

WINPEC Working Paper Series No. E2616
October 2026

Underinsured by Precision: Weather Forecasts, Contract Timing, and Risk Sharing under Liquidity Constraints

Ryo Takahashi

Waseda INstitute of Political EConomy
Waseda University
Tokyo, Japan

Underinsured by Precision: Weather Forecasts, Contract Timing, and Risk Sharing under Liquidity Constraints*

Ryo Takahashi†

September 2026

Abstract

Climate forecasts and index insurance are often viewed as complementary tools for managing climate risk, especially in developing countries. I show how this complementarity can fail when climate information changes the terms on which insurance is available. When that information enters insurance pricing, adverse news raises actuarially fair premiums, exposing insured households to premium risk and liquidity-constrained households to underinsurance. A model quantified with data from Ethiopia's Index-Based Livestock Insurance program shows that, under an illustrative high-precision forecast, 54.1 percent of households cannot finance a standardized coverage package, while only 1.5 percent cannot finance one tropical livestock unit. With adaptation and exposure adjustment disabled, the repricing loss equals about half the value of full pooled insurance. Forecasts also improve adaptation, creating a trade-off between decision gains and lost risk sharing. Earlier contracting can preserve pooling while retaining the decision value of information, making adaptation capacity and contract timing jointly central to the value of better forecasts.

Keywords: Value of information; weather forecasts; index-based insurance; liquidity constraints; Hirshleifer effect.
JEL codes: D83, G22, O13, Q54.

*I thank Christopher Costello, Kathy Baylis, Robert Heilmayr, Olivier Deschenes, Andrew Plantinga, Claudia Kelsall, and Keisaku Higashida for helpful comments, suggestions, and supports. I am also grateful to emLab, especially Michaela Clemence, Jennifer Bone, and Erin O'Reilly, for organizing the retreat at Channel Islands National Park, where discussions helped inspire this research.

†Bren School of Environmental Science & Management, University of California, Santa Barbara; Faculty of Political Science and Economics, Waseda University.

1 Introduction

Climate forecasts and index insurance offer two complementary approaches to managing risk in a changing climate. Forecasts help households adjust production, investment, and consumption before losses occur (Rosenzweig and Udry, 2014; Burlig et al., 2024), while insurance transfers part of the remaining risk (Cole and Xiong, 2017; Carter et al., 2017). These tools are especially important in developing countries, where exposure to climate shocks is substantial and formal risk-management opportunities are limited (Morton, 2007; Cohn et al., 2017). Their importance is likely to grow as climate change increases the frequency or severity of extreme weather (Grant et al., 2025). Yet the information that helps households prepare can also change the price at which insurance is sold.

This paper asks how better public forecasts interact with insurance markets to shape coverage and welfare. I focus on a key institutional margin: whether new climate information changes the terms on which insurance is available, and how the timing of contracting affects the resulting risk sharing. I study these mechanisms in a model with heterogeneous household liquidity and adaptation opportunities, and quantify their importance using data from Ethiopia's Index-Based Livestock Insurance program.

The central result is that improved information can make adequate coverage harder to finance precisely when expected losses are highest. This creates an important distinction between insurance participation and insurance adequacy: take-up can remain high even as effective coverage deteriorates. Households need not leave the insurance market to lose protection; they can remain insured while carrying substantially less coverage after adverse news.

The mechanism is straightforward. Public forecasts change the loss probabilities on which competitive insurers price coverage. Adverse news raises premiums, while favorable news lowers them. Insurance purchased after the forecast can smooth residual weather risk, but it cannot pool the risk already revealed by the forecast. Even a household that remains fully insured therefore faces a lottery over premiums. Liquidity constraints create an additional problem: some households cannot finance enough coverage after adverse news to maintain the same level of protection. Contract divisibility then determines whether this adjustment appears primarily as reduced coverage or as complete market exit.

The welfare consequences depend on whether households can also use the forecast to adapt, for example by moving livestock to alternative grazing areas or adjusting crop choice and planting timing. Without adaptation, informative post-signal pricing reduces the value of insurance even for households whose liquidity constraint does not bind; liquidity con-

straints add further welfare losses by leaving more risk uninsured after adverse news. With adaptation, however, information can also improve decisions and reduce expected losses. Unequal access to effective adaptation can therefore widen the welfare gap between households that can act on a forecast and those that can neither adapt nor maintain adequate coverage. The aggregate welfare effect depends on the balance between these decision and risk-sharing channels, which also makes the timing of insurance contracting central to the analysis.

To gauge the quantitative importance of these mechanisms, I use household data from the Index-Based Livestock Insurance (IBLI) program in southern Ethiopia (Takahashi et al., 2016, 2019). Assets, savings, borrowing, and an incentivized risk-preference exercise discipline heterogeneity in households' ability to finance insurance and bear risk, while administrative purchase records and randomized premium discounts provide supporting evidence on contract sizes and price sensitivity. The application is a counterfactual pricing exercise rather than an estimate of historical forecast-driven repricing by IBLI. Forecast precision is therefore varied rather than estimated, because the available inputs do not include a matched historical forecast–forage-index series.

The quantitative results reinforce the distinction between insurance participation and insurance adequacy. In an illustrative high-precision scenario, 54.1% of households cannot finance the standardized coverage package after adverse news, whereas only 1.5% cannot finance a contract covering one tropical livestock unit (TLU). With adaptation and exposure adjustment disabled, the loss from repricing is equivalent to about 50% of the risk-sharing value of full pooled insurance. A separate specification lets households choose coverage against their own exposure and sell livestock to finance premiums. At the reference sale-recovery rate, protection again shifts toward favorable-signal states and the static certainty-equivalent loss is small relative to mean herd value, although it remains substantial relative to the pooled-insurance value attainable in that specification. Because sales also reduce the exposure base in this specification, they create an additional destocking margin that I separate from the pure insurance channel in the robustness analysis.

The insurance-side mechanism also becomes more consequential when underlying climate risk is higher. In the high-precision scenario, doubling the baseline probability of the insured event from 15 to 30 percent raises the probability of adverse news from 22 to 34 percent, increases the share of households unable to finance the standardized coverage package from 54.1 to 63.6 percent, and increases the population certainty-equivalent loss with adaptation disabled from 0.95 to 1.70 percent of mean herd value. This sensitivity illustrates how higher underlying climate risk can amplify the insurance consequences of forecast-contingent pricing.

These insurance-side losses do not by themselves determine the welfare value of better forecasts, because households may also use the information to reduce expected losses. Under the reference scenario, which assigns the adaptation technology to the top asset quintile, the divisible-coverage effect is -0.45% of mean herd value in the high-precision case, but alternative assumptions about access and effectiveness can reverse its sign. Because the survey's herd-mobility measures capture observed behavior rather than forecast-induced loss reduction, I treat adaptation opportunities as a sensitivity margin rather than a directly estimated technology. This distinction matters for contract timing: forecasts can retain their decision value even when insurance is arranged earlier, allowing the adaptation benefit of information to be separated from the loss of ex ante risk sharing.

Preserving that risk sharing, however, requires arrangements that remain viable once the forecast becomes public. A fixed spot price needs support across signal states, especially if favorable-signal households stop buying. Holding the observed IBLI premium scale fixed, the illustrative high-precision calculation gives an adverse-signal insurer loss ratio—expected indemnity divided by premium—of 2.05. This accounting motivates treating earlier contracting separately from post-signal premium support. Earlier contracting also creates an implementation problem: if households can cancel after favorable news and buy replacement coverage more cheaply, pooled contracts require a non-refundable prepayment or another commitment device. I characterize that requirement without assuming that households necessarily have fewer resources before the forecast than afterward.

The paper contributes to four literatures. First, it connects the Hirshleifer effect to work on reclassification risk and dynamic insurance. Public information arriving before contracting can reduce gains from risk sharing (Hirshleifer, 1971; Schlee, 2001), while related work studies predictive information, selection, and pooling (Jensen et al., 2018; Solomon, 2026). When new information changes insurance terms between contracting dates, households face a form of reclassification risk analogous to that studied in dynamic insurance (Cochrane, 1995; Pauly et al., 1995; Handel et al., 2015). Dynamic contracts can protect future terms but require commitment when buyers can lapse after favorable news (Hendel and Lizzeri, 2003; Finkelstein et al., 2005). I show how household liquidity and contract divisibility shape the incidence of this loss, and characterize the prepayment needed to sustain pooled coverage when later replacement insurance becomes cheaper.

Second, research on agricultural insurance has studied how seasonal climate information affects insurance markets. Early-season weather information can generate intertemporal adverse selection (Luo et al., 1994), affect insurer returns (Ker and McGowan, 2000), and inform production, contract design, and rate setting (Carriquiry and Osgood, 2012; Yi et al.,

2020). Related work on property insurance shows that more granular climate-risk information can reshape insurance pricing and availability when risk classification differs across insurers (Boomhower et al., 2024). My setting instead considers public information observed by both households and insurers, so the mechanism does not rely on either side having an informational advantage. I focus on how forecast-contingent pricing changes the risk-sharing opportunities available to households and who bears the resulting loss.

Third, research on index insurance demand emphasizes liquidity, trust, basis risk, and payment timing at prevailing prices (Belissa et al., 2019; Clarke, 2016; McIntosh et al., 2019). Delaying premium payment can substantially increase take-up (Casaburi and Willis, 2018; Liu et al., 2020), while panel evidence from the Borena program traces demand across sales periods (Takahashi et al., 2020). Insurance and state-contingent credit can also relax investment constraints and improve resilience (Carter et al., 2016; Hill et al., 2019; Lane, 2024; Boucher et al., 2024). I shift attention from whether liquidity-constrained households purchase insurance at a given price to whether they can maintain adequate protection when public information changes that price. This distinction introduces an important quantity margin: take-up can remain high even as coverage deteriorates after adverse news. Insurance participation alone can therefore understate the loss of protection when contracts are divisible.

Fourth, research on climate information values forecasts through the production, labor-allocation, investment, adaptation, and mortality outcomes they improve (Rosenzweig and Udry, 2014; Burlig et al., 2024; Shrader et al., 2026), and documents substantial differences in forecast quality across settings (Linsenmeier and Shrader, 2023). Related work shows that insurance prices can themselves provide incentives and information for private self-protection (Boomhower et al., 2023). Among pastoralists in southern Ethiopia and northern Kenya, forecasts reach only a subset of households (Luseno et al., 2003), and recipients revise rainfall expectations in response (Lybbert et al., 2007). This paper identifies a complementary market-mediated channel: the same information that improves household decisions can also weaken ex ante risk sharing when insurance is priced after the forecast. The value of better information therefore depends not only on predictive accuracy and households' ability to act on it, but also on the contracting environment through which residual risk is shared.

Section 2 sets out the model and its benchmarks. Section 3 establishes the coverage and welfare results, and Section 4 presents the IBLI application. Section 5 examines contract timing and implementation. Section 6 evaluates alternative assumptions, and Section 7 concludes.

2 Model and Benchmarks

The model separates the information used to price insurance from the resources households can use to buy it. Insurers are competitive and risk neutral. Households are risk averse, differ in liquidity, and may have access to an adaptation technology. I first define their choices and the pooled benchmark, then study how public information changes coverage and welfare. The benchmark loss is common across households; the quantitative application also considers household-specific exposure.

2.1 Environment and timing

Let $D \in \{0, 1\}$ indicate the insured loss event. Before information about the current season arrives, households and insurers share the prior probability $p = \mathbb{P}(D = 1) \in (0, 1)$, which represents the long-run base rate of the insured event. A public seasonal forecast then takes values $s \in \{B, G\}$, where B denotes adverse news and G favorable news, and updates the probability of the event for the current insurance period. The timing of the spot benchmark reflects the sequencing of seasonal climate information and insurance sales:

$t = 0$	Households and insurers share the prior. No forecast has been observed and no insurance contract is written.
$t = 1$	The forecast arrives. Insurers set conditional prices; households choose adaptation and insurance subject to their liquidity constraints.
$t = 2$	The insured event is realized, indemnities are paid, and households consume.

This sequencing is empirically relevant for seasonal index insurance. In the original IBLI design, contracts were marketed in short windows before fixed coverage seasons running from March to September and from October to February (Vrieling et al., 2016), while regional climate outlooks are issued ahead of the rainy seasons. For the October–December rains, for example, outlooks are typically produced in late August (Walker et al., 2019). The resulting overlap means that information about the coming season may already be available when insurance is priced. The institutional comparison in Section 5 instead opens contracting at $t = 0$, before the forecast, while leaving the information available for subsequent decisions. Earlier contracting, however, requires commitment because favorable news can make replacement coverage cheaper after the signal.

To capture this information structure parsimoniously, I model the forecast as a symmetric

binary signal with precision $\theta \in [1/2, 1]$:

$$\mathbb{P}(B \mid D = 1) = \mathbb{P}(G \mid D = 0) = \theta. \tag{1}$$

Equation (1) defines accuracy for the insured event, not for a generic rainfall category. Bayes' rule gives

$$p_B(\theta) = \frac{\theta p}{\theta p + (1-\theta)(1-p)}, \qquad p_G(\theta) = \frac{(1-\theta)p}{(1-\theta)p + \theta(1-p)}. \tag{2}$$

These assumptions imply signal probabilities $q_B(\theta) = \theta p + (1-\theta)(1-p)$ and $q_G(\theta) = (1-\theta)p + \theta(1-p)$. I write $p_s(\theta)$ for the posterior associated with signal s and, for any signal-contingent quantity X_s, use $\mathbb{E}_s[X_s] = q_B(\theta)X_B + q_G(\theta)X_G$ for its expectation before the forecast. Greater precision raises p_B and lowers p_G, while preserving the prior on average: $\mathbb{E}_s[p_s(\theta)] = q_B p_B + q_G p_G = p$. Thus, precision redistributes perceived risk across signal states without changing the unconditional probability of loss.

The comparative statics in forecast precision hold the underlying event probability p fixed. A change in baseline climate risk is a distinct margin. For any informative signal, $\theta > 1/2$,

$$\frac{\partial q_B(\theta, p)}{\partial p} = 2\theta - 1 > 0,$$

so a higher baseline event probability makes adverse signals more frequent even when forecast technology is unchanged. It also raises the signal-contingent posterior risks and therefore actuarially fair premiums. Section 6.6 varies this margin separately from forecast precision.

Household i has asset wealth a_i and a common non-stochastic endowment $\bar{y}$, giving total resources $R_i = a_i + \bar{y}$. In the stylized model, resources available for up-front risk management are $w_i = a_i + \bar{b}$, with common borrowing capacity $\bar{b} \geq 0$. The application replaces this mapping with partially liquid livestock wealth, realized borrowing, and savings. Preferences over final consumption z are

$$u_i(z) = -\exp(-\gamma_i z), \qquad \gamma_i > 0, \tag{3}$$

where γ_i is household i's coefficient of absolute risk aversion. Constant absolute risk aversion in equation (3) makes non-stochastic wealth additive in the certainty equivalent. This isolates the liquidity channel from wealth effects on risk aversion, although the application allows γ_i to differ across households.

2.2 Adaptation, insurance, and the budget

Without adaptation, the household loses $L > 0$ when $D = 1$. Adaptation effort $e_i \geq 0$ is a reduced-form representation of loss-reducing actions taken in response to the forecast, such as moving livestock to alternative grazing areas or adjusting crop choice and planting timing. It reduces the loss to $h(e_i)$, where $h(0) = L$, $h(e_i) > 0$, $h' < 0$, and $h'' > 0$. Its cost is

$$\psi_i(e_i) = \kappa_i \mathbb{1}\{e_i > 0\} + \chi(e_i), \qquad \chi(0) = 0, \quad \chi' > 0, \quad \chi'' \geq 0. \tag{4}$$

Here $\kappa_i \geq 0$ is household i's fixed cost of activating adaptation, while $\chi(e_i)$ is the variable cost of adaptation effort. The fixed component distinguishes access to scalable adaptation from an initially costly emergency response. I use *inactive adaptation* to mean $e_i = 0$, so that $h(e_i) = L$ and $\psi_i(e_i) = 0$.

Assumption 1 (Lumpy emergency adaptation)**.** *Households with scalable adaptation have $\kappa_i = 0$. Other households face a fixed activation cost $\kappa_i = \kappa > L$.*

For households without scalable adaptation, activation costs more than the largest loss it could prevent. Inactive adaptation provides a transparent endpoint for the modeled emergency response. Section 6 relaxes it by allowing continuous adaptation within the constrained group's budget.

Alongside adaptation, households can use insurance to manage the remaining loss exposure. The basic insurance menu is divisible. Household i buys indemnity $m_i \in [0, L]$, paid if $D = 1$. I write its unit price as $p_s(\theta)$, anticipating the competitive fair-pricing result in Section 2.3 that equates this price to the posterior loss probability. For comparison, a lumpy menu restricts this choice to $m_i \in \{0, M\}$.

Assumption 2 (Coverage of the benchmark loss)**.** *The standardized package has maximum indemnity $M = L$.*

I write $I_i = m_i/M \in \{0, 1\}$ for its purchase indicator and $\pi_s = p_s M$ for its premium. This restriction converts a coverage choice into a participation choice while retaining the same repricing mechanism.

Because both adaptation and insurance must be financed before the loss is realized, they draw on the same up-front budget:

$$\psi_i(e_i) + p_s(\theta) m_i \leq w_i. \tag{5}$$

Given these choices, final consumption is

$$z_i = R_i - \psi_i(e_i) - p_s(\theta)m_i - D\big[h(e_i) - m_i\big]. \tag{6}$$

In equation (6), adaptation expenditure and insurance premiums are paid regardless of whether the loss occurs, while $h(e_i) - m_i$ is the household's remaining loss exposure when $D = 1$. Insurance payouts depend on the insured event rather than on the household's adaptation choice. The theoretical baseline aligns the index event with the household loss event; the quantitative sensitivity analysis introduces basis risk.

2.3 Pricing and household choice

Because the signal is public, the expected payment per unit of indemnity is $p_s(\theta)$. Competition and zero expected profit therefore equate the unit price to that posterior probability, giving the full-package premium

$$\pi_s(\theta) = p_s(\theta)M. \tag{7}$$

The premium in equation (7) is fair after the signal. It need not be constant before the signal. This distinction between conditional fairness and ex-ante pooling is the source of the welfare trade-off.

Lemma 1 (Insurance demand under fair pricing)**.** *Fix $p_s \in (0,1)$ and a feasible adaptation choice with $\psi_i(e_i) \leq w_i$.*

(i) With divisible coverage, the unique optimal indemnity is

$$m_i^* = \min\left\{h(e_i), \frac{w_i - \psi_i(e_i)}{p_s}\right\}.$$

Full coverage of residual loss is optimal when affordable; otherwise, the remaining insurance budget binds.

(ii) With inactive adaptation and the lumpy menu, the household purchases the full package if and only if $\pi_s \leq w_i$.

Proof. See Appendix B.2. □

Lemma 1 follows from a simple insurance comparison. At a fair price, changing coverage leaves conditional expected consumption unchanged. Full coverage makes consumption certain; every other choice adds dispersion around that mean. Liquidity can prevent this

smoothing without changing its desirability. With a divisible menu the household then buys less coverage, whereas a lumpy menu may force it to buy none. At degenerate posteriors of zero or one, strict uniqueness need not hold; the reported coverage formulas use their limiting convention.

Lemma 2 (Substitution between self-protection and insurance)**.** *For a fixed posterior $p_s \in (0,1)$ and effort e_i, unconstrained willingness to pay for full residual-loss coverage in excess of its fair premium $p_s h(e_i)$ is*

$$\rho_i^{\text{res}}(e_i; p_s) = \frac{1}{\gamma_i} \ln\left[1 - p_s + p_s e^{\gamma_i h(e_i)}\right] - p_s h(e_i). \tag{8}$$

It is strictly decreasing in e_i when $h(e_i) > 0$ and $h'(e_i) < 0$.

Proof. See Appendix B.3. □

Lemma 2 separates physical protection from financial protection. Reducing the physical loss lowers the risk premium associated with leaving that loss uninsured. This substitution does not imply that forecasts and insurance must be substitutes overall: the forecast also changes prices and the circumstances in which households choose either response. Lemma 3 allows both posterior probabilities and adaptation to change with forecast precision, bringing this substitution into the value-of-information comparison.

2.4 Welfare and the pooled benchmark

Conditional and ex-ante welfare are

$$\begin{aligned} V_i(s;\theta) &= \max_{e_i \geq 0,\, m_i \in [0,L]} \mathbb{E}[u_i(z_i) \mid s] \quad \text{s.t. (5)}, \\ W_i(\theta) &= q_B(\theta) V_i(B;\theta) + q_G(\theta) V_i(G;\theta). \end{aligned} \tag{9}$$

The lumpy comparison changes the coverage menu in equation (9) to $\{0, M\}$. Define $\text{CE}_i(\theta)$ by $u_i(\text{CE}_i(\theta)) = W_i(\theta)$, and let $\Delta\text{CE}_i(\theta) = \text{CE}_i(\theta) - \text{CE}_i(1/2)$. A negative change is a welfare loss measured in consumption units.

The uninformative benchmark, $\theta = 1/2$, has $p_B = p_G = p$ and premium pL. A household with inactive adaptation and $w_i \geq pL$ fully insures and consumes $R_i - pL$ with certainty. A household with $w_i < pL$ instead buys partial coverage on the divisible menu or no coverage on the lumpy menu. Those are different starting allocations and must not be interchanged when evaluating information.

Under informative pricing, a household with inactive adaptation that maintains full coverage pays $\pi_B(\theta)$ or $\pi_G(\theta)$ according to the signal. I call this ex-ante uncertainty in premium outlays the *premium lottery*; it remains even though the realized physical loss is fully insured.

For a household whose budget permits optimal adaptation and full residual-loss coverage, define

$$e^*(\mu) \in \arg\min_{e\geq 0}\{\psi(e) + \mu h(e)\}, \qquad g(\mu) = \min_{e\geq 0}\{\psi(e) + \mu h(e)\}. \tag{10}$$

Here μ denotes the posterior loss probability. Equation (10) minimizes adaptation expenditure plus the fair premium on the remaining loss. Its pooled certainty equivalent is $R_i - g(p)$, and its post-signal consumption is $R_i - g(p_s)$. Although g is increasing, it is concave: adaptation limits the increase in minimized expenditure as expected losses rise.

Remark 1 (Scope of the welfare accounting). Risk-neutral, zero-profit insurers isolate the household risk-sharing channel. The model excludes forecast-induced changes in insurer capital requirements, reinsurance prices, administration, and reserve management. Lower supply-side costs would add a benefit not measured here; higher costs would work in the opposite direction. The quantitative aggregates are therefore household-side money-metric comparisons, not a complete accounting of social surplus.

3 Information, Liquidity, and Welfare

The analysis proceeds in three steps. I first identify when adverse-signal prices exceed available liquidity. I then compare informative pricing with pooled insurance, distinguishing households that can adapt from those that cannot. Finally, I separate the distributional effect of liquidity from the additional effect of unequal adaptation opportunities. Proofs are collected in Appendix B.

3.1 When does precision reduce coverage?

Proposition 1 (Dynamic pricing and the liquidity threshold)**.** *The adverse-signal premium $\pi_B(\theta)$ is strictly increasing in forecast precision. Consequently, the liquidity required to finance the full benchmark package, $w_{\min}(\theta) = \pi_B(\theta)$, is strictly increasing in θ.*

Proof. See Appendix B.4. □

Proposition 1 concerns the price of a fixed amount of protection. A more accurate adverse signal implies a higher conditional probability of payment, which a competitive

insurer incorporates into its premium. For a household with fixed liquid resources and no adaptation, the resulting coverage choice is $\min\{L, w_i/p_B(\theta)\}$. Coverage remains at L until the budget binds and declines thereafter.

Proposition 2 (Loss of access to the benchmark package)**.** *Under the stylized mapping $w_i = a_i + \bar{b}$, the asset threshold for financing the benchmark package after an adverse signal is*

$$a^*(\theta) = \pi_B(\theta) - \bar{b}.$$

Households with $a_i < a^*(\theta)$ cannot finance the package, and $a^*(\theta)$ increases strictly with precision.

Proof. See Appendix B.4. □

Proposition 2 expresses the same affordability condition in asset units. Crossing the threshold does not imply leaving a divisible insurance market: positive liquidity still finances positive coverage. Complete exit follows only when the menu rules out a smaller contract. The distinction is important empirically, because participation and adequate protection measure different outcomes.

Remark 2 (Endogenous liquidity). The baseline holds borrowing capacity and liquidation values fixed across signals. Adverse news could also tighten credit or lower livestock sale prices. In pastoral livestock markets, sale prices are sensitive to rainfall and drought conditions (Barrett et al., 2003), so widespread destocking following an adverse forecast could reduce the value recovered from livestock sales. The sensitivity analysis therefore varies the liquidity mapping and, in the own-exposure exercise, allows the sale-recovery rate to fall specifically after an adverse forecast.

3.2 Can adaptation offset the premium lottery?

An unconstrained household can fully insure its remaining physical loss, but still pays different amounts after favorable and adverse news. Its welfare is $W_U(\theta) = \mathbb{E}_s[u_i(R_i - g(p_s))]$, with g defined in equation (10). More precise information improves the targeting of adaptation, while making the resulting expenditure more variable.

Lemma 3 (Value of information for an unconstrained household)**.** *Fix an unconstrained household i able to finance optimal adaptation and full residual-loss coverage at every posterior under consideration. Let $f(\mu) = u_i(R_i - g(\mu))$, so that $W_U(\theta) = \mathbb{E}_s[f(p_s(\theta))]$.*

(i) If adaptation is inactive throughout the range, $g(\mu) = \mu L$ and $W_U(\theta)$ decreases strictly with precision.

(ii) Fix $\theta_2 \in (1/2, 1)$. Suppose the adaptation optimum is unique and interior and g is twice differentiable throughout $[p_G(\theta_2), p_B(\theta_2)]$. Then W_U is nondecreasing on $[1/2, \theta_2]$ if

$$\gamma_i \leq \frac{-g''(\mu)}{g'(\mu)^2} = \frac{h'(e^*(\mu))^2}{h(e^*(\mu))^2[\psi''(e^*(\mu)) + \mu h''(e^*(\mu))]} \tag{11}$$

for every $\mu \in [p_G(\theta_2), p_B(\theta_2)]$.

(iii) More generally, $W_U(\theta) \geq W_U(1/2)$ if and only if

$$g(p) - \mathbb{E}_s[g(p_s(\theta))] \geq \rho_i^{\exp}(\theta), \tag{12}$$

where $\rho_i^{\exp}(\theta) \geq 0$ is the risk premium on the expenditure lottery, defined by

$$\mathbb{E}_s u_i(R_i - g(p_s)) = u_i\big(R_i - \mathbb{E}_s[g(p_s)] - \rho_i^{\exp}(\theta)\big).$$

Proof. See Appendix B.5. □

Lemma 3 identifies the two opposing effects. The left side of equation (12) is the resource saving from choosing adaptation after observing the signal. The right side is the compensation required for variable risk-management expenditure. Fair pricing does not eliminate the second term. Even households that retain full residual-loss coverage may therefore lose from more information.

Condition (11) gives a sufficient curvature condition under which the decision gain dominates along the precision path. At a smooth interior optimum, $\psi'(e^*) + \mu h'(e^*) = 0$ and the envelope theorem gives $g'(\mu) = h(e^*(\mu))$. Differentiating the first-order condition and then g' yields

$$\frac{\mathrm{d}e^*(\mu)}{\mathrm{d}\mu} = \frac{-h'(e^*(\mu))}{\psi''(e^*(\mu)) + \mu h''(e^*(\mu))}, \qquad g''(\mu) = \frac{-h'(e^*(\mu))^2}{\psi''(e^*(\mu)) + \mu h''(e^*(\mu))}.$$

These identities give the second equality in condition (11). Greater precision is valuable when adjusting adaptation saves enough to offset expenditure risk, which is why access alone does not determine the welfare sign.

Remark 3 (Activation with exponential loss reduction). For $\kappa_i = 0$, $\chi(e) = e$, and $h(e) =$

$Le^{-\beta e}$ with $\beta > 0$, an unconstrained household at posterior $\mu > 0$ chooses

$$e^*(\mu) = \max\left\{0, \frac{\ln(\mu\beta L)}{\beta}\right\}, \qquad g(\mu) = \begin{cases} \mu L, & \mu\beta L \le 1, \\ [1 + \ln(\mu\beta L)]/\beta, & \mu\beta L > 1. \end{cases}$$

On the interior, $g'(\mu) = 1/(\beta\mu)$ and $g''(\mu) = -1/(\beta\mu^2)$, so condition (11) simplifies to $\gamma_i \le \beta$. The interior restriction remains essential: adaptation is active only when $\mu\beta L > 1$.

3.3 Welfare of households that become constrained

Consider a household with inactive adaptation and $w_i \in [pL, L)$. It can afford full pooled insurance but not full coverage after sufficiently adverse information. Since $\pi_B(\theta)$ increases continuously and strictly from pL at $\theta = 1/2$ to L at $\theta = 1$, there is a unique $\bar{\theta}_i \in [1/2, 1)$ such that $\pi_B(\bar{\theta}_i) = w_i$. Favorable-signal coverage remains affordable at every informative precision because $\pi_G(\theta) < pL \le w_i$. The lumpy menu makes the consequences of crossing $\bar{\theta}_i$ especially transparent.

Proposition 3 (Welfare of the marginally insured)**.** *Under Assumptions 1–2, consider a household with $w_i \in [pL, L)$ and inactive adaptation facing the lumpy menu $m_i \in \{0, L\}$. It has the following welfare comparisons:*

(i) On $(1/2, \bar{\theta}_i]$, it remains insured after both signals, but welfare decreases strictly with precision.

(ii) Moving just above $\bar{\theta}_i$ causes a discrete loss:

$$\lim_{\theta \downarrow \bar{\theta}_i} W_i(\theta) = W_i(\bar{\theta}_i) - q_B(\bar{\theta}_i) J_i, \tag{13}$$

where

$$J_i = u_i(R_i - w_i) - \left[\left(1 - \frac{w_i}{L}\right) u_i(R_i) + \frac{w_i}{L} u_i(R_i - L)\right] > 0.$$

(iii) If $p \le 1/2$, welfare also decreases strictly on $(\bar{\theta}_i, 1)$.

(iv) For any prior $p \in (0, 1)$ and any $\theta > 1/2$, welfare is below the pooled benchmark: $W_i(\theta) < W_i(1/2)$.

Proof. See Appendix B.6. □

Proposition 3 separates the loss from information from the loss of the full package. Welfare already declines while the household remains insured, because its premium becomes uncertain. Equation (13) adds a discontinuity when an arbitrarily small price increase makes the entire package unaffordable. That discontinuity is a property of the menu, not a necessary implication of better forecasting.

Beyond $\bar{\theta}_i$, the household is insured only after favorable news. Greater precision reduces the ex-ante premium outlay on that insured branch, $q_G\pi_G = (1-\theta)pL$, while increasing the expected loss borne without insurance on the adverse branch, $q_B p_B L = \theta pL$. Expected consumption remains fixed, but more of the loss is borne in the uninsured branch; as $\theta \to 1$, consumption approaches the uninsured lottery $R_i - DL$. Under constant absolute risk aversion (CARA), $p \leq 1/2$ is sufficient for this reallocation to reduce welfare at every precision above the threshold. This sufficient restriction is satisfied by the calibration; the global comparison in part (iv) holds for every prior.

The global ranking in part (iv) has a simpler explanation than the local comparative statics. Fair pricing leaves expected consumption at $R_i - pL$. Pooled full coverage delivers that amount with certainty, whereas informative spot pricing produces a nondegenerate lottery with the same mean. The next result shows why allowing smaller contracts does not remove this comparison.

Corollary 1 (Divisible coverage)**.** *Keep adaptation inactive and allow $m_s \in [0, L]$ subject to $p_s m_s \leq w_i$. If the household can fully insure at $\theta = 1/2$, its ex-ante welfare at any informative precision is weakly below the pooled benchmark, and strictly below it whenever consumption is nondegenerate. Divisibility removes the discrete exit in Proposition 3, but not this global ranking.*

Proof. For every signal and coverage choice, conditional expected consumption is $R_i - p_s m_s - p_s(L - m_s) = R_i - p_s L$. Averaging across signals gives $R_i - pL$. Full pooled insurance makes consumption certain at this mean. Any nondegenerate consumption lottery with the same mean is strictly worse by Jensen's inequality. □

The corollary does not use a liquidity constraint to establish the welfare sign. It recovers the pure Hirshleifer comparison. Liquidity determines the extra risk left uninsured and its incidence across households. Unlike part (iii) of the lumpy result, the corollary does not claim monotonic welfare at every adjacent precision level.

Remark 4 (Households below the pooled-price threshold). A household with $w_i < pL$ begins uninsured under the lumpy menu. Once favorable-signal coverage becomes affordable, it can replace a conditional lottery with its mean and gain weakly from information. Under divisible

insurance it already has partial pooled coverage, so this argument no longer establishes a gain. The quantitative analysis reports both cases rather than imposing a common ranking.

Remark 5 (Divisible adaptation). Allowing adaptation leaves the affordability test for a fixed no-adaptation package unchanged, because $e_i = 0$ remains feasible. However, it can change optimal protection and welfare. A sufficiently effective adaptation technology can offset the loss from reduced insurance. The quantitative analysis therefore varies both adaptation opportunities and effectiveness.

3.4 Who bears the loss?

To isolate the role of liquidity, I first compare households with the same loss exposure and risk preferences, assuming inactive adaptation. Differences in welfare within this comparison therefore reflect liquidity alone. The quantitative application later allows liquidity, risk preferences, and adaptation access to differ across households.

Proposition 4 (Information-induced inequality)**.** *Under Assumptions 1–2, consider two households facing fair prices and the divisible menu with the same loss L. Household U is unconstrained, with sufficient resources to finance its optimum and full residual-loss coverage throughout the posterior range under consideration. Household C is liquidity constrained and has inactive adaptation, with $w_C \in [pL, L)$. Its affordability threshold $\bar{\theta}_C$ is defined by $\pi_B(\bar{\theta}_C) = w_C$.*

For each household $j \in \{U, C\}$, define the information-induced welfare loss relative to the uninformative benchmark as

$$\Lambda_j(\theta) = \mathrm{CE}_j(1/2) - \mathrm{CE}_j(\theta).$$

Thus, $\Lambda_j(\theta) > 0$ indicates a welfare loss from informative pricing relative to pooled pricing.

(i) If adaptation is inactive for both, $\gamma_U = \gamma_C$, and C can finance only $m_B < L$, then

$$\Lambda_C(\theta) > \Lambda_U(\theta) > 0, \qquad \theta \in (\bar{\theta}_C, 1). \tag{14}$$

(ii) If U has $\kappa_U = 0$ and satisfies the interiority, differentiability, and curvature conditions of Lemma 3(ii) for some $\theta_2 \in (1/2, 1)$, while adaptation is inactive for C, then

$$\mathrm{CE}_U(\theta) - \mathrm{CE}_C(\theta) > \mathrm{CE}_U(1/2) - \mathrm{CE}_C(1/2) \tag{15}$$

for every $\theta \in (1/2, \theta_2]$. This comparison does not require equal risk-aversion coefficients.

Proof. See Appendix B.7. □

Equation (14) isolates the additional burden of constrained insurance. Both households face the premium lottery, but C also retains weather risk after an adverse signal. Its net consumption lottery is a mean-preserving spread of U's. Holding risk aversion fixed isolates liquidity in this comparison. With unrestricted preference differences, the lottery ordering need not rank household-specific certainty-equivalent losses.

Equation (15) adds unequal adaptation opportunities. When U's decision gains offset its premium risk while C continues to lose, their money-metric gap widens relative to pooled insurance. This does not imply that every wealthy household gains, or that the gap increases at every step in precision. Nor does it establish that the poorest households always gain: that conclusion depended on starting uninsured under a lumpy menu. The empirical comparisons below retain these distinctions.

4 Quantitative Application: Livestock Index Insurance in Ethiopia

The quantitative application maps the theoretical mechanism to Ethiopia's Index-Based Livestock Insurance (IBLI) program. The setting is useful for the analysis because households differ substantially in livestock wealth and liquid resources, while insurance can be purchased in different quantities. I use household data to discipline the distribution of resources and risk preferences, and then compare pooled and forecast-contingent pricing in counterfactual exercises. Administrative purchase records and randomized premium discounts provide supporting evidence on observed contract sizes and price responsiveness, rather than identifying historical forecast-driven repricing by IBLI. Forecast skill, adaptation effectiveness, and liquidation capacity remain scenario inputs.

4.1 Institutional background and data

IBLI was introduced in southern Ethiopia in 2012 to protect pastoral households against drought-related livestock losses (Takahashi et al., 2016). Rather than conditioning indemnities on an individual household's realized livestock losses, the product uses a satellite-based vegetation index designed to capture forage scarcity associated with drought. Payouts are triggered when the index falls below a threshold based on the historical distribution of local vegetation conditions. Households choose how much livestock to insure, and premiums

depend on the insured herd value and location-specific premium rates. These features make IBLI particularly useful for distinguishing access to some insurance from the ability to finance an economically meaningful level of protection.

The household survey began with a pre-intervention baseline in March–April 2012, before the first IBLI sales period in August–September 2012 (Takahashi et al., 2016). The calibration sample used here contains 506 households with positive livestock holdings. The survey records livestock holdings, cash savings, realized borrowing, and choices in an incentivized six-choice risk experiment. Livestock holdings are aggregated into tropical livestock units, and the risk-experiment choices are mapped ordinally into a CARA risk-aversion index rather than treated as precise estimates of household-specific utility curvature. Table 1 reports the resulting sample characteristics.

The original IBLI evaluation also generated exogenous variation in insurance prices through randomly assigned premium-discount coupons (Takahashi et al., 2016). Within study sites, households were randomly offered discounts on coverage for the first 15 TLU insured, generating substantial variation in the effective price of IBLI. I do not use these discounts to estimate the forecast-pricing counterfactual. Instead, the randomized variation provides supporting evidence on whether actual insurance demand responds to price; Appendix Table A19 reports the corresponding take-up and insured-quantity regressions. Administrative purchase records are also used to benchmark the quantities of insurance households purchased in practice.

Let H_i denote household i's herd holdings in TLU, and $\bar{H}$ is the survey-weighted mean. For the baseline calibration, I use a common valuation of 6,000 Ethiopian birr (approximately 340 USD at the 2012 exchange rate) per TLU to express borrowing and savings in mean-herd-value units. Because the same valuation applies to both household and mean herd holdings, normalized livestock wealth is simply

$$a_i = \frac{H_i}{\bar{H}}.$$

Similarly, let b_i and σ_i denote realized borrowing and cash savings, respectively, expressed in the same units. The baseline measure of resources available to finance insurance premiums is

$$w_i = \alpha a_i + b_i + \sigma_i, \qquad \alpha = 0.5. \tag{16}$$

The parameter α captures the assumption that only part of livestock wealth can be converted into funds for immediate premium payments. In the baseline, one half of livestock wealth is treated as liquid, while realized borrowing and cash savings are fully available. This is

a measurement model rather than a direct measure of household liquidity: the survey does not observe unused credit limits or signal-specific liquidation values.

The main welfare and affordability calculations use the survey weights from the household identification module. Sampling was conducted within selected Reera-level study sites and stratified by livestock holdings. For livestock class j in study site c, the weight is the inverse sampling fraction, $\omega_{jc} = N_{jc}/n_{jc}$, using the IBLI team's 2012 population and livestock census. Because the study sites were not randomly sampled from the Borena Zone as a whole, these weights support representation at the Reera and study-site levels, not at broader administrative levels. Table 1 contains unweighted sample moments; the outcome calculations use survey weights unless otherwise indicated.

4.2 Product mapping and comparison regimes

The baseline prior event probability is set to $p = 0.15$, motivated by the product's 15th-percentile index trigger. The standardized loss and indemnity are set to $L = M = 0.5$ of mean herd value, implying an actuarially fair pooled premium of $pL = 7.5\%$ of mean herd value. This common loss scale makes the comparison of insurance affordability across households transparent. The main exercises hold p fixed and vary forecast precision, thereby isolating changes in forecast information from changes in underlying event risk. Section 6.6 separately varies p to examine how the same pricing mechanism changes as baseline climate risk increases.

Two differences from the marketed product govern interpretation. First, the baseline specification treats a trigger as paying the maximum indemnity, whereas actual IBLI indemnities vary with index severity after a trigger. To gauge the importance of this simplification, I also consider illustrative payout ramps in which expected indemnities conditional on a trigger are below the maximum. These scenarios imply actuarially fair rates of 3.75%–5.00%, compared with 7.5% in the baseline. The lower rates therefore reflect less expected protection, not a lower price for identical coverage. Second, marketed coverage can be purchased in small quantities. The lumpy package is therefore a benchmark for a quantity restriction, not a description of every available IBLI contract.

The application consequently reports three complementary measures. The standardized package measures affordability at a common protection scale. Small-contract affordability measures whether households can finance any positive amount of insurance. The own-exposure specification allows households to choose insurance coverage and livestock sales relative to a household-specific exposure target. Its purpose is to assess how protection

changes when households can adjust both their exposure and the resources used to finance premiums. These three measures distinguish adequate protection, access to positive coverage, and protection under endogenous exposure and financing.

For the adaptation-access group, the reference technology is $\chi(e) = e$ and $h(e) = Le^{-\beta e}$, with $\beta = 10$. An expenditure of $e = 0.05$ reduces the conditional physical loss by $1 - e^{-0.05\beta}$: approximately 22, 39, 63, and 86 percent for $\beta = 5, 10, 20, 40$. These expenditure–loss mappings give the assumed effectiveness levels an economic scale. The reference assignment grants access to the top asset quintile. In the reported lumpy comparisons, non-adapting households face the lumpy package, while the adaptation-access group retains the residual-loss insurance calculation. Section 6.4 varies the technology and the assignment, and separately allows constrained households to choose adaptation and insurance jointly.

In the reference welfare comparisons, I report the adaptation-access group separately and classify the remaining households by liquidity relative to the standardized package: below the pooled-price threshold ($w_i < pL$), the marginal band ($pL \leq w_i < L$), and the high-liquidity non-adapting group ($w_i \geq L$). Thus, the marginal band has no adaptation access by construction and corresponds directly to the household class studied in Proposition 3.

To move beyond the common-scale benchmark, the own-exposure specification allows households to finance premiums by selling livestock while recognizing that those sales also reduce the exposure that remains to be insured. This avoids treating the same livestock as both retained exposure and freely available cash. Let x_{is} be livestock sales in mean-herd-value units, let $r \in (0, 1]$ be the fraction of their value recovered, and write $\ell_i = b_i + \sigma_i$ for cash savings plus realized borrowing. With H_i denoting the pre-sale herd in TLU, remaining insurable loss is

$$L_i(x) = \frac{1}{2\bar{H}} \min\{H_i - x\bar{H}, 15\}. \tag{17}$$

Equation (17) removes sold animals from the exposure base. Holding modeled adaptation effort at $e_i = 0$, the household solves

$$\begin{aligned} &\max_{0 \leq x_{is} \leq H_i/(2\bar{H}),\, 0 \leq m_{is} \leq L_i(x_{is})} \mathbb{E}[u_i(z_{is}) \mid s] \quad \text{s.t. } p_s m_{is} \leq \ell_i + r x_{is}, \\ &z_{is}(D) = R_i - p_s m_{is} - (1 - r)x_{is} - D[L_i(x_{is}) - m_{is}]. \end{aligned} \tag{18}$$

In equation (18), sale proceeds finance premiums, $(1-r)x_{is}$ is the resource cost of selling, and sales are capped at half of the original herd. This prevents counting the same animal both as retained exposure and as cash raised by selling it. The sale-financing constraint replaces the implicit liquidation budget in equation (16). Although modeled adaptation effort is zero, sales can still adjust physical exposure.

To scale the welfare effect by the insurance opportunity available in the same specification, let

$$\Gamma_i^{\text{pool}}(r) = \text{CE}_i^{\text{pool,ins}}(r) - \text{CE}_i^{\text{pool,no ins}}(r),$$
$$\Xi(\theta, r) = \frac{-\sum_i \omega_i [\text{CE}_i(\theta, r) - \text{CE}_i^{\text{pool,ins}}(r)]}{\sum_i \omega_i \Gamma_i^{\text{pool}}(r)}. \tag{19}$$

Both pooled choices in equation (19) optimize livestock sales under the same recovery rate, sales cap, and exposure rule; only insurance availability changes. This avoids counting the value of selling livestock as an insurance benefit. The ratio uses population-weighted sums rather than averaging household ratios. It is reported as unavailable when the aggregate denominator is nonpositive, and remains signed when information produces a gain.

4.3 State-contingent protection and benchmark adequacy

The own-exposure exercise starts with the target $\min\{H_i, 15\}$ TLU and removes livestock sold to pay premiums from the remaining exposure. The mechanical financing calculation in Appendix Figure A1 gives mean protected shares of 91.5% under pooled pricing, 84.2% at $\theta = 0.7$, and 71.1% at $\theta = 0.9$, at a sale-recovery rate of one half. This calculation measures protection under a prescribed financing rule rather than an optimized household choice.

Table 2 optimizes livestock sales and insurance coverage over sale-recovery rates from one half to one. At the reference recovery rate and $\theta = 0.9$, optimal protection is 5.7% of own exposure after adverse news and 38.0% after favorable news. The resulting certainty-equivalent loss is 0.09% of mean herd value. Under pooled pricing, coverage is 14.8%, and insurance raises welfare by 0.159% of mean herd value relative to the same sales problem without insurance. The repricing loss therefore amounts to 56.4% of the attainable pooled-insurance value. Thus, a small absolute welfare effect can still represent a substantial loss relative to the insurance value available in this specification.

The own-exposure specification differs from the standardized benchmark because livestock sales affect both financing and exposure. Households can respond to higher premiums by purchasing less insurance, but selling livestock also reduces the loss that remains to be insured. This endogenous destocking margin can generate positive information effects at high recovery rates and does not contradict Corollary 1, which holds physical exposure fixed. Appendix Table A2 shows the connection to the standardized benchmark. With full sale recovery and a half-herd sales cap, the maximum insurance budget equals the baseline liquidity measure; if exposure is held fixed, the resulting allocation reproduces the standardized divis-

ible benchmark. Allowing sales to reduce exposure adds the destocking margin. Section 6.3 examines how these results change when sale recovery is costly and can deteriorate after adverse news. The exercise remains static, leaving dynamic herd adjustment and associated wealth effects outside the calculation (Lybbert et al., 2004; Barrett and Santos, 2014).

Returning to the common-scale benchmark, the protection target is held fixed across households. Figure 1 locates household liquidity relative to the pooled and adverse-signal premiums. Define

$$S_i(\theta) = \mathbb{1}\{w_i < \pi_B(\theta)\}, \qquad A_i(\theta) = \min\left\{1, \frac{w_i}{\pi_B(\theta)}\right\}. \tag{20}$$

Here, $S_i(\theta)$ indicates whether household i cannot finance the full standardized package after adverse news, while $A_i(\theta)$ gives the share of that package it can finance. Consistent with Proposition 1, the full-package shortfall rises and the financeable share falls as forecast precision increases. The adverse-signal premium crosses median household liquidity near $\theta = 0.89$. At $\theta = 0.9$, 54.1% cannot finance the full package and the mean financeable share is 72.9%, whereas only 1.5% cannot finance a one-TLU contract (Figure 2 and Table 3). The contrast highlights the distinction between insurance participation and adequate protection. Informative pricing operates primarily on the intensive margin in this calibration: most households can still finance a small amount of insurance, even though many can no longer finance economically meaningful coverage.

To assess the welfare magnitude of this insurance-side effect, define the gain from *full* pooled insurance relative to no insurance as

$$\Gamma_i^{\text{full}} = \frac{1}{\gamma_i} \ln\left(1 - p + pe^{\gamma_i L}\right) - pL. \tag{21}$$

The quantity Γ_i^{full} is the risk premium eliminated by actuarially fair full pooled coverage. Appendix Table A3 scales the mean welfare loss from repricing, with modeled adaptation disabled and physical exposure fixed, by the mean of Γ_i^{full}. The resulting ratio measures the share of the welfare value of full pooled insurance lost under forecast-contingent pricing. At $\theta = 0.9$, this share is 50% for the population, 51% for the marginal band, and 58% for households below the pooled-price threshold. This common denominator allows the insurance-side welfare effects to be compared across household groups. The own-exposure analysis in Table 2 and Appendix Table A2 instead scales welfare changes by the pooled-insurance value attainable within each financing specification. Although modeled adaptation is also disabled in those exercises, endogenous livestock sales reduce post-sale exposure, so they are not pure fixed-exposure Hirshleifer comparisons.

Finally, among households matched to purchase records, only 1.7% cannot finance their first observed quantity at $\theta = 0.9$. Small observed purchases may already reflect subsidies, trust, basis risk, and other demand constraints. Their affordability therefore complements, rather than replaces, the coverage-adequacy measures.

4.4 Distributional incidence

This subsection examines how forecast-contingent pricing affects households differently depending on their liquidity and access to adaptation. Figure 3 relates these distributional patterns to the theoretical household types. In the lumpy comparison, the marginally insured group loses as forecast precision increases. The adaptation-access group initially experiences a small loss but gains at higher precision, while households initially below the pooled-price threshold gain slightly once favorable-signal coverage becomes affordable. The latter pattern follows from their initially uninsured position under the lumpy menu described in Remark 4. The reversal for the adaptation-access group instead reflects the activation of the adaptation technology.

At the reference values $p = 0.15$, $L = 0.5$, and $\beta = 10$, adaptation is inactive at the prior and becomes active after adverse news only when $\theta > 0.586$, as characterized in Remark 3. Before adaptation becomes active, the premium lottery generates a small welfare loss. Once adaptation is activated, the expenditure-saving channel in Lemma 3(iii) begins to offset that loss, although the net information effect need not become positive immediately. This activation pattern explains the initial decline and eventual gain of the adaptation-access group in Figure 3.

For the marginal band, the loss at $\theta = 0.7$ is 0.44% under the lumpy package and 0.17% under divisible coverage. At $\theta = 0.9$, the respective losses are 1.45% and 0.99%. Divisible coverage therefore cushions the welfare loss by preserving partial insurance, consistent with Corollary 1, rather than eliminating the effect of forecast-contingent pricing.

Figure 4 examines how these distributional effects relate to household liquidity. Panel A reports the critical affordability thresholds $\bar{\theta}_i$, defined by $\pi_B(\bar{\theta}_i) = w_i$ for households with $pL \leq w_i < L$. Households already unable to finance pooled full coverage are assigned the plotting threshold $1/2$, while those with $w_i \geq L$ have no binding threshold and are assigned $+\infty$. Panel B reports the change in the certainty-equivalent gap between the adaptation-access group and the marginal band. Because these are observed household groups under the reference assignment, the comparison combines differences in resources, risk preferences, and adaptation opportunities. It therefore complements, rather than replicates, the controlled

pure-liquidity comparison in Proposition 4(i).

4.5 Aggregate money-metric effects

Figure 5 first aggregates household certainty-equivalent changes with survey weights. These are changes relative to the corresponding pooled allocation, expressed as percentages of mean herd value, not percentages of each household's income. Under the reference top-quintile, $\beta = 10$ scenario, the divisible-coverage aggregate at $\theta = 0.9$ is -0.45%. Its sign reflects the net effect of decision gains and lost insurance smoothing under these assumptions.

The figure also reports Atkinson equally distributed equivalents (EDEs), which attach greater weight to lower consumption levels. For each contract specification, write $\widetilde{\omega}_i = \omega_i / \sum_j \omega_j$ and aggregate the positive certainty-equivalent consumption levels $\mathrm{CE}_i(\theta)$ as

$$\mathrm{EDE}_\varepsilon(\theta) = \begin{cases} \left[\sum_i \widetilde{\omega}_i \mathrm{CE}_i(\theta)^{1-\varepsilon}\right]^{1/(1-\varepsilon)}, & \varepsilon \geq 0,\ \varepsilon \neq 1, \\ \exp\left[\sum_i \widetilde{\omega}_i \ln \mathrm{CE}_i(\theta)\right], & \varepsilon = 1. \end{cases} \tag{22}$$

The parameter ε measures inequality aversion; $\varepsilon = 0$ recovers the weighted mean, and the reported inequality-sensitive comparisons use $\varepsilon = 1$ or 2. The inputs are household consumption *levels*, including $R_i = a_i + \bar{y}$, rather than $\Delta\mathrm{CE}_i$ or realized state-specific consumption. The reported change is $\mathrm{EDE}_\varepsilon(\theta) - \mathrm{EDE}_\varepsilon(1/2)$, expressed as a percentage of mean herd value. Its larger loss in the reference scenario reflects greater weight on losses at low consumption levels. These calculations use $\bar{y} = 1$ mean-herd-value unit in the reference scenario and are sensitive to that baseline, whereas the mean certainty-equivalent change is invariant to a common additive endowment under CARA.

Within the competitive risk-neutral benchmark, expected insurer profits remain zero, so the household loss represents foregone risk sharing rather than a transfer to insurers. Section 6.8 discusses the additional inputs needed for a broader social-surplus evaluation.

5 Contract Timing and Implementation

The previous results arise because information arrives before the opportunity to insure. This section distinguishes three questions: what an enforceable earlier contract can protect, what sustains household participation in that contract, and how it would be financed. Treating these questions separately avoids interpreting an idealized timing comparison as evidence that households can implement it unaided.

5.1 Enforceable pre-signal coverage

I begin with the first question: how much of the pooled insurance allocation can be preserved when coverage can be locked in before the forecast is observed. Let $\lambda \in [0, 1]$ denote the share of benchmark coverage locked before the forecast at the pooled price. The remaining share is purchased afterward. Spot contracting has $\lambda = 0$ and full commitment has $\lambda = 1$. For a non-adapting household seeking full protection, total premiums are

$$\lambda pL + (1 - \lambda)p_s(\theta)L. \tag{23}$$

Equation (23) shows the two effects of commitment. It reduces premium variation across signals and lowers the adverse-signal top-up that must be financed. When $w_i \geq pL$, full commitment can restore the certain pooled consumption outcome for a household without adaptation.

Figure 6 applies this comparison to the marginal band. Partial commitment reduces its full-coverage shortfall and its information-related welfare loss. The forecast still arrives before the physical loss, so changing the insurance date need not remove the opportunity to use information for adaptation. Preserving that opportunity, however, does not require specifying an optimal committed contract for every adapting household.

The aggregate exercise therefore takes a deliberately narrower approach. Appendix Table A4 retains the adaptation-access group's forecast response and applies the insurance lock only to the standardized allocation of the remaining households. Each lock regime is evaluated against its own pooled counterpart. At full commitment, the information effect for non-adapting households is zero by construction, so the aggregate effect equals the weighted access share multiplied by the adapting group's average retained information effect. The reference value of $+0.35\%$ therefore represents the retained contribution of the adapting group, rather than a separate estimate of the gain from introducing earlier contracts. The sign of this aggregate effect therefore depends on whether forecast information is valuable to the adapting group, as characterized by Lemma 3.

This aggregate exercise should not be interpreted as a fully financed policy comparison. The committed block is supplied under the allocation convention used in Appendix Table A4, while the separate liquidity exercise assesses households' capacity to finance it. A fully financed policy evaluation would instead require a common reference allocation, explicit accounting for any financing gap, and joint reoptimization of adaptation and the committed insurance contract. Section 6.8 discusses these distinctions in more detail.

5.2 What sustains household commitment?

An earlier pooled contract creates a retention problem after favorable news. A household may prefer to cancel the contract and purchase the same protection at the lower spot price. The following result characterizes the minimum prepayment needed to prevent this particular lapse decision; it does not address where the household obtains the funds for that prepayment.

Proposition 5 (Minimum prepayment under unsecured commitment)**.** *Under Assumptions 1–2, consider a household with $w_i \in [pL, L)$ and inactive adaptation. It locks a share $\lambda \in (0, 1]$ at pooled premium λpL, pays a non-refundable deposit $d \in [0, \lambda pL]$, and owes $\lambda pL - d$ after the forecast. Any additional coverage is purchased at the spot price.*

(i) Honoring the contract is weakly preferred to lapsing and replacing the locked coverage after favorable news if and only if

$$d \geq d^*_\lambda(\theta) \equiv \lambda L[p - p_G(\theta)]. \tag{24}$$

For $\lambda > 0$, this minimum increases strictly with precision, from zero at $\theta = 1/2$ to λpL at $\theta = 1$.

(ii) Under actuarially fair pricing, the same amount is the break-even deposit for an insurer that retains adverse-signal buyers but loses every favorable-signal buyer.

(iii) If the deposit is drawn from the resources otherwise available at $t = 1$, full-package feasibility remains

$$\lambda pL + (1 - \lambda)\pi_s \leq w_i. \tag{25}$$

Proof. See Appendix B.8. □

Equation (24) prices the temptation to lapse: the deposit must cover the saving from replacing the locked block at the favorable-signal spot price. At equality, the household is indifferent between honoring and replacing the contract; strict retention therefore requires a larger deposit or an appropriate cancellation provision. When $\lambda = 0$, there is no locked coverage and hence no lapse problem.

Under fair pricing, the favorable-signal buyer's saving from lapse, $\lambda L(p - p_G)$, is exactly the excess of its pooled premium over its expected indemnity. Across signal states,

$$q_G \lambda L(p - p_G) = q_B \lambda L(p_B - p).$$

A forfeited deposit of this amount replaces the cross-subsidy lost when favorable-signal buyers lapse, which explains why the household-retention and insurer break-even thresholds coincide.

Equation (25) shows that the deposit does not raise the total premium obligation when it is paid out of resources otherwise available at $t = 1$: it reduces both the outstanding premium bill and the remaining household budget by the same amount. The required deposit as a share of the locked premium depends only on the prior and forecast precision, while its level also scales with the amount of locked coverage. At $p = 0.15$, the minimum deposit is 53.1% of the locked premium at $\theta = 0.7$ and 87.2% at $\theta = 0.9$.

Remark 6 (Pre-signal resources). The incentive condition above is distinct from the household's ability to finance the deposit before the forecast arrives. Let ℓ_i^0 denote resources available at that date. For informative signals, the deposit constraint implies

$$\lambda_i^{\max}(\theta) = \min\left\{1, \frac{\ell_i^0}{L[p - p_G(\theta)]}\right\}.$$

This is only a deposit-feasibility bound, not a complete financing condition. At $\theta = 0.9$, full commitment requires a deposit of 6.54 percent of mean herd value and a total premium of 7.5 percent, compared with an adverse-signal spot premium of approximately 30.7 percent. If the date-1 liquidity measure were available before the signal, every household in the marginal band could finance full commitment.

Across $\theta \in \{0.6, 0.7, 0.8, 0.9\}$, the share of marginal-band households whose cash savings, or savings plus realized borrowing, can fund the minimum full-commitment deposit ranges from 1.6 to 4.7%. These calculations do not imply that assets are intrinsically more difficult to liquidate before news arrives. Because the survey does not identify date-0 resources, household-funded commitment capacity cannot be directly calibrated.

Commitment, enforcement, and financing are therefore distinct institutional problems. Earlier enrollment, lender-linked or group contracts, and cancellation provisions illustrate possible arrangements, but they are not evaluated interventions in this paper. The exact deposit result applies to the stated replacement comparison for households in the marginal band. For households below that band, constrained partial replacement requires a separate participation calculation, so no corresponding deposit ordering is inferred here.

5.3 Targeted support after adverse news

Targeted premium support addresses the immediate coverage constraint while leaving forecast-contingent pricing in place. I first consider the lumpy benchmark. Restoring full coverage for a household with $w_i < \pi_B(\theta)$ requires a transfer equal to the adverse-signal financing gap,

$$\tau_i^B(\theta) = \pi_B(\theta) - w_i. \tag{26}$$

Because the transfer is paid only after an adverse signal, its expected fiscal cost is $q_B(\theta)\tau_i^B(\theta)$. Table 4 reports the resulting household certainty-equivalent gain per unit of fiscal expenditure, with the corresponding equal-cost unconditional transfer normalized to one.

With divisible insurance, support instead operates on the intensive margin by purchasing additional adverse-signal coverage. Baseline adverse-signal coverage is $m_{iB}^0 = \min\{L, w_i/p_B\}$. For households with $m_{iB}^0 < L$, I consider a small adverse-signal grant $\tau_0 = 10^{-4}$ in mean-herd-value units. Holding the household's own premium contribution fixed, the grant raises coverage to

$$m_{iB}^+ = \min\{L, m_{iB}^0 + \tau_0/p_B\}. \tag{27}$$

The corresponding benefit-cost ratio is $[\mathrm{CE}_i^{\mathrm{sup}}(\theta;\tau_0) - \mathrm{CE}_i^{\mathrm{sup}}(\theta;0)]/(q_B\tau_0)$. When the increment does not reach the coverage cap, this finite difference approximates the marginal welfare gain from targeted support. The fiscal denominator charges the full grant even when the cap binds.

To isolate this adverse-signal margin, the exercise holds the favorable-signal allocation fixed at the allocation used in the lumpy support comparison: full coverage if $w_i \geq \pi_G$ and no coverage otherwise. Thus, "divisible, marginal" refers only to the adverse-signal coverage increment, not to a reoptimization of insurance in both signal states. Table 4 reports unweighted household percentiles of these individual benefit-cost ratios within each eligible group.

The comparison shows that the value of targeting depends on contract divisibility. Under the lumpy benchmark, filling a small financing gap near the affordability threshold can restore an entire insurance package; at $\theta = 0.9$, the marginal band's median benefit-cost ratio is 1.5. With divisible insurance, the same type of support purchases coverage continuously rather than generating a discrete restoration of the package, and the below-threshold group has a slightly higher median ratio. These calculations measure household welfare gains per unit of fiscal expenditure, not net social returns after administrative costs or the marginal cost of public funds. They therefore illustrate why the evaluation of targeted premium support cannot be separated from the insurance menu.

6 Robustness, Sensitivity, and Scope

The baseline separates a general pricing mechanism from a conditional numerical application. This section asks which conclusions survive alternative measurement and technology assumptions. Sampling intervals describe uncertainty in the observed household distribution. The remaining exercises vary inputs that the survey does not identify, rather than treating all uncertainty as sampling error.

6.1 Sampling uncertainty

This subsection examines how sensitive the main quantitative results are to sampling variation across survey locations. Appendix Figure A2 and Table A5 report percentile intervals from a block bootstrap over the eight Woreda (district-level) sampling clusters. The full-package shortfall has a 95% interval of $[21.4, 32.2]$ at $\theta = 0.7$ and $[46.0, 59.7]$ at $\theta = 0.9$. The corresponding intervals for the lumpy marginal-band welfare loss are $[0.38, 0.53]$ and $[1.36, 1.50]$. With only eight clusters, these intervals are necessarily coarse. They capture sampling variation across locations while holding the model parameters fixed, and therefore do not reflect uncertainty about adaptation effectiveness, forecast skill, or liquidation capacity.

6.2 Liquidity, contract size, and payout severity

This subsection examines how the underinsurance result depends on three features of the insurance environment: the liquidity households can mobilize, the size of the available contract, and the severity of payouts conditional on a trigger.

I first vary α, the fraction of livestock value that can be mobilized to finance premiums. Appendix Figure A3 and Table A6 retain the same forecast-contingent pricing mechanism while changing this liquidity assumption. Across the reported values of α, adverse-state protection declines as forecast precision increases, although the level of coverage is sensitive to the assumed budget. These comparisons should therefore be interpreted as liquidity scenarios rather than estimates of unused credit or liquidation capacity.

Contract size changes the distinction between access to insurance and adequate protection. A two-TLU contract provides approximately 13 percent of the standardized 15-TLU benchmark. Relative to household-specific protection targets, its median coverage ratio is 24.6%, and 73.9% of households obtain less than half of their target protection (Appendix Tables A7 and A8). Thus, the ability to finance a small contract can coexist with substantial

underinsurance.

Finally, payout severity affects affordability through the premium required to finance expected indemnities. Appendix Table A9 holds the nominal maximum indemnity fixed while lowering expected indemnity conditional on a trigger. At $\theta = 0.9$, the full-package shortfall decreases from 54.1% to 37.5% when expected indemnity is two-thirds of the maximum, and to 29.1% when it is one-half. The lower shortfall reflects a lower premium because the contract provides less expected protection, not a lower price for the same coverage. These exercises hold conditional payout severity fixed across signals and therefore do not estimate the joint distribution of forecasts and payout severity.

6.3 Household-specific exposure and costly sales

The own-exposure problem in equations (17)–(18) allows livestock sales both to finance insurance and to reduce the exposure that remains at risk. This subsection examines how these two roles of sales shape the results and how they depend on the value recovered from liquidation.

For an interior solution financed partly through livestock sales, the uninsured gap satisfies

$$L_i(x_{is}) - m_{is} = \frac{1}{\gamma_i} \ln \left[\frac{1 - p_s}{r - p_s(1 - k_i)} \right], \qquad k_i = -L_i'(x_{is}), \tag{28}$$

subject to a positive denominator and feasibility of the interior choice. Here, k_i captures the marginal reduction in insured exposure generated by an additional sale. Where the 15-TLU target binds strictly, $k_i = 0$; where the remaining herd is below that target, $k_i = 1/2$. Equation (28) therefore shows why optimal coverage is sensitive to the sale-recovery rate. At the reference value $r = 1/2$, the below-target interior gap is independent of the insurance price, so variation in coverage comes from financing constraints and corner solutions rather than this interior condition. Table 2 reports the numerical optimum across recovery rates, allowing for the sale cap, the exposure kink, the full-coverage financing boundary, and interior solutions.

Livestock sales also create an exposure-adjustment margin that is absent from the fixed-exposure benchmark. Although modeled adaptation effort is set to $e_i = 0$, equation (17) removes sold animals from drought exposure. When the remaining herd is below the 15-TLU target, so that $k_i = 1/2$, the direct effect of a marginal sale on conditional expected consumption, holding insurance coverage fixed, is positive when

$$p_s > 2(1 - r).$$

An adverse posterior can therefore make destocking attractive even without the separate adaptation technology. This mechanism explains the positive information effects at high constant recovery rates in Appendix Table A10: for the reported cases with $r \in \{0.7, 0.8, 0.9\}$, the adverse posterior can cross the destocking threshold while the favorable posterior does not. The condition identifies the direct incentive to reduce exposure, however, not the sign of the total welfare effect, because forecast-contingent premiums and financing responses remain present.

To separate the financing and exposure channels, Appendix Table A2 modifies the baseline in stages. With a fixed standardized loss, full recovery $r = 1$, and maximum sales equal to one half of the herd, the available budget is $\ell_i + 0.5a_i = w_i$, and sales impose no resource loss; this reproduces the baseline divisible allocation. Lowering r introduces costly liquidation. Replacing the fixed loss with the capped own-herd target then makes exposure household specific, and removing sold animals from that target adds the destocking margin. Because this final margin can itself respond to forecast information, the sign of the own-exposure welfare effect need not coincide with the pure fixed-exposure Hirshleifer comparison.

The positive effects under high constant recovery also depend on liquidation values remaining unaffected by adverse public news. Appendix Table A10 relaxes this assumption by setting the pooled and favorable-signal recovery rates to a normal value r_N and lowering only the adverse-signal rate, $r_B = r_N - \delta$, with $\delta \geq 0$. A 0.10 decline in adverse-state recovery turns three of the four positive constant-recovery cases negative and reduces the remaining gain, at $r_N = 0.9$ and $\theta = 0.9$, from +0.678 to +0.102 percent of mean herd value. A 0.20 decline makes all reported information effects negative. The pooled-insurance-value denominator remains fixed within each r_N because the pooled comparison continues to use the normal recovery rate. These calculations are reduced-form sensitivities to worse liquidation terms after common adverse news, not an equilibrium model of livestock prices. All own-exposure comparisons are also static, leaving dynamic herd accumulation and general-equilibrium livestock-market responses outside the calculation.

6.4 Adaptation access and effectiveness

Adaptation is the main source of uncertainty about the aggregate welfare sign. With divisible insurance at $\theta = 0.9$, the reported changes are -0.95% with no access, -0.45% under the top-quintile reference, -0.11% with top-tercile access, and +1.70% with universal access. At $\beta = 10$, assigning access from the top of the asset distribution produces a zero crossing near 38% access. This is a conditional model threshold, not an empirical estimate or a universal

access requirement.

Appendix Figures A5 and A6 vary access and effectiveness separately. As Remark 3 shows, a higher β lowers the cost of physical protection but also changes whether adaptation is active at the pooled benchmark, so the *incremental* forecast value can be non-monotone. The comparison is always with the same technology under uninformative pricing. Appendix Table A12 keeps liquidity groups fixed when access is reassigned.

Pre-treatment mobility provides a behavioral diagnostic of the adaptation assignment. In the livestock module, 47.8 percent of households move some animals to a satellite camp, 38.6 percent move at least 10 percent, 25.2 percent at least one quarter, and 11.8 percent at least one half; 18.2 percent move lactating animals (Appendix Table A13). These intensity measures use recorded animal counts rather than value-weighted exposure. Their interpretation is observed movement, leaving the option to move, forecast responsiveness, and causal loss reduction unidentified.

Assigning the modeled technology to the households selected by each proxy tests the consequences of that assignment. For any mobility, aggregate welfare at $\theta = 0.9$ is -0.63 percent with $\beta = 5$ and +0.30 percent with $\beta = 10$. For movement of at least 10 percent, the corresponding values are -0.69 and +0.06 percent; for at least one quarter, they are -0.78 and -0.29 percent. Both composition and effectiveness matter. Appendix Table A13 and the generated outputs use the same calibration households and survey weights. The assignment grants the modeled scalable technology; the accompanying diagnostic also records whether assigned households can fund its use.

A separate exercise allows the baseline marginal band to choose insurance and adaptation jointly within its liquidity constraint. Its change at $\theta = 0.9$ is -0.99% without adaptation, -0.36% with $\beta_{\text{poor}} = 5$, and +1.66% with $\beta_{\text{poor}} = 10$ (Appendix Table A14). The loss survives the less effective technology but reverses under the stronger one. This is a conditional result about feasible self-protection, not evidence that a chosen mobility threshold measures effective adaptation.

6.5 Basis risk, preferences, and aggregation

This subsection examines whether the welfare results are sensitive to three features omitted or normalized in the reference calibration: basis risk, risk preferences and premium loadings, and the baseline consumption level used for inequality-sensitive aggregation.

Basis risk weakens both the protection provided by insurance and the welfare cost of losing that protection. Field evidence from index-based livestock insurance shows that insured

pastoralists can retain substantial residual risk in high-loss states (Jensen et al., 2016), while basis risk can also alter how insurance demand varies with risk aversion (Clarke, 2016). At $\theta = 0.9$, reducing payout reliability from $\varphi = 1$ to $\varphi = 0.7$ lowers the marginal band's welfare loss from 1.45% to 0.54%, while model-predicted pooled full-package take-up remains 88.3%. At still lower reliability, voluntary nonparticipation further reduces the insurance value exposed to repricing, shrinking the measured loss toward zero (Appendix Figure A8 and Table A15). This numerical pattern is specific to the modeled joint exposure and payout structure and does not imply a general sign result for arbitrary forms of basis risk.

Risk preferences and premium loadings affect different margins. Rescaling risk aversion changes welfare magnitudes without changing the mechanical affordability condition. At $\theta = 0.9$, halving reference risk aversion reduces the marginal-band welfare loss to 0.68%, while doubling it raises the loss to 3.25%; the full-package shortfall and mean affordable share remain 54.1% and 72.9%, respectively. Premium loadings, by contrast, affect both affordability and willingness to pay. Moderate loadings preserve pooled demand for the benchmark package over the reported grid, while sufficiently large loadings eliminate pooled participation and therefore leave little insurance value for forecast-contingent repricing to destroy (Appendix Figure A9 and Table A16). These exercises are regime checks rather than estimates of the marketed product's effective loading, which would require information on its payout distribution and insurer costs.

Finally, Appendix Figure A10 and Table A17 vary baseline non-livestock income to assess how inequality-sensitive aggregation depends on consumption levels. The mean certainty-equivalent change remains fixed at -0.56% of mean herd value across the reported income levels, as implied by CARA preferences. The Atkinson EDE loss, by contrast, becomes smaller in magnitude as baseline income rises: for $\varepsilon = 1$, it changes from -1.01% at a baseline income of 0.6 mean-herd units to -0.71% at 5.0 units, while for $\varepsilon = 2$ it changes from -1.12% to -0.82%. The inequality-sensitive effects nevertheless remain negative and larger in magnitude than the mean certainty-equivalent loss throughout the reported range. Thus, the stronger losses under inequality-sensitive aggregation reflect the greater weight placed on lower consumption levels, while their quantitative magnitude depends on the baseline consumption normalization.

6.6 Baseline climate risk and forecast skill

Climate change can alter the underlying frequency and severity of extreme weather in developing economies (Grant et al., 2025). The model's prior p captures the frequency margin rather than changes in loss severity. To distinguish this margin from forecast precision, Ap-

pendix Table A18 repeats the quantitative exercise for $p \in \{0.10, 0.15, 0.20, 0.25, 0.30\}$ while holding the forecast technology fixed. A higher p raises the pooled premium and both signal-contingent posterior risks; for $\theta > 1/2$, it also raises the probability of adverse news because $\partial q_B / \partial p = 2\theta - 1 > 0$.

The sensitivity shows that higher baseline risk strengthens the insurance-side channel. At $\theta = 0.9$, increasing p from the baseline 0.15 to 0.30 raises the probability of an adverse signal from 22 to 34 percent and the posterior event probability after adverse news from 61.4 to 79.4 percent. The full-package shortfall rises from 54.1 to 63.6 percent, while mean affordable coverage falls from 72.9 to 65.6 percent. With adaptation disabled, the population certainty-equivalent change falls from -0.95 to -1.70 percent of mean herd value, and the share of full pooled-insurance value destroyed rises from 50 to 58 percent. Under the reference adaptation assignment, the corresponding aggregate change moves more modestly from -0.45 to -0.59 percent.

The aggregate effect need not be monotone in baseline risk once adaptation is active: at $\theta = 0.7$, the no-adaptation loss grows from -0.19 to -0.41 percent as p rises from 0.15 to 0.30, whereas the reference-adaptation aggregate changes from -0.13 to -0.07 percent. Higher baseline risk therefore robustly increases exposure to high-premium states and worsens insurance adequacy in this calibration, but its total welfare effect remains conditional on adaptation.

For the symmetric signal in equation (1), the hit rate and correct-rejection rate both equal θ. The area under the piecewise-linear ROC curve implied by this single operating point therefore equals θ, while the Peirce skill score equals $2\theta - 1$ (Mason and Graham, 2002). The scenarios $\theta = 0.7$, 0.8, and 0.9 thus correspond to Peirce scores of 0.4, 0.6, and 0.8. These mappings place the illustrative scenarios on standard forecast-verification scales; they do not calibrate forecast skill for the insured forage-loss event. Dynamical forecasts for East Africa generally show stronger skill for the short rains than for the long rains (Walker et al., 2019; Young and Klingaman, 2020), while the forage index itself becomes increasingly predictable within the season as shorter integration windows capture much of its interannual variation (Vrieling et al., 2016). The latter also bears on how late an insurance sales window can close before information about the insured season becomes available.

A direct calibration would match archived forecasts at the relevant purchase date with the realized woreda-level forage index and estimate $\Pr(B \mid D = 1)$ and $\Pr(G \mid D = 0)$ out of sample. The code can ingest these rates separately, but the current inputs do not include that matched series. Consequently, $\theta = 0.7, 0.8, 0.9$ are illustrative technology scenarios rather than estimates of currently available forecast skill. Uncertainty on this margin affects

quantitative relevance even when the mechanism remains well defined.

6.7 Supporting evidence on demand and fixed pricing

This subsection provides two pieces of supporting evidence around the central repricing counterfactual. The first asks whether observed insurance demand responds to price changes at prevailing, non-forecast-contingent prices. The second examines what an insurer would face if premiums instead remained fixed as forecast information changed expected payouts.

Randomized premium discounts provide evidence on the first margin. Appendix Table A19 reports take-up and unconditional insured-quantity regressions with Woreda fixed effects and HC1 standard errors. A 10-percentage-point discount increases take-up by approximately 4.5 percentage points and unconditional insured quantity by 0.25 TLU. The experimental discounts generate exogenous price variation, but they do not span the high premiums implied by informative adverse signals. The results therefore support price responsiveness in insurance demand without identifying the forecast-driven repricing counterfactual studied in the model.

The second exercise considers the alternative of holding the charged premium fixed as forecast information changes expected payouts. Appendix Figure A11 uses a fixed premium of 0.075 per unit of insured value and an illustrative payout ramp with expected indemnity of 0.25 per unit conditional on a trigger. Define the adverse-signal *insurer loss ratio* as expected indemnity divided by premium:

$$LR_B(\theta) = \frac{0.25\, p_B(\theta)}{0.075}. \tag{29}$$

The 0.075 rate represents the observed premium scale used in this exercise, whereas 0.25 is an assumed conditional payout under the illustrative ramp. Both are expressed per unit of insured value and therefore differ from the standardized full-loss package priced at $p_s L$. Define $\theta^{\mathrm{LR}} = \inf\{\theta : LR_B(\theta) \geq 1\}$, which is distinct from the household affordability thresholds $\bar{\theta}_i$. In the reported grid, the adverse-signal loss ratio reaches break-even at approximately $\theta^{\mathrm{LR}} = 0.71$ and rises to $LR_B(0.9) = 2.05$.

The figure also reports a fair-rate comparison and an ex-ante loss ratio that allows participation to respond to the signal. Under the fixed observed-price scale, favorable-signal participation declines, weakening the cross-subsidization that could otherwise offset adverse-signal losses. These calculations illustrate the pressure that informative public signals can place on fixed pricing, but they do not establish that IBLI historically repriced on forecasts, identify an optimal restriction on sales timing, or determine which institutional constraint

prevents earlier contracting.

6.8 Scope of the aggregate and timing comparisons

The reported aggregates are household certainty-equivalent changes under specified prices, preferences, and adaptation opportunities. They exclude insurer capital and reinsurance responses, dynamic herd effects, and the financing and administrative costs of a policy intervention. Mobility assignments and universal access describe scenarios rather than statistical bounds. The fixed-price loss-ratio exercise is counterfactual accounting; the available data identify neither historical forecast-driven repricing nor which institutional constraint prevents earlier enrollment.

The full-commitment entries retain the adapting group's existing spot response and set the non-adapting information effect to zero relative to each regime's pooled counterpart. The resulting aggregate is a weighted group contribution. A policy gain would instead compare feasible allocations against a common benchmark and jointly specify committed coverage, subsequent adaptation, and the cost of any external finance. The marginal-band timing comparison is narrower: these non-adapting households can already finance pooled full coverage. Similarly, the premium-support ratios compare household benefits per fiscal unit, before administrative costs or the cost of public funds.

7 Conclusion

This paper asks how the value of better climate information depends on the financial and contractual environment in which households can respond to it. When a public forecast arrives before insurance is purchased, competitive pricing converts information about future losses into variation in premiums. Adverse news therefore makes protection more expensive precisely when it is most valuable. Liquidity constraints do not create this underlying risk-sharing loss, but they can amplify it by forcing households to reduce coverage after adverse news. Contract divisibility determines whether this adjustment takes the form of underinsurance or complete exclusion.

The quantitative application to livestock insurance in Ethiopia shows why participation alone can be a poor measure of protection. Small contracts remain affordable for most households even as meaningful coverage deteriorates, and forecast-contingent repricing can eliminate a substantial share of the insurance value available under pooled pricing. Adaptation creates an opposing benefit, so the aggregate welfare effect depends on who can adapt

and how effective those responses are. The same distinction matters as underlying climate risk rises: more frequent adverse events expose households more often to high-premium states and deepen the insurance-side loss, while the overall welfare effect remains conditional on adaptation opportunities.

Contract timing provides a way to separate these two roles of information. Insurance purchased before the forecast can preserve pooled risk sharing while still allowing households to use subsequent information for adaptation. The numerical commitment exercise shows the potential of this arrangement, but also its financing requirements: the resources needed to sustain a pre-forecast contract need not be available to the households that benefit most from it. Participation, financing, and contractual commitment are therefore distinct implementation margins.

Important empirical uncertainties remain, including forecast skill for the insured event, the evolution of baseline climate risk, the effectiveness and distribution of adaptation opportunities, and the costs of earlier contracting. The model also abstracts from dynamic herd adjustment and equilibrium responses in insurance and livestock markets. These limitations affect magnitudes and institutional comparisons, but not the basic mechanism. Better forecasts should therefore be evaluated together with the insurance and financial arrangements through which households respond to them, rather than as an information technology in isolation.

References

Barrett, C. B., F. Chabari, D. Bailey, P. D. Little, and D. L. Coppock (2003). Livestock pricing in the northern Kenyan rangelands. *Journal of African Economies 12*(2), 127–155.

Barrett, C. B. and P. Santos (2014). The impact of changing rainfall variability on resource-dependent wealth dynamics. *Ecological Economics 105*, 48–54.

Belissa, T., E. Bulte, F. Cecchi, S. Gangopadhyay, and R. Lensink (2019). Liquidity constraints, informal institutions, and the adoption of weather insurance: A randomized controlled trial in Ethiopia. *Journal of Development Economics 140*, 269–278.

Boomhower, J., M. Fowlie, J. Gellman, and A. Plantinga (2024). How are insurance markets adapting to climate change? Risk classification and pricing in the market for homeowners insurance. NBER Working Paper 32625, National Bureau of Economic Research.

Boomhower, J., M. Fowlie, and A. J. Plantinga (2023). Wildfire insurance, information, and self-protection. In *AEA Papers and Proceedings*, Volume 113, pp. 310–315.

Boucher, S. R., M. R. Carter, J. E. Flatnes, T. J. Lybbert, J. G. Malacarne, P. P. Marenya, and L. A. Paul (2024). Bundling genetic and financial technologies for more resilient and productive small-scale farmers in Africa. *Economic Journal 134*(662), 2321–2350.

Burlig, F., A. Jina, E. M. Kelley, G. V. Lane, and H. Sahai (2024). The value of forecasts: Experimental evidence from India. NBER Working Paper 32173, National Bureau of Economic Research.

Carriquiry, M. A. and D. E. Osgood (2012). Index insurance, probabilistic climate forecasts, and production. *Journal of Risk and Insurance 79*(1), 287–300.

Carter, M., A. De Janvry, E. Sadoulet, and A. Sarris (2017). Index insurance for developing country agriculture: A reassessment. *Annual Review of Resource Economics 9*, 421–438.

Carter, M. R., L. Cheng, and A. Sarris (2016). Where and how index insurance can boost the adoption of improved agricultural technologies. *Journal of Development Economics 118*, 59–71.

Casaburi, L. and J. Willis (2018). Time versus state in insurance: Experimental evidence from contract farming in Kenya. *American Economic Review 108*(12), 3778–3813.

Clarke, D. J. (2016). A theory of rational demand for index insurance. *American Economic Journal: Microeconomics 8*(1), 283–306.

Cochrane, J. H. (1995). Time-consistent health insurance. *Journal of Political Economy 103*(3), 445–473.

Cohn, A. S., P. Newton, J. D. B. Gil, L. Kuhl, L. Samberg, V. Ricciardi, J. R. Manly, and S. Northrop (2017). Smallholder agriculture and climate change. *Annual Review of Environment and Resources 42*(1), 347–375.

Cole, S. A. and W. Xiong (2017). Agricultural insurance and economic development. *Annual Review of Economics 9*, 235–262.

Finkelstein, A., K. McGarry, and A. Sufi (2005). Dynamic inefficiencies in insurance markets: Evidence from long-term care insurance. *American Economic Review 95*(2), 224–228.

Grant, L., I. Vanderkelen, L. Gudmundsson, E. Fischer, S. I. Seneviratne, and W. Thiery (2025). Global emergence of unprecedented lifetime exposure to climate extremes. *Nature 641*(8062), 374–379.

Handel, B., I. Hendel, and M. D. Whinston (2015). Equilibria in health exchanges: Adverse selection versus reclassification risk. *Econometrica 83*(4), 1261–1313.

Hendel, I. and A. Lizzeri (2003). The role of commitment in dynamic contracts: Evidence from life insurance. *Quarterly Journal of Economics 118*(1), 299–328.

Hill, R. V., N. Kumar, N. Magnan, S. Makhija, F. de Nicola, D. J. Spielman, and P. S. Ward (2019). Ex ante and ex post effects of hybrid index insurance in Bangladesh. *Journal of Development Economics 136*, 1–17.

Hirshleifer, J. (1971). The private and social value of information and the reward to inventive activity. *American Economic Review 61*(4), 561–574.

Jensen, N. D., C. B. Barrett, and A. G. Mude (2016). Index insurance quality and basis risk: Evidence from northern Kenya. *American Journal of Agricultural Economics 98*(5), 1450–1469.

Jensen, N. D., A. G. Mude, and C. B. Barrett (2018). How basis risk and spatiotemporal adverse selection influence demand for index insurance: Evidence from northern Kenya. *Food Policy 74*, 172–198.

Ker, A. P. and P. McGowan (2000). Weather-based adverse selection and the US crop insurance program: The private insurance company perspective. *Journal of Agricultural and Resource Economics 25*(2), 386–410.

Lane, G. (2024). Adapting to climate risk with guaranteed credit: Evidence from Bangladesh. *Econometrica 92*(2), 355–386.

Linsenmeier, M. and J. G. Shrader (2023). Global inequalities in weather forecasts. SocArXiv Working Paper 7e2jf, Center for Open Science.

Liu, Y., K. Chen, and R. V. Hill (2020). Delayed premium payment, insurance adoption, and household investment in rural China. *American Journal of Agricultural Economics 102*(4), 1177–1197.

Luo, H., J. R. Skees, and M. A. Marchant (1994). Weather information and the potential for intertemporal adverse selection in crop insurance. *Review of Agricultural Economics 16*(3), 441–451.

Luseno, W. K., J. G. McPeak, C. B. Barrett, P. D. Little, and G. Gebru (2003). Assessing the value of climate forecast information for pastoralists: Evidence from southern Ethiopia and northern Kenya. *World Development 31*(9), 1477–1494.

Lybbert, T. J., C. B. Barrett, S. Desta, and D. L. Coppock (2004). Stochastic wealth dynamics and risk management among a poor population. *Economic Journal 114*(498), 750–777.

Lybbert, T. J., C. B. Barrett, J. G. McPeak, and W. K. Luseno (2007). Bayesian herders: Updating of rainfall beliefs in response to external forecasts. *World Development 35*(3), 480–497.

Mason, S. J. and N. E. Graham (2002). Areas beneath the relative operating characteristics (ROC) and relative operating levels (ROL) curves: Statistical significance and interpretation. *Quarterly Journal of the Royal Meteorological Society 128*(584), 2145–2166.

McIntosh, C., F. Povel, and E. Sadoulet (2019). Utility, risk and demand for incomplete insurance: Lab experiments with Guatemalan co-operatives. *Economic Journal 129*(622), 2581–2607.

Morton, J. F. (2007). The impact of climate change on smallholder and subsistence agriculture. *Proceedings of the National Academy of Sciences 104*(50), 19680–19685.

Pauly, M. V., H. Kunreuther, and R. Hirth (1995). Guaranteed renewability in insurance. *Journal of Risk and Uncertainty 10*(2), 143–156.

Rosenzweig, M. R. and C. Udry (2014). Rainfall forecasts, weather, and wages over the agricultural production cycle. *American Economic Review 104*(5), 278–283.

Schlee, E. E. (2001). The value of information in efficient risk-sharing arrangements. *American Economic Review 91*(3), 509–524.

Shrader, J. G., S. Thies, L. Bakkensen, M. Linsenmeier, and D. Lemoine (2026). Weather forecasts become more important for reducing mortality as the climate warms. *Proceedings of the National Academy of Sciences 123*(16), e2523372123.

Solomon, A. (2026). Imperfect private information in insurance markets. *Review of Economics and Statistics 108*(2), 485–503.

Takahashi, K., C. B. Barrett, and M. Ikegami (2019). Does index insurance crowd in or crowd out informal risk sharing? Evidence from rural Ethiopia. *American Journal of Agricultural Economics 101*(3), 672–691.

Takahashi, K., M. Ikegami, M. Sheahan, and C. B. Barrett (2016). Experimental evidence on the drivers of index-based livestock insurance demand in southern Ethiopia. *World Development 78*, 324–340.

Takahashi, K., Y. Noritomo, M. Ikegami, and N. D. Jensen (2020). Understanding pastoralists' dynamic insurance uptake decisions: Evidence from four-year panel data in Ethiopia. *Food Policy 95*, 101910.

Vrieling, A., M. Meroni, A. G. Mude, S. Chantarat, C. C. Ummenhofer, and C. A. J. M. de Bie (2016). Early assessment of seasonal forage availability for mitigating the impact of drought on East African pastoralists. *Remote Sensing of Environment 174*, 44–55.

Walker, D. P., C. E. Birch, J. H. Marsham, A. A. Scaife, R. J. Graham, and Z. T. Segele (2019). Skill of dynamical and GHACOF consensus seasonal forecasts of East African rainfall. *Climate Dynamics 53*(7–8), 4911–4935.

Yi, F., M. Zhou, and Y. Y. Zhang (2020). Value of incorporating ENSO forecast in crop insurance programs. *American Journal of Agricultural Economics 102*(2), 439–457.

Young, H. R. and N. P. Klingaman (2020). Skill of seasonal rainfall and temperature forecasts for East Africa. *Weather and Forecasting 35*(5), 1783–1800.

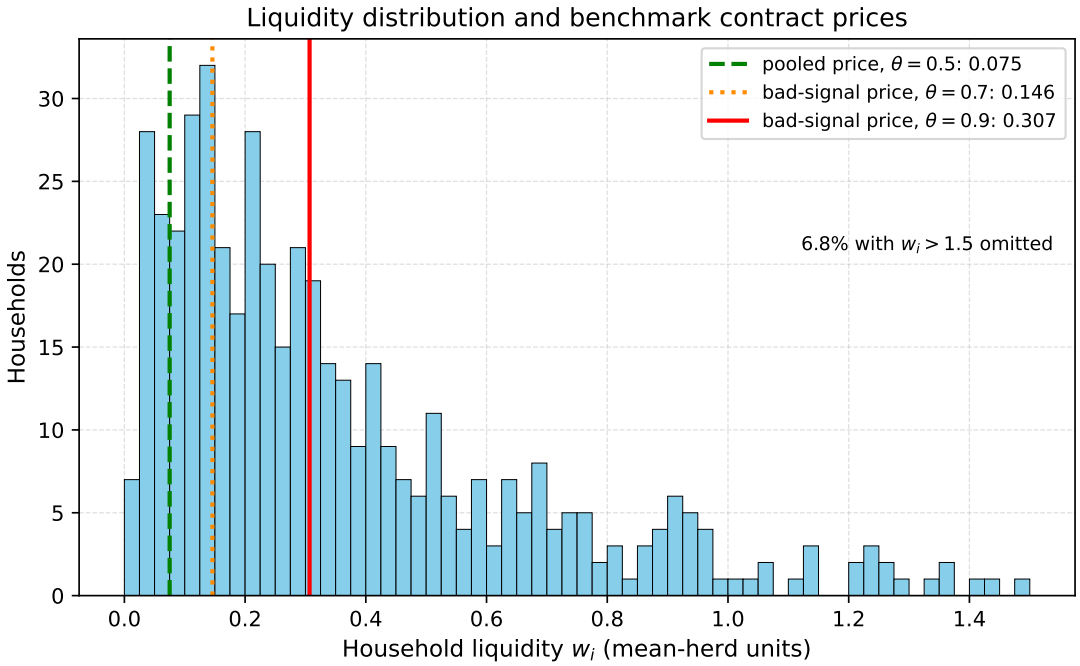


Figure 1: Household Liquidity and Signal-Contingent Benchmark Prices

Notes: The histogram reports household liquidity in mean-herd-value units. Vertical lines mark pooled and adverse-signal prices of the standardized package. Precision levels are illustrative, not estimates of observed forecast skill. The histogram counts households; reported outcome shares elsewhere use survey weights.

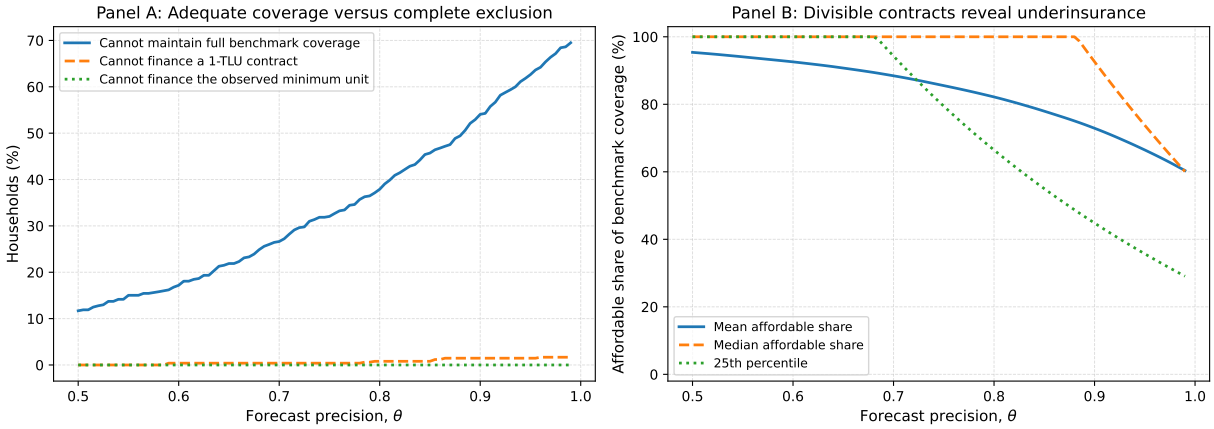


Figure 2: Common-Scale Adequacy and Small-Contract Affordability

Notes: Panel A distinguishes inability to finance the complete standardized package from inability to purchase a small positive contract. Panel B reports financeable shares of standardized coverage with divisible insurance. These are common-scale affordability measures, not optimal own-exposure coverage ratios.

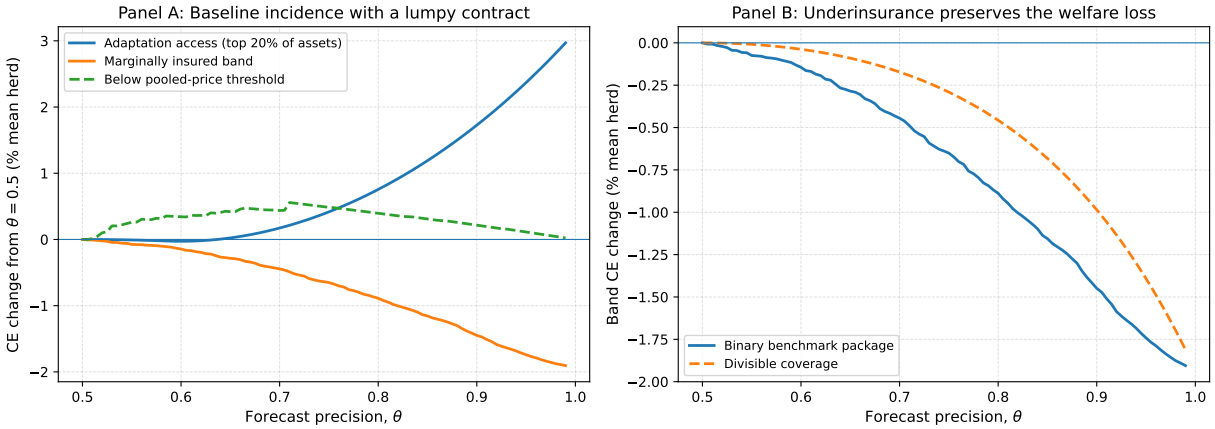


Figure 3: Welfare Incidence under Lumpy and Divisible Coverage

Notes: Panel A compares the reference adaptation-access group, the marginal band, and households below the pooled-price threshold in the lumpy comparison; the figure label "binary" denotes that same two-point insurance menu. The adaptation-access group retains its residual-loss insurance calculation. Panel B compares the same non-adapting marginal band under lumpy and divisible coverage. Changes are relative to each specification's pooled counterpart and measured as percentages of mean herd value.

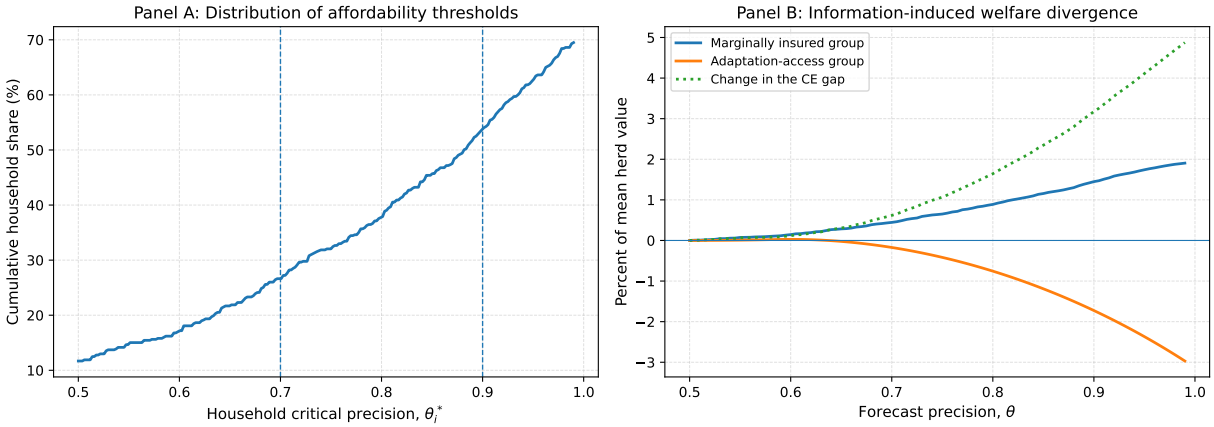


Figure 4: Affordability Thresholds and Welfare-Gap Changes

Notes: Panel A shows the household affordability thresholds $\bar{\theta}_i$ defined in Section 4.4; the plotted label θ_i^* denotes the same threshold. Panel B compares certainty-equivalent changes for the reference adaptation-access group and the marginal band in the lumpy comparison. Its gap combines differences in liquidity, adaptation access, and risk preferences; it is not the common-preference liquidity-only comparison in Proposition 4(i).

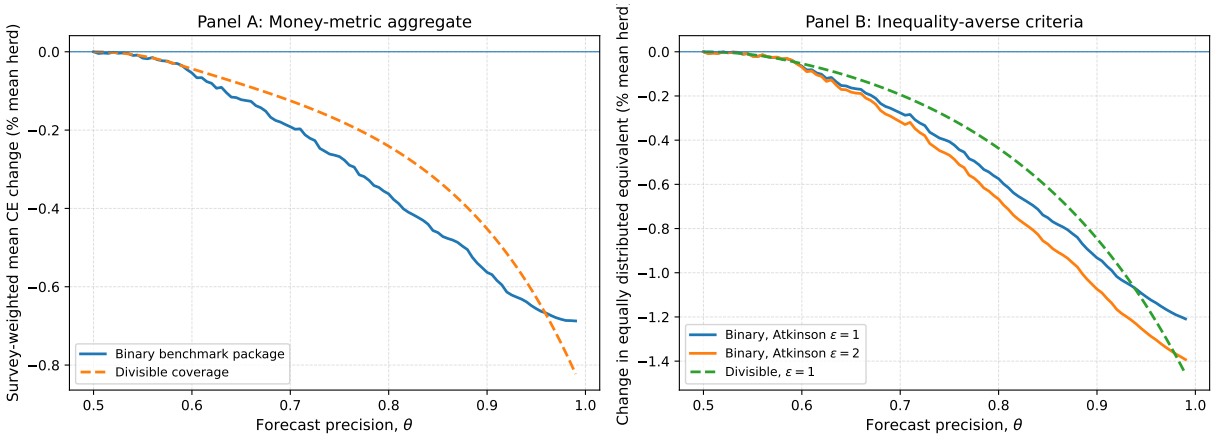


Figure 5: Aggregate Money-Metric Effects under the Reference Adaptation Scenario

Notes: Panel A reports survey-weighted mean certainty-equivalent changes; Panel B reports Atkinson equally distributed equivalents from equation (22), applied to positive household certainty-equivalent consumption levels, for the labeled contract specifications. Access is assigned to the top asset quintile with reference effectiveness $\beta = 10$. Neither is estimated. All changes are relative to pooled pricing with the same technology and are measured in percent of mean herd value. Baseline non-livestock income equals one mean herd value; Appendix Figure A10 varies that normalization.

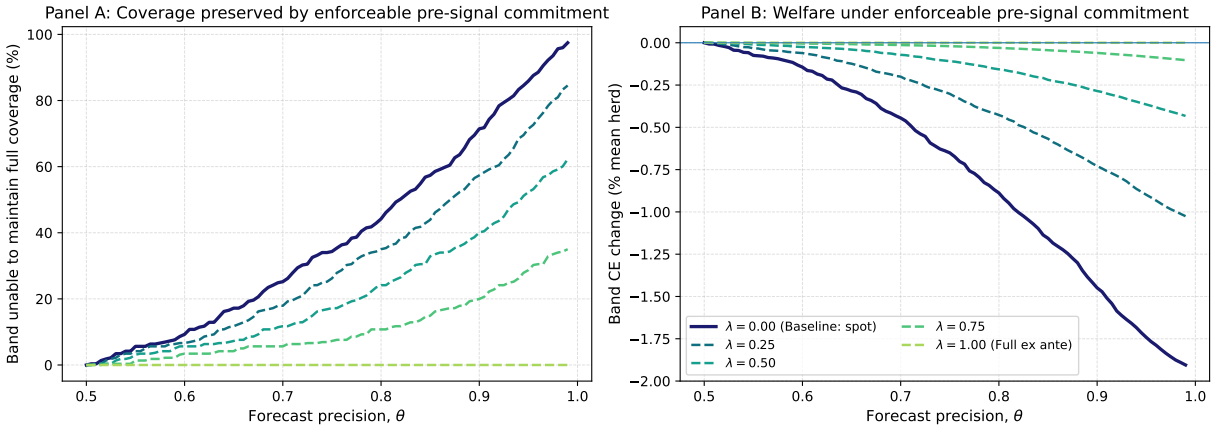


Figure 6: Marginal-Band Protection under Enforceable Pre-Signal Commitment

Notes: A share λ of the standardized package is committed at the pooled price, with any remainder purchased after the forecast. Panel A shows the marginal band's full-coverage shortfall; Panel B shows its certainty-equivalent change under the lumpy comparison. These households can finance pooled full coverage and have inactive adaptation. The figure therefore isolates their timing comparison, not the financing or adaptation problem for the entire population.

Table 1: Household Sample Characteristics

Variable	Mean	Standard deviation	Minimum	Maximum
Livestock assets (TLU)	15.03	22.23	0.20	266.44
Borrowing (ETB)	253.82	940.43	0.00	10000.00
Cash savings (ETB)	1083.58	7320.33	0.00	100000.00
Risk aversion index	1.03	0.46	0.50	2.00
Liquidity (mean-herd units)	0.53	0.80	0.01	10.22

Notes: Unweighted sample moments from the 506-household calibration sample. The model's outcome calculations use survey weights unless stated otherwise. One tropical livestock unit corresponds to one cow, 0.7 camel, or ten sheep/goats under the data conversion. Liquidity is measured in units of the survey-weighted mean herd value.

Table 2: Own-Exposure Coverage and Welfare across Sale-Recovery Rates

r	θ	Coverage (%)			Pooled insurance value (%)	Information ΔCE (%)	Loss / pooled value (%)
		Pooled	Adverse	Favorable			
0.50	0.7	14.8	9.7	22.3	0.159	-0.019	11.8
0.50	0.9	14.8	5.7	38.0	0.159	-0.090	56.4
0.60	0.7	16.6	11.3	24.2	0.179	-0.022	12.5
0.60	0.9	16.6	11.0	39.9	0.179	-0.086	47.8
0.70	0.7	23.1	19.2	29.5	0.249	-0.027	10.8
0.70	0.9	23.1	77.3	44.0	0.249	+0.008	-3.2
0.80	0.7	36.0	45.7	38.9	0.394	-0.014	3.5
0.80	0.9	36.0	82.7	50.7	0.394	+0.338	-85.7
0.90	0.7	75.1	92.8	63.4	0.604	+0.253	-41.9
0.90	0.9	75.1	90.1	66.3	0.604	+0.678	-112.2
1.00	0.7	100.0	100.0	100.0	0.423	-0.034	8.0
1.00	0.9	100.0	100.0	100.0	0.423	-0.190	45.0

Notes: Households jointly choose sales and coverage in equation (18), with modeled adaptation effort fixed at zero and sales capped at half the herd. Recovery r is the fraction of sale value retained and is held constant across forecast states in this table. Because sold livestock leave the post-sale exposure base, sales also constitute an endogenous exposure-adjustment (destocking) margin; these rows are therefore not pure fixed-exposure no-adaptation comparisons. Coverage columns are survey-weighted means of household protected shares of post-sale exposure. Pooled insurance value is the CE gain from insurance relative to no insurance, optimizing sales under the same recovery and exposure rule in both cases. Information ΔCE compares informative pricing with that specification's pooled allocation; both CE columns are percentages of mean herd value. The final column is minus the weighted mean information change divided by the weighted mean pooled insurance value, as in equation (19). Negative ratios denote information gains; a dash denotes a nonpositive denominator.

Table 3: Coverage Adequacy and Welfare at Selected Precision Levels

θ	Full benchmark shortfall (%)	Mean affordable share (%)	Cannot finance 1 TLU (%)	Band CE change	
				Lumpy (%)	Divisible (%)
0.7	26.6	88.4	0.4	-0.44	-0.17
0.8	37.9	82.2	0.8	-0.89	-0.46
0.9	54.1	72.9	1.5	-1.45	-0.99

Notes: The standardized package covers a loss equal to half of mean herd value. The first three outcome columns distinguish full-package affordability, financeable coverage, and a one-tropical-livestock-unit purchase. The marginal band's certainty-equivalent changes are relative to pooled pricing and expressed as percentages of mean herd value. Negative entries are losses. Forecast precision is illustrative.

Table 4: Value of Bad-Signal Premium Support per Unit of Expected Fiscal Cost

Group (contract form)	θ	p25	Median	p75
Marginal band (lumpy)	0.8	1.4	1.7	2.1
Marginal band (lumpy)	0.9	1.3	1.5	1.8
Below pooled threshold (lumpy)	0.8	1.2	1.3	1.5
Below pooled threshold (lumpy)	0.9	1.1	1.3	1.4
Marginal band (divisible)	0.9	1.2	1.4	1.5
Below pooled threshold (divisible)	0.9	1.2	1.5	1.7

Notes: Household certainty-equivalent gains per unit of expected fiscal expenditure; an equal-cost unconditional transfer is normalized to one. Groups exclude the reference adaptation-access households as defined in Section 4.2, and entries are unweighted percentiles within eligible groups. "Below pooled threshold" denotes $w_i < pL$. Lumpy support restores a complete package and can have a high ratio near the affordability threshold. Divisible, marginal support uses the grant $\tau_0 = 10^{-4}$ and coverage increment in equation (27), charging expected fiscal cost $q_B\tau_0$ and holding favorable-signal coverage at the lumpy support benchmark. These ratios do not deduct administration costs or the social cost of raising public funds.

A Additional Figures and Tables

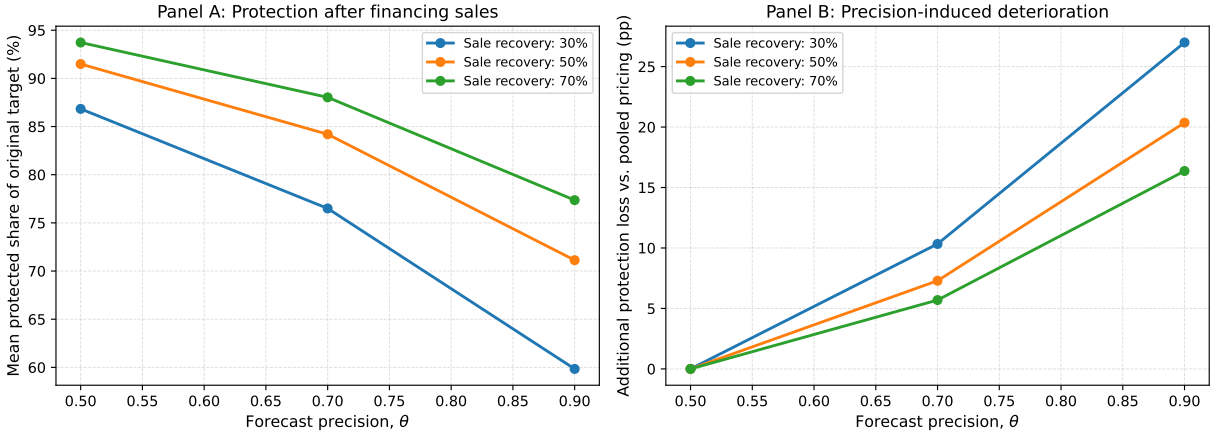


Figure A1: Own-Target Protection under the Mechanical Financing Rule

Notes: Mean protected share of the original target $\min\{H_i, 15\}$ tropical livestock units under the mechanical financing rule. Cash and realized credit are used first, followed by sales at recovery $r \in \{0.3, 0.5, 0.7\}$. Sold livestock leave the insurable exposure base and sales are capped at half the herd. At $r = 0.5$, shares are 91.5% under pooled pricing, 84.2% at $\theta = 0.7$, and 71.1% at $\theta = 0.9$. This is distinct from the joint optimization in Table 2.

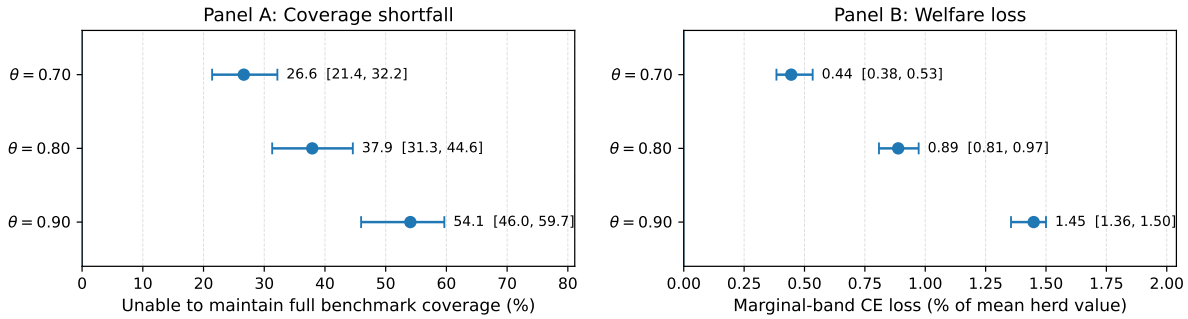


Figure A2: Sampling Uncertainty in Selected Benchmark Outcomes

Notes: Points are original-sample estimates and bars are 95% percentile intervals from 1,000 block-bootstrap replications over eight Woreda clusters. Structural and technology parameters are held fixed. With few clusters, the intervals are coarse and should not be interpreted as covering uncertainty in forecast skill or adaptation effectiveness.

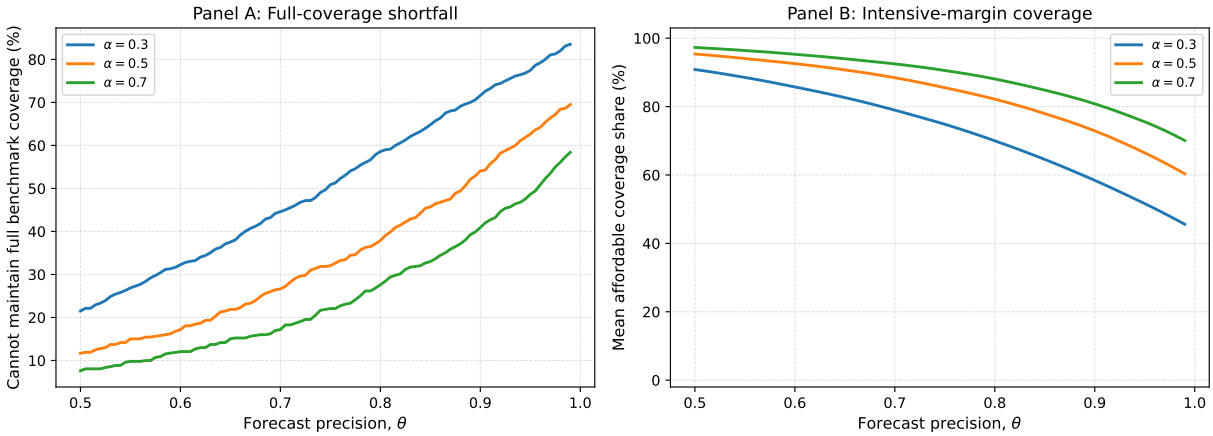


Figure A3: Sensitivity to the Liquidatable Share of Livestock Wealth

Notes: Full-package shortfall and financeable coverage under alternative liquidatable shares $\alpha \in \{0.3, 0.5, 0.7\}$. The reference is $\alpha = 0.5$. Liquidity is $w_i = \alpha a_i + b_i + \sigma_i$ in mean-herd-value units. Groups are reconstituted for each liquidity mapping and outcome statistics use survey weights.

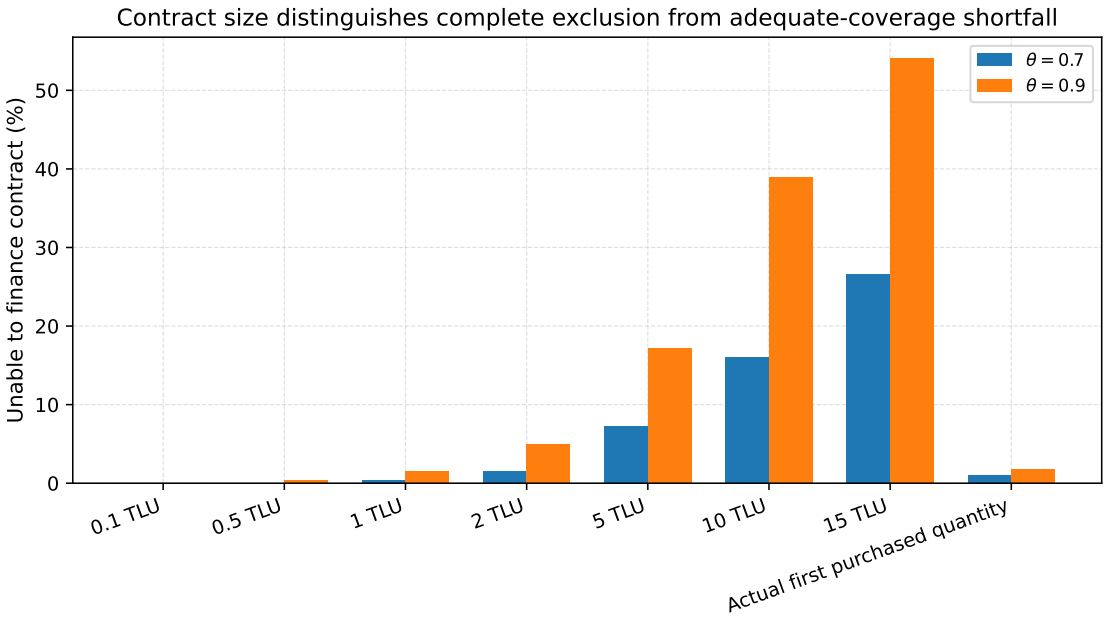


Figure A4: Affordability of Small and Benchmark-Scale Contracts

Notes: Shares unable to finance fixed contract quantities after an adverse signal. Fifteen tropical livestock units supply the standardized quantity benchmark; the actual-purchase comparison uses each matched purchaser's first observed quantity. A small affordable contract establishes access to positive coverage, not adequacy of protection.

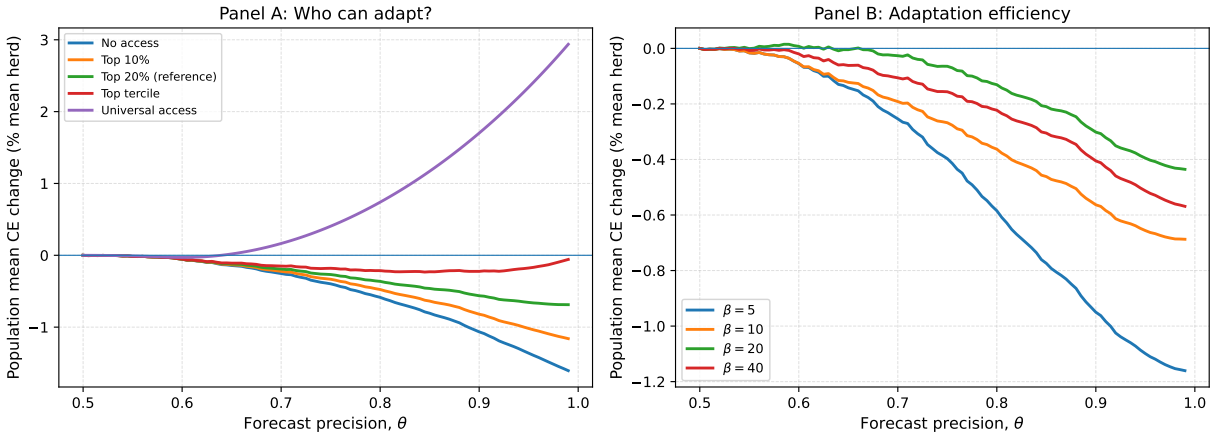


Figure A5: Adaptation Access and Effectiveness in the Lumpy Comparison

Notes: Survey-weighted mean certainty-equivalent changes in the lumpy comparison under alternative access assignments and effectiveness levels. The adaptation-access group retains its residual-loss insurance calculation. Universal access assigns the same reference technology to every household. Figure A6 uses divisible coverage for the other households.

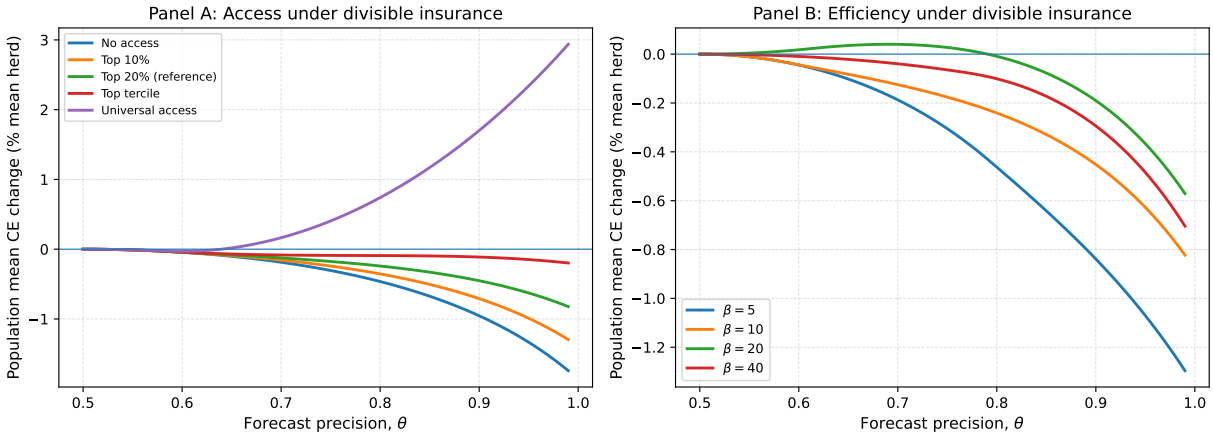


Figure A6: Adaptation Access and Effectiveness with Divisible Insurance

Notes: Panel A varies asset-ranked access at reference effectiveness $\beta = 10$; Panel B varies effectiveness at the top-quintile reference assignment. These are separate sensitivities, not a jointly estimated welfare surface. Spending 5% of mean herd value reduces modeled conditional loss by about 22, 39, 63, and 86 percent at $\beta = 5, 10, 20, 40$. Each information effect is measured relative to pooled pricing with the same technology.

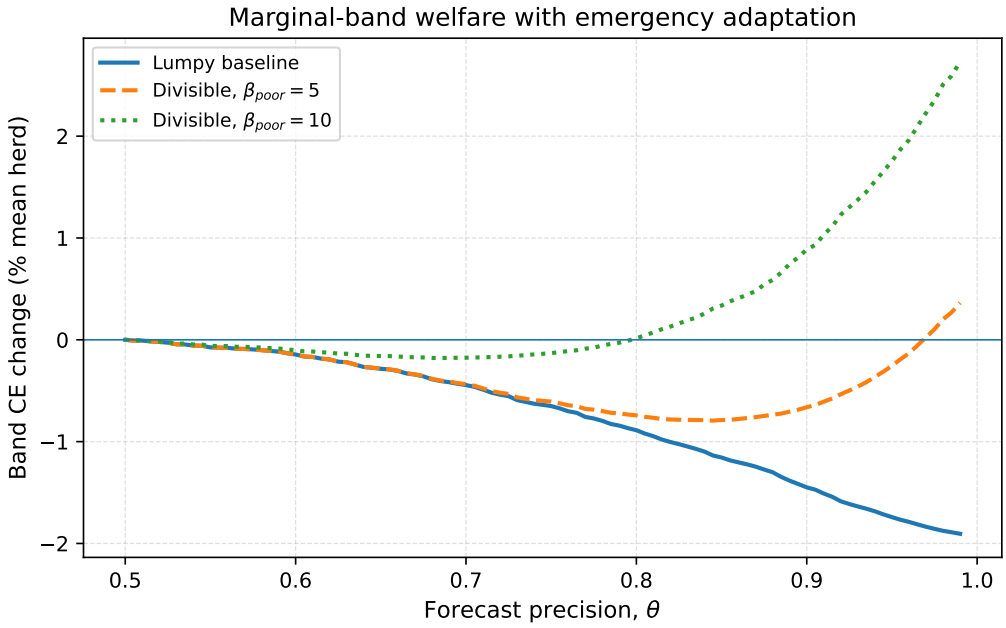


Figure A7: Marginal-Band Adaptation with Lumpy Insurance

Notes: Welfare of the fixed baseline marginal band when continuous emergency adaptation is available but insurance remains a lumpy package. Table A14 instead permits both adaptation and insurance to be chosen continuously within the same liquidity budget.

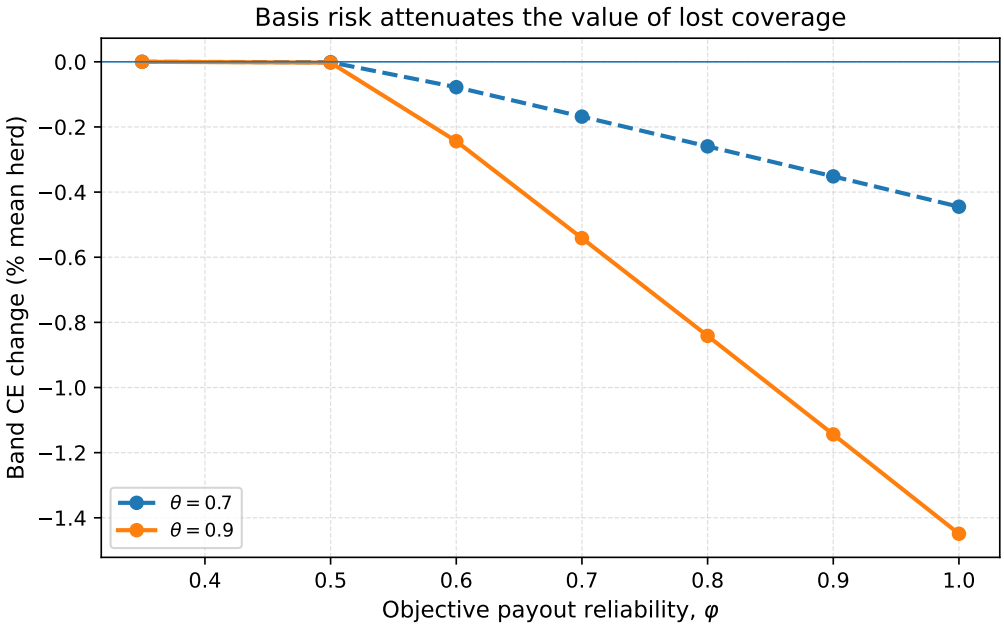


Figure A8: Sensitivity to Index Protection

Notes: Sensitivity to the degree of index protection with voluntary participation. In this specification, lower protection reduces both insurance demand and the value lost through adverse-signal exclusion. The plotted pattern is not a theorem for arbitrary joint distributions of individual loss and index payouts.

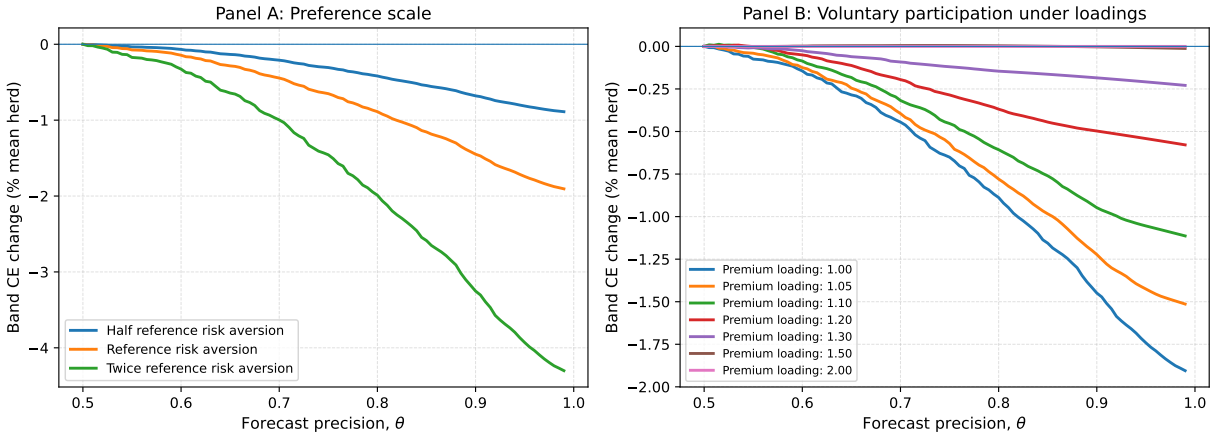


Figure A9: Risk Preferences and Premium Loadings

Notes: Preference rescaling and loadings on the constrained group's standardized package. Loadings change willingness to pay as well as affordability; high loadings can eliminate pooled participation. The exercise is a partial-equilibrium regime check rather than a calibration of the marketed product's effective loading.

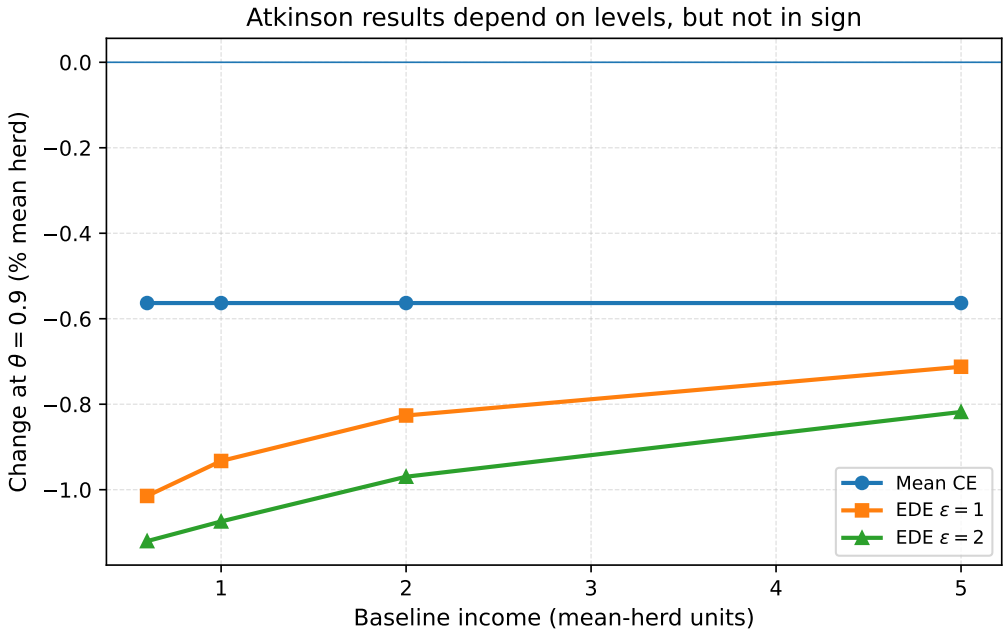


Figure A10: Baseline Income and Inequality-Sensitive Aggregation

Notes: Mean CE and Atkinson equally distributed equivalent changes at $\theta = 0.9$ under the displayed baseline-income assumptions. The comparison uses the lumpy reference scenario. CARA makes mean CE changes invariant to a common additive endowment; inequality-sensitive changes depend on the resulting consumption levels.

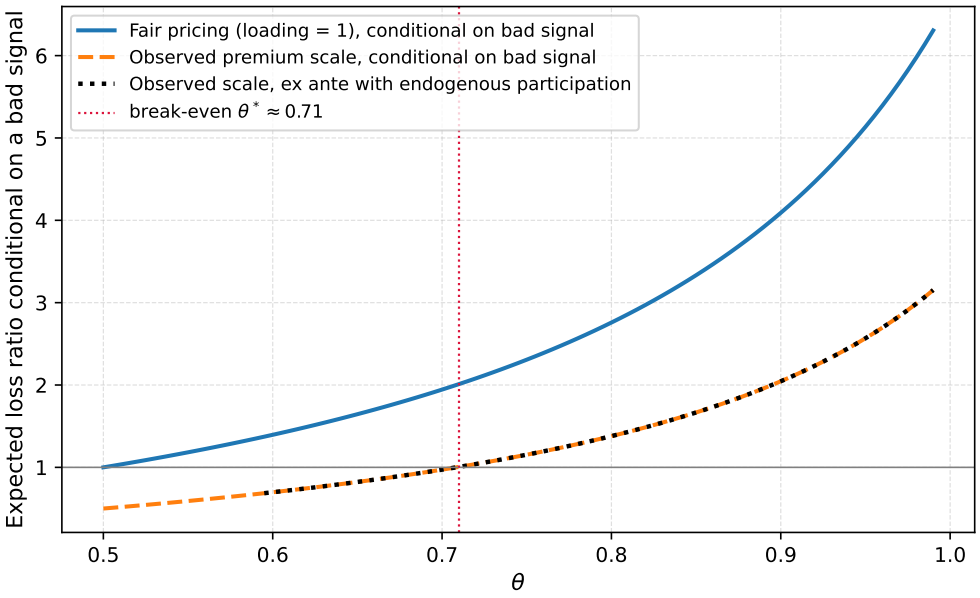


Figure A11: Fixed-Price Insurer Accounting across Forecast Scenarios

Notes: Fixed-price insurer accounting. The solid line uses the ex-ante fair rate. The dashed line is $LR_B(\theta)$ in equation (29): the observed 7.5% charged rate and an illustrative expected indemnity of 0.25 of insured value conditional on a trigger. The plotted break-even label θ^* denotes θ^{LR}, not a household affordability threshold. The dotted line is an ex-ante ratio allowing participation to respond. Favorable-signal demand declines in that calculation, reducing scope for pooling across signal states. These insurer payout-to-premium ratios differ from the shares of pooled-insurance value lost in Table A3.

Table A1: Own-Target Protection under Alternative Sale-Recovery Rates

Sale recovery	$\theta = 0.5$	$\theta = 0.7$	$\theta = 0.9$	Decline 0.5 to 0.9
0.3	86.8	76.5	59.8	27.0
0.5	91.5	84.2	71.1	20.4
0.7	93.7	88.0	77.4	16.4

Notes: Weighted mean protected shares of the original own-herd target under the mechanical financing rule. Livestock sold to raise cash no longer contribute to the exposure that remains insurable. The maximum sale share is one half. These are not the joint optimizing results in Table 2.

Table A2: From Standardized Coverage to Own Exposure: Recovery and Exposure Comparisons

Exposure rule	r	Coverage (%)			Pooled insurance value (%)	Information ΔCE (%)	Loss / pooled value (%)
		Pooled	Adverse	Favorable			
Standardized loss	1.00	95.4	72.9	99.9	1.862	-0.954	51.3
Standardized loss	0.50	9.9	3.5	29.2	0.238	-0.125	52.4
Fixed own exposure	1.00	100.0	100.0	100.0	0.847	-0.377	44.5
Fixed own exposure	0.50	14.7	5.6	38.0	0.158	-0.089	56.4
Sale-adjusted own exposure	1.00	100.0	100.0	100.0	0.423	-0.190	45.0
Sale-adjusted own exposure	0.50	14.8	5.7	38.0	0.159	-0.090	56.4

Notes: All rows use $\theta = 0.9$, fair prices, no modeled adaptation effort, and a half-herd sales cap. Standardized exposure keeps L fixed; fixed own exposure uses the capped original herd; sale-adjusted own exposure removes sold animals and therefore retains an endogenous destocking margin. With standardized exposure and $r = 1$, the budget and allocation coincide with the baseline $\alpha = 0.5$ liquidity model. Within each exposure rule, changing r changes the resource cost of finance. The last three columns use the definitions in Table 2. In particular, the pooled-value denominator is attained under the stated financing constraints, whereas Table A3 uses the value of full pooled coverage.

Table A3: Divisible Coverage and Loss Relative to Full Pooled-Insurance Gains

	Coverage of L (%)			Full pooled value lost (%)		
θ	Pooled	Adverse	Favorable	Population	Band	Below threshold
0.7	95.4	88.4	99.0	10	9	19
0.8	95.4	82.2	99.5	24	24	37
0.9	95.4	72.9	99.9	50	51	58

Notes: Coverage is $m_s = \min\{L, w_i/p_s\}$ with modeled adaptation disabled and physical exposure fixed. The share of full pooled-insurance value lost divides the weighted mean CE loss from repricing by the weighted mean value of full pooled insurance, $\Gamma_i^{\text{full}} = \gamma_i^{-1} \ln(1 - p + pe^{\gamma_i L}) - pL$. That denominator averages 1.90% of mean herd value. It differs from the gain actually attained by a household that can afford only partial pooled coverage. These ratios do not use the reference adaptation scenario and are not percentages of total household welfare.

Table A4: Information Effects under Alternative Access and Commitment Scenarios

Adaptation	Lumpy benchmark		Divisible coverage	
access (%)	Spot ($\lambda = 0$)	Committed ($\lambda = 1$)	Spot ($\lambda = 0$)	Committed ($\lambda = 1$)
0.0	-1.07	+0.00	-0.95	+0.00
10.0	-0.82	+0.17	-0.71	+0.17
20.0	-0.56	+0.35	-0.45	+0.35
33.3	-0.22	+0.57	-0.11	+0.57
42.0	+0.01	+0.72	+0.12	+0.72
60.0	+0.57	+1.03	+0.57	+1.03
100.0	+1.70	+1.70	+1.70	+1.70

Notes: Reported information effects at $\theta = 0.9$ in percent of mean herd value. The reference $\beta = 10$ technology is assigned by asset rank. Access rows use nominal shares, with the 33.3-percent row set to exactly one third, as in Tables A11 and A12; the realized survey-weighted shares are retained in the output data. The implementation applies the insurance lock to non-adapting households and retains the adaptation group's existing forecast response. At full commitment the former's information effect is zero, so the aggregate equals the realized weighted access share times the latter's mean retained effect. Each lock regime uses its own pooled counterpart, so differences between columns need not equal a common-baseline policy gain. Supplying the locked block to households unable to fund it requires financing not costed here. The nonnegative full-commitment entries describe this reference grid, not every possible technology or financed policy.

Table A5: Bootstrap Intervals

Outcome	θ	Estimate	95% interval
Marginal-band CE loss: lumpy contract	0.7	0.44	[0.38, 0.53]
Marginal-band CE loss: lumpy contract	0.8	0.89	[0.81, 0.97]
Marginal-band CE loss: lumpy contract	0.9	1.45	[1.36, 1.50]
Population CE change: lumpy contract	0.7	-0.19	[-0.23, -0.15]
Population CE change: lumpy contract	0.8	-0.36	[-0.43, -0.32]
Population CE change: lumpy contract	0.9	-0.56	[-0.64, -0.50]
Full benchmark coverage shortfall	0.7	26.63	[21.41, 32.16]
Full benchmark coverage shortfall	0.8	37.90	[31.31, 44.59]
Full benchmark coverage shortfall	0.9	54.06	[45.95, 59.66]
Mean affordable coverage share	0.7	88.44	[85.96, 91.37]
Mean affordable coverage share	0.8	82.18	[78.85, 85.89]
Mean affordable coverage share	0.9	72.90	[69.02, 77.35]
Marginal-band CE loss: divisible coverage	0.7	0.17	[0.16, 0.19]
Marginal-band CE loss: divisible coverage	0.8	0.46	[0.42, 0.49]
Marginal-band CE loss: divisible coverage	0.9	0.99	[0.92, 1.04]
Population CE change: divisible coverage	0.7	-0.13	[-0.14, -0.11]
Population CE change: divisible coverage	0.8	-0.24	[-0.27, -0.21]
Population CE change: divisible coverage	0.9	-0.45	[-0.48, -0.39]

Notes: Original-sample outcomes and 95% percentile block-bootstrap intervals. Coverage outcomes are percentages of households; CE losses are percentages of mean herd value. Eight Woreda clusters are resampled, holding the structural scenario fixed.

Table A6: Liquidity Sensitivity

α	$\theta = 0.7$			$\theta = 0.9$		
	Shortfall	Mean share	Band loss	Shortfall	Mean share	Band loss
0.3	44.6	79.0	0.58	71.6	58.4	1.58
0.5	26.6	88.4	0.44	54.1	72.9	1.45
0.7	17.2	92.5	0.39	40.9	80.8	1.42

Notes: Shortfall is the weighted share unable to finance the full package; mean share is weighted financeable coverage. Band loss uses the lumpy comparison and is measured as a percentage of mean herd value. Varying α changes both liquidity and the composition of the marginal band.

Table A7: Affordability by Contract Size

Contract	$\theta = 0.5$	$\theta = 0.7$	$\theta = 0.9$
0.1 TLU	0.0	0.0	0.0
0.5 TLU	0.0	0.0	0.4
1 TLU	0.0	0.4	1.5
2 TLU	0.4	1.5	5.0
5 TLU	1.5	7.2	17.2
10 TLU	7.4	16.0	38.9
15 TLU	11.7	26.6	54.1
Actual first purchased quantity	0.4	1.1	1.7

Notes: Weighted shares unable to finance the indicated quantity after an adverse signal. At $\theta = 0.5$ the signal is uninformative and prices are pooled. The last row uses each matched purchaser's first observed insured quantity.

Table A8: Coverage Adequacy by Contract Size

Contract	Mean coverage ratio (%)	Median coverage ratio (%)	Share below 50% (%)	Share below 25% (%)
1 TLU	22.3	12.3	87.9	73.9
2 TLU	37.5	24.6	73.9	50.6
5 TLU	66.1	61.4	40.4	0.0
10 TLU	88.7	100.0	0.0	0.0
15 TLU	100.0	100.0	0.0	0.0

Notes: For a contract of K tropical livestock units, the ratio is $\min\{K, T_i\}/T_i$ with $T_i = \min\{H_i, 15\}$. These weighted ratios measure coverage relative to the household target, independently of price or sales. Table A7 measures affordability instead.

Table A9: Affordability under Illustrative Payout-Ramp Scenarios

$\mathbb{E}$[indemnity | trigger] (% of maximum)	Pooled fair rate (%)	θ	Full shortfall (%)	Mean affordable share (%)
50	3.75	0.7	11.7	95.6
50	3.75	0.9	29.1	87.6
67	5.00	0.7	15.6	93.2
67	5.00	0.9	37.5	82.4
100	7.50	0.7	26.6	88.4
100	7.50	0.9	54.1	72.9

Notes: The nominal maximum is fixed while expected payout conditional on a trigger varies. Smaller expected payouts reduce both premiums and protection. Conditional severity is held fixed across forecast signals, so the rows are loss-scale sensitivities, not estimates of forecast-specific product severity.

Table A10: Own-Exposure Welfare under Constant and State-Dependent Sale Recovery

r_N	r_B	θ	Coverage (%)			Pooled insurance value (%)	Information ΔCE (%)	Loss / pooled value (%)
			Pooled	Adverse	Favorable			
Panel A: Constant sale recovery across forecast states								
0.70	0.70	0.7	23.1	19.2	29.5	0.249	-0.027	10.8
0.70	0.70	0.9	23.1	77.3	44.0	0.249	+0.008	-3.2
0.80	0.80	0.7	36.0	45.7	38.9	0.394	-0.014	3.5
0.80	0.80	0.9	36.0	82.7	50.7	0.394	+0.338	-85.7
0.90	0.90	0.7	75.1	92.8	63.4	0.604	+0.253	-41.9
0.90	0.90	0.9	75.1	90.1	66.3	0.604	+0.678	-112.2
Panel B: Lower sale recovery after an adverse forecast								
0.70	0.60	0.7	23.1	11.3	29.5	0.249	-0.072	28.7
0.70	0.60	0.9	23.1	11.0	44.0	0.249	-0.150	60.0
0.70	0.50	0.7	23.1	9.7	29.5	0.249	-0.080	32.1
0.70	0.50	0.9	23.1	5.7	44.0	0.249	-0.171	68.5
0.80	0.70	0.7	36.0	19.2	38.9	0.394	-0.134	34.1
0.80	0.70	0.9	36.0	77.3	50.7	0.394	-0.127	32.2
0.80	0.60	0.7	36.0	11.3	38.9	0.394	-0.179	45.4
0.80	0.60	0.9	36.0	11.0	50.7	0.394	-0.285	72.2
0.90	0.80	0.7	75.1	45.7	63.4	0.604	-0.204	33.7
0.90	0.80	0.9	75.1	82.7	66.3	0.604	+0.102	-16.9
0.90	0.70	0.7	75.1	19.2	63.4	0.604	-0.325	53.7
0.90	0.70	0.9	75.1	77.3	66.3	0.604	-0.363	60.0

Notes: Households jointly choose livestock sales and insurance coverage with modeled adaptation effort fixed at zero and a half-herd sales cap. Panel A holds the sale-recovery rate constant across pooled, adverse, and favorable states. Panel B sets the pooled and favorable-signal recovery rates to the normal value r_N and lowers only the adverse-signal recovery rate to $r_B < r_N$. Sold animals leave the exposure base, so sales are also an endogenous destocking response. Pooled insurance value is evaluated at r_N and is therefore unchanged across Panel B rows with the same r_N. Information ΔCE compares informative pricing and state-contingent recovery with the pooled allocation at r_N. The final column is minus the information change divided by the attained pooled insurance value; negative ratios denote information gains. Panel B is a reduced-form sensitivity to lower liquidation values after adverse public news, not an equilibrium livestock-price calculation.

Table A11: Adaptation Sensitivity under the Lumpy Benchmark at $\theta = 0.9$

Dimension	Welfare group	Specification	CE change
Access share	Population	No access	-1.07
Access share	Population	Top 10%	-0.82
Access share	Population	Top 20% (reference)	-0.56
Access share	Population	Top tercile	-0.22
Access share	Population	Universal access	+1.70
Adaptation efficiency	Population	$\beta = 5$	-0.95
Adaptation efficiency	Population	$\beta = 10$	-0.56
Adaptation efficiency	Population	$\beta = 20$	-0.30
Adaptation efficiency	Population	$\beta = 40$	-0.41
Marginal-group technology	Marginal band	Lumpy baseline	-1.45
Marginal-group technology	Marginal band	Divisible, $\beta_{\text{poor}} = 5$	-0.66
Marginal-group technology	Marginal band	Divisible, $\beta_{\text{poor}} = 10$	+0.89

Notes: CE changes in percent of mean herd value. Access and effectiveness are scenario inputs. In the lumpy comparison, the adaptation-access group retains its residual-loss insurance calculation. Information gains need not increase monotonically with β: at high effectiveness, adaptation is already active in the pooled benchmark, reducing the additional value of the forecast.

Table A12: Divisible-Coverage Welfare Incidence across Adaptation Access Assignments at $\theta = 0.9$

Adaptation access	Population	Below pooled price	Pooled-affordable band	High liquidity
No access	-0.95	-1.18	-0.99	-0.80
Top 10%	-0.71	-1.18	-0.99	+0.05
Top 20% (reference)	-0.45	-1.18	-0.99	+0.93
Top tercile	-0.11	-1.18	-0.79	+1.70
Universal access	+1.70	+1.68	+1.70	+1.72

Notes: Weighted CE changes under divisible insurance, in percent of mean herd value. The liquidity groups remain fixed across access assignments: below the pooled full-package price, between that price and the maximum loss, and at or above the maximum loss. The top-quintile assignment and universal access are scenarios, not observed effective-adaptation regimes.

Table A13: Baseline Mobility Proxies and Conditional Adaptation Assignments

Mobility proxy	N	Weighted share (%)	Population ΔCE (%)	
			$\beta = 5$	$\beta = 10$
Any movement	506	47.8	-0.63	+0.30
At least 10 percent moved	506	38.6	-0.69	+0.06
At least one quarter moved	506	25.2	-0.78	-0.29
At least one half moved	506	11.8	-0.87	-0.64
Any lactating animal moved	506	18.2	-0.83	-0.48

Notes: Round-1 livestock records matched to the calibration households. Movement intensity is the fraction of recorded animal counts moved, without species-value weights. The welfare columns assign the modeled scalable adaptation technology to exactly the households satisfying each proxy and evaluate divisible insurance at $\theta = 0.9$. Entries are survey-weighted CE changes relative to pooled pricing with the same assignment, as percentages of mean herd value. They use the same calculations as the numerical macros in Section 6.4. A dash indicates that a complete-sample proxy assignment is unavailable. Assignment presumes access to the unconstrained adaptation technology; financeability diagnostics are also exported.

Table A14: Marginal-Group Welfare with Jointly Divisible Insurance and Adaptation

Marginal-group adaptation technology	θ	Band CE change
No marginal-group adaptation	0.9	-0.99
Divisible adaptation, $\beta_{\text{poor}} = 5$	0.9	-0.36
Divisible adaptation, $\beta_{\text{poor}} = 10$	0.9	+1.66

Notes: CE changes for the fixed baseline marginal band, in percent of mean herd value. Households choose both adaptation and insurance continuously within the same budget. The values $\beta_{\text{poor}} = 5$ and 10 describe alternative effectiveness assumptions; neither is estimated from realized loss reduction.

Table A15: Index Protection, Participation, and the Value of Lost Coverage

φ	Band loss $\theta = 0.7$	Band loss $\theta = 0.9$	Pooled full-package take-up
1.00	0.44	1.45	88.3
0.90	0.35	1.14	88.3
0.80	0.26	0.84	88.3
0.70	0.17	0.54	88.3
0.60	0.08	0.24	88.3
0.50	0.00	0.00	6.0
0.35	0.00	0.00	0.0

Notes: Band CE losses are percentages of mean herd value. The last column is model-predicted purchase of the full standardized package at pooled prices, not observed program take-up. Falling demand reduces the insurance value exposed to repricing.

Table A16: Robustness Summary at $\theta = 0.9$

Specification	Shortfall (%)	Affordable share (%)	Band CE loss (%)	Adaptation ΔCE (%)	Pooled take-up (%)
Baseline	54.1	72.9	1.45	+1.73	88.3
Risk aversion: half reference	54.1	72.9	0.68	+1.91	88.3
Risk aversion: twice reference	54.1	72.9	3.25	+1.34	88.3
Premium loading: 1.05	55.9	71.6	1.22	—	87.5
Premium loading: 1.10	58.2	70.3	0.95	—	86.3
Premium loading: 1.20	61.3	67.8	0.50	—	48.4
Premium loading: 1.30	63.6	65.4	0.18	—	26.6
Premium loading: 1.50	68.6	61.2	0.00	—	5.4
Premium loading: 2.00	76.8	52.5	0.00	—	0.0
Birr per livestock unit: 5,000	53.9	73.0	1.45	+1.73	88.3
Birr per livestock unit: 7,000	54.1	72.8	1.44	+1.73	87.9
Unweighted	54.0	72.8	1.43	+1.72	88.3

Notes: Loading rows apply a partial-equilibrium markup to the constrained group's standardized package; adaptation results are omitted for those rows. The final column is predicted pooled full-package take-up, not observed take-up. At fair pricing it equals affordability, whereas under loadings it also reflects willingness to pay. The reference adaptation assignment and effectiveness are not estimated.

Table A17: Baseline-Income Sensitivity at $\theta = 0.9$

Baseline income	Mean CE change	EDE, $\varepsilon = 1$	EDE, $\varepsilon = 2$
0.6	-0.56	-1.01	-1.12
1.0	-0.56	-0.93	-1.07
2.0	-0.56	-0.83	-0.97
5.0	-0.56	-0.71	-0.82

Notes: Lumpy-reference changes in percent of mean herd value. Mean CE changes are invariant to the common additive endowment under CARA; the Atkinson measures also depend on the cross-household distribution of consumption levels. These values must not be relabeled as the corresponding divisible-insurance aggregates.

Table A18: Sensitivity to Baseline Event Risk

p (%)	θ	Pr(B) (%)	p_B (%)	Full-package shortfall (%)	Mean affordable coverage (%)	ΔCE no adaptation (%)	ΔCE reference adaptation (%)	Full pooled value lost (%)
10	0.7	34.0	20.6	16.8	92.7	-0.10	-0.10	7
10	0.9	18.0	50.0	45.7	78.0	-0.58	-0.33	42
15	0.7	36.0	29.2	26.6	88.4	-0.19	-0.13	10
15	0.9	22.0	61.4	54.1	72.9	-0.95	-0.45	50
20	0.7	38.0	36.8	34.1	84.4	-0.28	-0.09	12
20	0.9	26.0	69.2	59.2	69.5	-1.27	-0.51	55
25	0.7	40.0	43.7	40.9	81.0	-0.35	-0.06	13
25	0.9	30.0	75.0	61.8	67.2	-1.52	-0.54	57
30	0.7	42.0	50.0	45.7	78.0	-0.41	-0.07	14
30	0.9	34.0	79.4	63.6	65.6	-1.70	-0.59	58

Notes: Each row varies the prior event probability p while holding the household sample, liquidity mapping, contract parameters, and forecast precision θ fixed. Welfare changes compare informative post-signal pricing with pooled pricing under the same value of p, so they measure the information-related effect rather than the direct welfare cost of a higher underlying event probability. "Reference adaptation" assigns the modeled scalable adaptation technology to the top asset quintile. "Full pooled value lost" disables adaptation and divides the population certainty-equivalent loss from repricing by the value of full pooled insurance under the same prior.

Table A19: Randomized Discounts, Take-Up, and Insured Quantity

	Take-up	Unconditional insured TLU
Randomized discount	0.451***	2.497***
	(0.079)	(0.596)
Liquidity	0.040	-0.175
	(0.035)	(0.141)
Discount–liquidity interaction	-0.047	0.831
	(0.096)	(0.719)
Risk aversion	-0.069*	-0.334
	(0.040)	(0.233)
Log herd size	-0.026	0.246**
	(0.028)	(0.116)
Woreda fixed effects	Yes	Yes
Observations	502	502
R^2	0.132	0.192

Notes: Linear probability and OLS estimates with Woreda fixed effects and HC1 standard errors in parentheses. Nonbuyers have zero insured quantity. The discount regressor is a fraction between zero and one: 0.10 denotes a 10-percentage-point reduction in the premium. Its coefficient and interaction coefficient are therefore per unit of that fraction; a 10-percentage-point change multiplies the corresponding effect by 0.10.

B Proofs

Throughout, $R = R_i = a_i + \bar{y}$ denotes non-stochastic resources. I suppress household subscripts when the comparison concerns one household, and use $\pi_s = p_s L$, $q'_G = 1 - 2p$, and $q_G \pi_G = (1 - \theta)pL$. Cross-household comparisons in Proposition 4(i) use a common risk-aversion coefficient; part (ii) instead compares each household with its own pooled benchmark.

B.1 Preliminary result: two-point mean-preserving spreads

The following observation supplies the posterior-dispersion argument used below.

Lemma B.1 (Straddling two-point lotteries). Let X take values $x_1 < x_2$ with probabilities $(q, 1-q)$, and Y take values $y_1 \le x_1 < x_2 \le y_2$ with probabilities $(r, 1-r)$, where $q, r \in (0, 1)$. If the supports differ and $\mathbb{E}[Y] = \mathbb{E}[X]$, then Y is a mean-preserving spread of X. Hence $\mathbb{E}[v(Y)] < \mathbb{E}[v(X)]$ for strictly concave v, and $\mathbb{E}[\phi(Y)] \ge \mathbb{E}[\phi(X)]$ for convex ϕ.

Proof. Let F_X, F_Y be the distribution functions and define

$$S(x) = \int_{-\infty}^{x} [F_Y(t) - F_X(t)] \, \mathrm{d}t.$$

Equal means imply $S(y_2) = 0$. On $[y_1, x_1)$, $F_Y - F_X = r \ge 0$, so $S(x_1) \ge 0$. On $[x_2, y_2)$, $F_Y - F_X = r - 1 \le 0$; working back from $S(y_2) = 0$ gives $S(x_2) \ge 0$. Between x_1 and x_2, S is linear and cannot fall below the smaller endpoint value. Outside the supports it is zero. Thus $S \ge 0$ everywhere, which is the integrated-distribution condition for a mean-preserving spread. Distinct nondegenerate supports make the concave-utility comparison strict. □

The posterior lottery has outcomes $p_G(\theta), p_B(\theta)$ and invariant mean p. Increasing precision moves its lower outcome down and its upper outcome up. Lemma B.1 therefore applies between informative precision levels. At $\theta = 1/2$ the posterior is constant, so the corresponding comparison follows directly from Jensen's inequality. The same argument applies to the premium lottery when adaptation is inactive.

B.2 Proof of Lemma 1: insurance demand

Part (i). Fix feasible e_i, write $\ell = h(e_i)$, and let $A = R_i - \psi_i(e_i)$. Coverage m gives consumption $A - p_s m$ without a loss and $A - p_s m - (\ell - m)$ with a loss. Conditional expected

consumption is $A-p_s\ell$ for every m. Only $m=\ell$ makes consumption certain. Strict concavity therefore makes full residual-loss coverage uniquely optimal when it is affordable.

If $p_s\ell > w_i - \psi_i(e_i)$, the feasible coverage interval ends below ℓ. Expected utility is strictly increasing in coverage on that interval, since each increase reduces dispersion around the same mean. The budget binds, giving $m=(w_i-\psi_i(e_i))/p_s$. Coverage exceeding ℓ creates the opposite dispersion and is dominated by $m=\ell$ whenever feasible.

Part (ii). With $e_i=0$ and $M=L$, purchasing the lumpy package yields certain consumption $R-\pi_s$. Remaining uninsured yields R and $R-L$ with mean $R-p_sL=R-\pi_s$. Because $p_s\in(0,1)$,

$$u(R-\pi_s) > (1-p_s)u(R)+p_s u(R-L).$$

The household therefore buys whenever $\pi_s \le w_i$. Otherwise no positive coverage is available on the restricted menu. □

B.3 Proof of Lemma 2: self-protection and insurance

Fix $p_s\in(0,1)$. After effort e, uninsured consumption is $R-\psi(e)$ without a loss and $R-\psi(e)-h(e)$ with a loss. CARA utility gives

$$\mathbb{E}[u(z)] = -e^{-\gamma[R-\psi(e)]}\big[1-p_s+p_s e^{\gamma h(e)}\big].$$

The certainty equivalent is consequently

$$\mathrm{CE}(e) = R-\psi(e)-\frac{1}{\gamma}\ln\big(1-p_s+p_s e^{\gamma h(e)}\big).$$

Subtracting it from expected consumption produces the willingness to pay above the fair premium:

$$\rho^{\mathrm{res}}(e;p_s) = \frac{1}{\gamma}\ln\big(1-p_s+p_s e^{\gamma h(e)}\big) - p_s h(e).$$

Its derivative is

$$\frac{\partial \rho^{\mathrm{res}}(e;p_s)}{\partial e} = p_s h'(e)\frac{(1-p_s)(e^{\gamma h(e)}-1)}{1-p_s+p_s e^{\gamma h(e)}} < 0.$$

The fraction is positive and $h'(e)<0$. The decline is driven by the smaller physical loss, not by an income effect on risk aversion. □

B.4 Proof of Propositions 1 and 2: affordability

Differentiating the adverse posterior in equation (2) gives

$$p'_B(\theta) = \frac{p(1-p)}{[\theta p + (1-\theta)(1-p)]^2} > 0.$$

Because $\pi_B = Mp_B$, both the premium and the full-package liquidity threshold rise strictly with precision. This proves Proposition 1.

Under the stylized mapping, a non-adapting household can finance the package if and only if $a_i + \bar{b} \geq \pi_B$. Hence $a^*(\theta) = Mp_B(\theta) - \bar{b}$ and

$$\frac{\mathrm{d}a^*}{\mathrm{d}\theta} = Mp'_B(\theta) > 0.$$

The share below this threshold is nondecreasing in precision, and strictly increasing wherever the asset distribution has positive density at the threshold. This proves Proposition 2. □

B.5 Proof of Lemma 3: unconstrained welfare

For a household able to finance the optimum, Lemma 1 implies $m = h(e)$. Conditional consumption is certain and the problem reduces to

$$f(\mu) = u(R - g(\mu)), \qquad g(\mu) = \min_{e \geq 0}\{\psi(e) + \mu h(e)\}.$$

The minimum of affine functions of μ is concave; g is nondecreasing because $h \geq 0$. With a unique minimizer, the envelope theorem gives $g'(\mu) = h(e^*(\mu))$. At a smooth interior optimum,

$$\psi'(e^*) + \mu h'(e^*) = 0, \qquad \frac{\mathrm{d}e^*}{\mathrm{d}\mu} = \frac{-h'(e^*)}{\psi''(e^*) + \mu h''(e^*)} > 0,$$

and therefore

$$g''(\mu) = -\frac{h'(e^*)^2}{\psi''(e^*) + \mu h''(e^*)} < 0.$$

These derivatives show how adaptation limits the increase in expenditure when expected losses rise.

Part (i). With adaptation inactive throughout the posterior range, $g(\mu) = \mu L$. Thus $f(\mu) = u(R - \mu L)$ satisfies $f''(\mu) = u''(R - \mu L)L^2 < 0$. Greater precision spreads posterior beliefs around the same mean, so Lemma B.1 and Jensen's inequality imply a strict decline

in $\mathbb{E}_s[f(p_s)]$.

Part (ii). Fix the θ_2 and posterior interval in the lemma. At a smooth interior solution,

$$\begin{aligned} f'(\mu) &= -g'(\mu)u'(R - g(\mu)), \\ f''(\mu) &= u''(R - g(\mu))g'(\mu)^2 - u'(R - g(\mu))g''(\mu) \\ &= u'(R - g(\mu))[-\gamma g'(\mu)^2 - g''(\mu)]. \end{aligned}$$

The last equality uses CARA. Consequently, f is convex wherever $\gamma \leq -g''/g'^2$. Substitution of the envelope derivatives yields equation (11). If that condition holds throughout $[p_G(\theta_2), p_B(\theta_2)]$, the nested posterior ranges imply nondecreasing welfare on $[1/2, \theta_2]$.

Part (iii). Define $\rho^{\exp}(\theta)$ through

$$\mathbb{E}_s[u(R - g(p_s))] = u\big(R - \mathbb{E}_s[g(p_s)] - \rho^{\exp}(\theta)\big).$$

Jensen's inequality gives $\rho^{\exp}(\theta) \geq 0$. Since u is increasing and pooled welfare is $u(R - g(p))$,

$$W_U(\theta) \geq W_U(1/2) \quad \Longleftrightarrow \quad g(p) - \mathbb{E}_s[g(p_s)] \geq \rho^{\exp}(\theta).$$

The left side is nonnegative by concavity of g and measures expected expenditure savings. The right side prices expenditure risk. This establishes equation (12). □

B.6 Proof of Proposition 3: the marginally insured

Keep adaptation inactive and restrict coverage to $\{0, L\}$. A household with $w \in [pL, L)$ buys if and only if $\pi_s \leq w$. Since π_B increases continuously from pL to L, a unique $\bar{\theta}_i \in [1/2, 1)$ satisfies $\pi_B(\bar{\theta}_i) = w$. The household always affords favorable-signal coverage, because $\pi_G < pL$ for informative signals.

The invariant mean. If the household buys after both signals, expected consumption is $R - [q_B p_B + q_G p_G]L = R - pL$. If it exits after adverse news, then

$$\mathbb{E}[z] = R - q_G \pi_G - \theta pL = R - (1 - \theta)pL - \theta pL = R - pL. \tag{30}$$

Equation (30) makes clear that the welfare changes in this proposition concern risk around a fixed mean.

Part (i). While the household buys after both signals, its consumption lottery has outcomes $R - \pi_B$ and $R - \pi_G$. Higher precision spreads these outcomes while preserving their mean. Strict concavity and Lemma B.1 imply decreasing welfare on the stated interval.

Part (ii). At the threshold, the household buys and consumes $R - w$ after adverse news. Just above it, the same branch becomes uninsured, with limiting consumption R or $R - L$ and loss probability w/L. Its mean is still $R - w$, but strict concavity gives

$$J_i = u(R - w) - [(1 - w/L)u(R) + (w/L)u(R - L)] > 0.$$

Only this branch changes discontinuously. Multiplication of J_i by $q_B(\bar{\theta}_i)$ yields equation (13).

Part (iii). In the exclusion region,

$$\begin{aligned} W(w;\theta) = q_G u(R - \pi_G) + (1 - \theta)(1 - p)u(R) \\ + \theta p\, u(R - L). \end{aligned} \tag{31}$$

Using $q_G' = 1 - 2p$ and $q_G \pi_G' = -pL - (1 - 2p)\pi_G$, differentiation of equation (31) gives

$$\begin{aligned} \frac{\partial W}{\partial \theta} = (1 - 2p)u(R - \pi_G) + [pL + (1 - 2p)\pi_G]u'(R - \pi_G) \\ - (1 - p)u(R) + p\, u(R - L). \end{aligned} \tag{32}$$

Substitute CARA utility into equation (32). With temporary notation $t = \gamma \pi_G$ and $\eta = \gamma L$, the sign is that of

$$\varphi(t) = e^t[(1 - 2p)(t - 1) + p\eta] + (1 - p) - pe^{\eta}, \qquad 0 < t < p\eta.$$

For $p \leq 1/2$,

$$\varphi'(t) = e^t[(1 - 2p)t + p\eta] > 0.$$

It remains to sign its upper endpoint. Define

$$G(\eta) = \varphi(p\eta) = e^{p\eta}[2p(1 - p)\eta - (1 - 2p)] + (1 - p) - pe^{\eta}.$$

Then $G(0) = 0$ and

$$G'(\eta) = -pe^{p\eta}H(\eta),$$
$$H(\eta) = e^{(1-p)\eta} - 2p(1-p)\eta - 1.$$

Since $H(0) = 0$ and $H'(\eta) = (1-p)[e^{(1-p)\eta} - 2p] > 0$ for $\eta > 0$ and $p \leq 1/2$, $G'(\eta) < 0$ and $G(\eta) < 0$. Thus $\varphi(t) < \varphi(p\eta) < 0$, establishing the stated local monotonicity. The global comparison below does not require this restriction on the prior.

Part (iv). At pooled pricing, full insurance gives certain consumption $R - pL$. At any informative precision, the household has a nondegenerate consumption lottery with the same mean by equation (30) and its insured-region counterpart. Jensen's inequality gives

$$W(w;\theta) = \mathbb{E}[u(z)] < u(R - pL) = W(w;1/2).$$

This proves the global ranking for every $p \in (0,1)$. □

B.7 Proof of Proposition 4: distributional incidence

Part (i). Fix $\theta \in (\bar{\theta}_C, 1)$, inactive adaptation, common L, and $\gamma_U = \gamma_C = \gamma$. Subtract each household's non-stochastic resources. Under CARA this only shifts its certainty equivalent, so it leaves the information-related change unaffected. The unconstrained net lottery is

$$X_U = \begin{cases} -\pi_B & \text{with probability } q_B, \\ -\pi_G & \text{with probability } q_G. \end{cases}$$

The constrained household replaces the first branch by a lottery with coverage $m_B = w_C/p_B < L$:

$$X_C = \begin{cases} -\pi_G & \text{with probability } q_G, \\ -p_B m_B & \text{with probability } q_B(1-p_B), \\ -p_B m_B - (L - m_B) & \text{with probability } q_B p_B. \end{cases}$$

The substituted branch has mean $-p_B m_B - p_B(L - m_B) = -\pi_B$. It is nondegenerate because $m_B < L$ and $p_B < 1$. Thus X_C is a mean-preserving spread of X_U, and the common utility function gives $\mathrm{CE}(X_C) < \mathrm{CE}(X_U)$.

At pooled pricing both net certainty equivalents equal $-pL$. Consequently,

$$\Lambda_C = -pL - \mathrm{CE}(X_C) > -pL - \mathrm{CE}(X_U) = \Lambda_U > 0.$$

The final inequality follows from the unconstrained premium lottery. Under the lumpy restriction $m_B = 0$, the additional spread is larger. At perfect information $p_B = 1$, the conditional spread vanishes and the two net lotteries coincide; this explains the open upper endpoint in equation (14). With unequal risk aversion, the lottery ordering alone does not order household-specific CE losses.

Part (ii). Under the interiority and regularity conditions of Lemma 3(ii), for every $\theta \in (1/2, \theta_2]$ that lemma gives $\mathrm{CE}_U(\theta) \geq \mathrm{CE}_U(1/2)$. Corollary 1 gives $\mathrm{CE}_C(\theta) < \mathrm{CE}_C(1/2)$. Subtracting the second comparison from the first establishes equation (15). Each comparison uses the household's own preferences, so equal risk-aversion coefficients are not required in this part. No cardinal utility levels are compared across households. □

B.8 Proof of Proposition 5: prepayment and commitment

Fix $\lambda \in (0, 1]$ and $\theta > 1/2$. A non-adapting household with $w_i \geq pL$ commits to indemnity λL, pays d before the forecast, and owes $\lambda pL - d$ afterward. Favorable-signal full replacement is affordable because $p_G < p$.

Part (i). The deposit is sunk whether the household honors or cancels. Honoring costs the remaining balance; canceling costs $\lambda p_G L$ to replace the same locked block. Thus honoring is weakly preferred if and only if

$$\lambda pL - d \leq \lambda p_G L \quad \Longleftrightarrow \quad d \geq \lambda L(p - p_G).$$

The same coverage is delivered under both choices. Risk aversion does not affect this comparison, while liquidity matters through replacement feasibility. Adverse-signal buyers prefer honoring because $p_B > p$. The monotonicity and endpoints follow from $p_G' < 0$, $p_G(1/2) = p$, and $p_G(1) = 0$. Equality gives indifference, not strict retention.

Part (ii). Suppose every favorable-signal buyer lapses and forfeits d, while adverse-signal buyers pay the full locked premium. Expected receipts and indemnities on the locked block

satisfy the break-even condition

$$q_G d + q_B \lambda p L = \theta p \lambda L.$$

It gives

$$d = \frac{\lambda L p(\theta - q_B)}{q_G} = \frac{\lambda L p(1-p)(2\theta - 1)}{q_G} = \lambda L(p - p_G),$$

where the last equality uses $q_G - (1-\theta) = (1-p)(2\theta - 1)$. This equals the household participation threshold. The equality is specific to fair pricing. With a common multiplicative loading $\nu > 1$, the household threshold is νd^*_λ while the insurer's nonnegative break-even threshold is $\max\{0, \lambda L p[\theta - \nu q_B]/q_G\}$.

Part (iii). When the deposit is drawn from date-1 resources, the remaining budget is $w_i - d$. The remaining full-protection bill is the locked balance plus the spot top-up. Therefore

$$\lambda p L - d + (1-\lambda)\pi_s \leq w_i - d \quad \Longleftrightarrow \quad \lambda p L + (1-\lambda)\pi_s \leq w_i.$$

This proves equation (25). The additional condition is availability of the deposit before the signal, $d \leq \ell^0_i$; satisfying it alone need not fund all later premiums. The equivalence abstracts from interest, intervening income, and credit-line renewal. Separate pre-signal resources would change the date-1 constraint. □