muddy, and the corn was still in the ground. Relatives kept asking the same things: Have you harvested? How much did you sell? How much did you lose to mold? Insurance paid almost nothing, and the corn price had fallen again in recent days. I watched smallholder farmers suffer heavy losses from an extreme-weather disaster. Their worry and their sense that there was little they could do made me realize that climate change is not far away. It is a production and livelihood problem facing people here and now.

That is what made me ask why agricultural insurance paid so little after the disaster, and why the grain price fell when almost no one was buying. This seemed at odds with insurance theory and the supply-demand rule in my economics textbook, and I wanted to know why. After I returned to Beijing, the news showed similar cases in Henan Province and elsewhere. Books such as *Research on Agricultural Insurance in China* made the country's setting clearer: China is a large country with many small farms and an aging agricultural workforce, which may mean stronger risk aversion. I took these questions to Professor Rujia Lan, a finance professor at Beijing Normal University. She pointed me to the research literature and helped me see that smallholder farmers' insurance choices are an important issue, especially when climate disasters are becoming more frequent. After several rounds of discussion, we settled on this topic.

The paper relies on theory and numerical work, and the derivations were the hardest part for me. From the setup to the details of the proofs, no step was easy. During the research process, I also received guidance from Ms. Lu Yu, my economics teacher at the Experimental High School Attached to Beijing Normal University. Both teachers were patient. When I got stuck, they walked me through the argument step by step. Professor Lan also suggested adding aging and regional aggregation, which became Propositions 5 and 6. That suggestion deepened the paper and produced findings I would not have reached on my own. Throughout the whole

project, both teachers gave a great deal of their time and offered specific advice. I am grateful for their patience and for guidance given without holding anything back.

The division of work was clear. (1) The research question—insufficient agricultural-insurance demand under extreme climate disasters—came from what I saw when I went home. Professor Lan supervised the work and helped extend it from protection failure and distress sales to aging and regional aggregation. (2) I assembled the literature review. Professor Lan helped me tighten the logic. (3) I built and derived the theoretical model under the guidance of Professor Lan and Ms. Yu. (4) I completed the numerical work in Python with Professor Lan's help. (5) I wrote the paper. Professor Lan and Ms. Yu helped improve the results, structure, and wording.

During that difficult stretch, I began to apply textbook insurance theory and the supply-demand rule more carefully, and I stopped seeing agriculture only through books. I gained a closer sense of what Chinese farming looks like under climate change. The older people in the Hengshui courtyard showed me that how aging rural households sell grain and seek protection is both an economic question and a matter of livelihood for many people. If this work can do a little to improve agricultural protection as extreme weather becomes more frequent, it will not have wasted the worry of those neighbors standing in the rain.

I declare that I did not use AI to write the body of the paper, to invent references, or to fabricate numerical results. There is no undeclared use of AI, and no secret, sensitive, or unauthorized data were uploaded. The derivations, the numerical work, and the writing were completed under my teachers' supervision.