Fraud, Transaction Costs, and Artificial Intelligence in Livestock Insurance: Theory and Evidence from Cattle Facial Recognition in China*

(Very preliminary. Please do not circulate or cite without the authors' permission.)

Honglin Li† Yiming Cao‡ Fangzhou Lu§ Hao Zheng¶

June 5, 2026

Abstract

This paper studies how artificial intelligence affects fraud, transaction costs, and market coverae in livestock insurance. We examine the adoption of AI-based cattle facial recognition in a large Chinese livestock insurance market, where traditional ear-tag verification is costly, labor intensive, and vulnerable to manipulation. We develop a theoretical framework in which lack of identity verification induces fraudulent claims on uninsured cattle that distorts insurers' marginal costs. This distortion generates a downward-sloping marginal cost curve that mimics adverse selection, leading to inefficiently under-enrollment at the household level in equilibrium. In addition, high

*We thank Justin Sydnor, Stuart Craig, Phil Mulder, Dan Sacks, Ty Leverty, Kyohei Okumura, and the participants at the R&I Workshop at the UW-Madison for helpful comments and suggestions. All remaining errors are our own.

†University of Wisconsin-Madison. Email: hli434@wisc.edu.

‡University of Hong Kong. Email:ymcao@hku.hk.

§University of Hong Kong. Email: lufz@hku.hk.

¶University of International Business and Economics. Email: zhenghao@uibe.edu.cn.

transaction costs impose capacity constraints, limiting the scale of insurance provision. AI reduces this identity-verification friction and lowers underwriting transaction costs, relaxing capacity constraints. Using administrative cow-level data from 2021–2025 from a single insurer, covering 1.94 million cow-year observations and about 48,000 policyholders, we document the large-scale transition from traditional ear-tag verification to AI-based cattle facial recognition. Comparing early AI-exposed policyholders with later-exposed policyholders, we show that AI reduces claim rates by 14 percent and increases the average insured herd size by about 5 percent, consistent with our theoretical predictions. We also present preliminary evidence that AI lowers underwriting time per cow and increases underwriter productivity, allowing underwriters to process more cattle and reach more distant farmers. These findings suggest that reduced transaction costs relax operational constraints and expand insurance provision.

1 Introduction

Livestock are economically important assets worldwide. They generate income, provide food and nutrition, serve as stores of wealth and collateral, and act as informal insurance against shocks (Food and Agriculture Organization of the United Nations, 2009; World Bank, 2009). Globally, livestock support the livelihoods of approximately 1.3 billion people, account for roughly 40 percent of agricultural output, and provide more than one-third of the world's protein supply (Food and Agriculture Organization of the United Nations, 2009). Yet these assets are also exposed to substantial risks, especially mortality risk, which can impose large financial losses on households and farms.

In principle, financial contracts around livestock—including insurance, collateralized lending, and broader agricultural risk-management products—can help households manage these risks and unlock the value of biological assets. In practice, however, contracting on livestock is difficult because animals are hard to identify, verify, and monitor over time. This reflects a broader problem in developing-market finance: many contracts fail not because assets lack economic value, but because weak verification and enforcement make productive assets difficult to pledge, insure, or monitor (Besley, 1995; Giné and Yang, 2009; Karlan et al., 2014a). In livestock mortality insurance, this friction is particularly acute because insurers must verify whether a deceased animal is the same animal that was originally insured. Traditional identification technologies, such as ear tags and physical records, are costly, labor-intensive, and vulnerable to manipulation or loss. Recent advances in AI-based biometric identification may help overcome these frictions by improving verification and lowering transaction costs.

This paper studies the impact of AI, specifically cattle facial recognition, on fraud, transaction costs, and market outcomes in livestock insurance. We first develop a theoretical framework that formalizes the central contracting friction: imperfect and costly animal identity verification. When insurers cannot perfectly verify whether a claimed dead animal is the same animal that was insured, farmers may be able to claim losses on uninsured cattle. This identity friction distorts insurers' marginal costs because contracts are no longer perfectly tied to specific animal identities: a farmer can use a limited number of insurance contracts to claim losses on deaths occurring among a larger herd.[1] As a result, the marginal

[1] Identity verification is also necessary for accurately determining the number of cattle covered by insurance. Without the ability to uniquely identify individual animals, insurers cannot reliably determine which cows have already been insured, allowing farmers to conceal part of their herd during underwriting. As a

cost of the nth insurance contract is not simply the actuarially fair cost, but depends on the probability that the herd experiences at least n deaths. This generates a downward-sloping marginal cost curve that mimics adverse selection, even though cows are homogeneous in underlying mortality risk. The distortion therefore differs from standard adverse selection (e.g., Akerlof 1970; Einav et al. 2010): it arises not from unobserved risk heterogeneity, but from the insurer's inability to verify animal identity. The model also highlights a related friction from the same verification process: traditional underwriting is costly and time-intensive, imposing capacity constraints on underwriters and limiting the scale of insurance provision. A low-cost and accurate identity technology reduces both frictions. By tying coverage and claims to specific cows, AI reduces fraudulent claims on uninsured animals and under-enrollment; by simplifying verification, it lowers underwriting costs and relaxes operational capacity constraints.

Guided by these theoretical predictions, we study the empirical implications of adopting AI-based cattle facial recognition at a large livestock insurer in China. The setting is useful because it combines a sharp change in verification technology with rich administrative data and meaningful variation in exposure to the new technology. Under the traditional technology, underwriters visit farms, physically inspect cattle, restrain animals, attach or scan ear tags, and record identifying information. The AI technology instead uses computer vision to extract facial features from smartphone images and compare them with images in the insurer's database. The insurer began with a small pilot in 2022, made AI broadly available to underwriters in 2023, and moved to near-universal AI use in 2024 and 2025. This staggered transition generates variation in AI exposure across cows, households, underwriters, and years. This rollout generates variation in AI exposure across cows, households, underwriters, and years.

We combine administrative underwriting and claims data covering more than 2 million cow-year observations and approximately 48,000 policyholders with underwriter characteristics and geographic information. The data allow us to observe which cattle were processed using AI, whether and when they were claimed, insured herd size at the policyholder level, the underwriters responsible for enrollment, and the geographic reach of insurance provision. These linked data allow us to study both fraud and transaction costs through claims, herd size, underwriter productivity, and the distance of covered households.

result, insurers face an under-enrollment problem in which some animals remain uninsured while still potentially benefiting from weak identity verification at the claims stage.

We first document that AI exposure in 2023 is highly polarized at the policyholder level: about 44 percent of policyholders have no AI-processed cattle, while about 51 percent have all insured cattle processed with AI, leaving limited partial exposure. We then estimate a household-level event-study design based on each household's AI-processed share in 2023. This is not a standard staggered-adoption design. Instead, we classify households according to whether they were treated based on their early AI exposure in 2023, before AI processing became nearly universal. Both the pre-treatment and later-year estimates provide useful validation tests. The absence of differential effects before adoption would support the identifying assumption that households with different levels of 2023 AI exposure were on similar trajectories prior to treatment. Conversely, later-year estimates serve as a placebo-style check: as nearly all households become AI treated, the effect of initial 2023 exposure should diminish if the estimates are driven by differential AI treatment rather than underlying household differences.

We begin by examining the fraud channel through household claim rates. The event-study evidence shows no differential pre-trend in 2021. In 2023, AI exposure is associated with a 0.77 percentage-point decline in household claim rates, corresponding to a 13.95 percent reduction relative to the baseline. The effect becomes close to zero and statistically insignificant in 2024, consistent with convergence as AI becomes widespread. This pattern supports the model's fraud channel: AI reduces claims when it differentially improves identity verification, and the difference fades once verification technology becomes nearly universal.

We next examine insured herd size. Empirically, households with greater AI exposure in 2023 increase insured herd size by about 0.81 cows relative to 2022, a 4.84 percent increase relative to the baseline. The effect is concentrated in 2023 and diminishes in later years, again consistent with a transition-year response as AI relaxes constraints before becoming broadly available.

Finally, we study the productivity and distributional implications of lower transaction costs at the underwriter level. Preliminary time-stamp evidence suggests that AI reduces average underwriting time per cow from about 3 minutes under traditional underwriting to about 1 minute under AI-based underwriting, a decline of roughly 70 percent. Underwriters with higher 2023 AI exposure also serve more policyholders after adoption. In a balanced sample of continuing underwriters, high-AI underwriters serve roughly 27 ad-

ditional policyholders in 2023 relative to 2022. In the full underwriter-year panel, which includes entrants and exiters, the effect is about 20 additional policyholders in 2023 and remains positive in 2024. We also find evidence that AI helps underwriters reach more distant households: in the full sample, higher 2023 AI exposure is associated with serving households about 3 kilometers farther away in 2023, with positive effects persisting in later years. These findings suggest that AI relaxes operational constraints not only by making incumbent underwriters more productive, but also by changing market composition and expanding access to more remote policyholders.

Together, the results show that AI-based identity verification can reshape livestock insurance markets through two channels. First, it reduces identity-based fraud by making claims verifiable at the animal level. Second, it lowers transaction costs, relaxes underwriting capacity constraints, and expands coverage. The evidence highlights a broader role for AI in agricultural finance: by solving identity and verification problems, AI can make financial contracts around livestock more enforceable, scalable, and inclusive.

Literature Review This paper contributes to two strands of literature. A large literature studies how asymmetric information distorts financial and insurance markets, including adverse selection (Akerlof, 1970; Rothschild and Stiglitz, 1976) and moral hazard (Pauly, 1968; Shavell, 1979). Related work in agricultural insurance emphasizes similar frictions, including adverse selection, moral hazard, basis risk, and high administrative costs (Skees and Reed, 1986; Nelson and Loehman, 1987; Miranda, 1991; Smith and Goodwin, 1996). Our paper builds on this literature but emphasizes a *distinct* friction: imperfect verification of the insured object. In our setting, the insurer may be unable to verify whether a claimed dead animal is the same animal that was insured. This identity-verification problem creates scope for claims on uninsured animals, duplicate claims, and other forms of fraud, relating our paper to work on costly state verification and insurance fraud (Townsend, 1979; Picard, 1996; Crocker and Morgan, 1998).

This friction distorts marginal costs even when cows are homogeneous in underlying mortality risk. Because a contract on one cow may effectively cover deaths from a broader herd, the marginal cost of coverage can resemble the downward-sloping cost patterns often associated with adverse selection. The mechanism, however, is different: it arises from imperfect identification of the insured object rather than heterogeneity in risk types. Livestock insurance makes this friction particularly salient: unlike crop insurance, where verification

often concerns realized yields or weather-related losses, livestock mortality insurance requires verifying both death and animal identity.

A growing literature studies how digital technology and AI affect financial markets. Existing work shows that technology can improve portfolio choices through robo-advising (D'Acunto et al., 2019), reduce processing frictions in mortgage lending (Fuster et al., 2019), create value through financial innovation (Chen et al., 2019), and expand credit access and entrepreneurial growth (Hau et al., 2024). Closest to our setting, Liu (2024) studies an experiment in insurance distribution and shows that AI-generated demand predictions improve agents' sales by shifting attention toward high-intent consumers. This literature demonstrates that digital technologies can reshape financial markets by improving information processing, automating decisions, and lowering transaction costs. Much less is known about how AI changes insurance markets, especially agricultural and livestock insurance markets where verification of the insured asset is itself a central friction.

Our paper extends this literature to a setting in which AI changes the feasibility and cost of verifying the insured object. Recent work in computer vision and precision livestock farming shows that deep learning can identify individual cattle using facial images, coat patterns, and video data (Andrew et al., 2021; Qiao et al., 2021). The Chinese cattle facial-recognition setting is particularly valuable because the technology is deployed in an actual insurance market. The empirical setting therefore allows us to study whether a technical improvement in animal identification translates into economically meaningful changes in claims, underwriting productivity, household coverage, and the spatial reach of insurance. This helps bridge the gap between engineering evidence on biometric identification and economic evidence on market design.

Finally, our paper speaks to financial inclusion and missing markets in developing economies. Livestock are among the most important assets held by rural households, yet markets that rely on these assets remain underdeveloped. A large literature argues that weak property rights, limited verifiability, and high monitoring costs constrain access to insurance, credit, and other financial services in developing countries, contributing to persistent missing markets and inefficient risk allocation (Stiglitz and Weiss, 1981; Townsend, 1994; Banerjee and Duflo, 2010). These frictions are particularly severe for biological and movable assets such as livestock, whose ownership and identity are difficult to verify at scale. By making livestock identity verifiable at low cost, AI can make these assets more con-

tractible and expand access to insurance, collateralized credit, and other risk-management products. Expanding such markets could generate substantial welfare gains. Improved access to insurance can encourage households to undertake higher-return but riskier investments (Karlan et al., 2014b; Cole et al., 2013), while greater access to collateralized credit can relax financing constraints and support productive investment (Banerjee and Newman, 1993; Buera et al., 2011). More broadly, our results suggest a role for AI in development finance: reducing missing markets by lowering the cost of contracting on valuable but hard-to-verify household assets, thereby improving risk sharing, investment, and economic resilience among rural households.

2 Conceptual Framework

This section develops a simple framework to clarify how AI-based cattle facial recognition affects livestock insurance markets. The model emphasizes two frictions. First, when the insurer cannot perfectly identify whether a dead cow is the same cow that was insured, farmers may insure only part of their herd while effectively claiming against deaths from the entire herd. This identity friction creates scope for insurance fraud and distorts the marginal cost of coverage. Second, identifying and underwriting cattle is costly. Underwriters must travel to farms, physically identify cattle, verify ownership and condition, and record identifying information. These transaction costs limit underwriting capacity and restrict market expansion.

AI-based cattle facial recognition can relax both frictions. By improving animal identity verification, AI reduces the ability to claim on uninsured cattle. By lowering the cost of underwriting and verification, AI also increases underwriter productivity and expands coverage.

2.1 Environment

Consider a farmer who owns N cows. Cow i dies with probability π_i, independently across cows. If cow i dies, the insurer pays indemnity L. The marginal cost of underwriting cow i by underwriter u is

$$MC_i = \pi_i L + c_i^u, \tag{1}$$

where $\pi_i L$ is the actuarially fair expected claim cost and c_i^u is the transaction cost of underwriting cow i by underwriter u. The transaction cost captures travel time, labor costs,

physical verification costs, and the administrative cost of recording identifying information.

Let $r \in \{T, A\}$ denote the identification technology. Technology T denotes the case in which the insurer cannot identify cow identity. Technology A denotes the case in which the insurer can perfectly identify cow identity, for example with the assistance of AI-based cattle facial recognition.

The distinction between T and A is a simplifying benchmark. In practice, pre-AI identification technologies may provide some information about animal identity, but they are costly to use and vulnerable to manipulation. The model captures the central economic difference between an environment with imperfect and costly identity verification and an environment with accurate and lower-cost identity verification.

To isolate the role of identity verification, suppose first that cows are homogeneous: $\pi_i = \pi \quad$ for all i. Let $X \sim \text{Binomial}(N, \pi)$ denote the total number of deaths in the herd.

2.2 Information/Indentify Frictions

2.2.1 Identity Frictions When Cow Identity Cannot Be Perfectly Verified

Suppose first that the insurer cannot perfectly identify cow identity at the time of claim. In this case, insurance contracts may not be effectively tied to specific animal identities. Thus, a farmer who owns N cows may purchase only $n < N$ insurance contracts and still use those contracts to claim against deaths from the entire herd. In this environment, n indexes the number of insurance contracts purchased, not the identity of a specific cow.

If the farmer buys n contracts, the insurer pays for up to the first n deaths in the herd. The insurer's expected claim cost is

$$TC^T(n) = \sum_{j=0}^{n-1} \Pr(X = j) \cdot jL + \Pr(X \geq n) \cdot nL. \tag{2}$$

The first term captures states in which fewer than n cows die, so the insurer pays only for the realized number of deaths. The second term captures states in which deaths exceed the number of insurance contracts, so the insurer pays the maximum amount nL.

The claim marginal cost of the nth contract is

$$mc_n^T = TC^T(n) - TC^T(n-1) = \Pr(X \geq n)L. \tag{3}$$

Ignoring transaction costs, the marginal cost is

$$MC_n^T = mc_n^T + c_n^{u,T} = \Pr(X \geq n)L + c_n^{u,T}. \tag{4}$$

Equation (3) is the key mechanism. When cow identity cannot be perfectly verified, the marginal claim cost of the nth contract is not πL. Instead, it is the probability that the herd experiences at least n deaths, multiplied by the indemnity. Since $\Pr(X \geq n)$ declines with n, the marginal cost curve is downward sloping.

This downward-sloping marginal cost does not arise from standard adverse selection. Cows are homogeneous in the model. The downward slope arises because the first few insurance contracts can be used to cover deaths from the entire herd. When only one cow is insured, the contract effectively pays if any cow in the herd dies. When two cows are insured, the second contract pays only if at least two cows die, and so on. Identity frictions therefore make partial insurance disproportionately costly for the insurer.

2.2.2 Accurate Identification (Assisted by AI)

AI-based cattle facial recognition changes the claim technology. When cow identity can be perfectly verified, an insurance contract is tied to the identity of a specific insured cow. Therefore, ignoring transaction costs, the claim marginal cost of an additional insured cow becomes

$$MC_n^A = \pi L. \tag{5}$$

Comparing equations (4) and (5) shows the fraud-reduction channel:

$$\Pr(X \geq n)L \longrightarrow \pi L. \tag{6}$$

AI removes the value of using insured contracts to claim on uninsured cattle. As a result, AI lowers paid claim rates by limiting identity-based fraud. This effect need not reflect a decline in underlying mortality risk.

2.2.3 Willingness to Pay

Let the farmer have wealth w and utility $u(\cdot)$. When cow identity cannot be perfectly verified, if the farmer buys n contracts and pays total willingness to pay $WTP^T(n)$, the contracts cover the first n deaths in the herd. The farmer's willingness to pay solves

$$\sum_{j=0}^{n-1} \Pr(X=j)u\left(w-WTP^{T}(n)\right)+\sum_{j=n}^{N} \Pr(X=j)u\left(w-(j-n)L-WTP^{T}(n)\right)$$
$$=\sum_{j=0}^{N} \Pr(X=j)u(w-jL). \tag{7}$$

The right-hand side is expected utility without insurance. The left-hand side is expected utility when the farmer buys n contracts. If fewer than n cows die, all losses are covered. If more than n cows die, the farmer bears only the residual loss $(j-n)L$.

Now suppose that cow identity can be perfectly verified, for example with the assistance of AI-based cattle facial recognition. If the farmer insures n cows, only deaths among those specific insured cows are covered. Deaths among the remaining $N-n$ uninsured cows are not reimbursed. Let

$$\widetilde{X} \sim \text{Binomial}(N-n,\pi) \tag{8}$$

denote deaths among uninsured cows. Because deaths among insured cows are fully indemnified, the farmer's residual risk comes only from the $N-n$ uninsured cows. The willingness to pay for insuring n identified cows, $WTP^{A}(n)$, solves

$$\sum_{j=0}^{N-n} \Pr(\widetilde{X}=j)u\left(w-jL-WTP^{A}(n)\right)=\sum_{j=0}^{N} \Pr(X=j)u(w-jL). \tag{9}$$

For the simulation and figure, define marginal willingness to pay as

$$MWTP^{r}(n)=WTP^{r}(n)-WTP^{r}(n-1), \qquad r\in\{T,A\}. \tag{10}$$

Equilibrium coverage can then be characterized by the largest n such that

$$MWTP^{r}(n)\geq MC_{n}^{u,r}. \tag{11}$$

This condition compares the farmer's incremental value of the nth contract with the insurer's full marginal cost of supplying that contract.

The left panel of Figure 1 compares marginal cost and willingness to pay when cow identity cannot be perfectly verified and when cow identity can be perfectly verified with

AI assistance. When identity cannot be perfectly verified, a small number of contracts can cover deaths from the entire herd. As a result, the claim marginal cost of the nth contract is $\Pr(X \geq n)L$, which declines with n. This generates a downward-sloping marginal cost curve even though cows are homogeneous. AI removes this distortion by making cow identity verifiable.

2.3 Transaction Costs and Capacity

The second channel operates through transaction costs. Identifying and underwriting cattle requires underwriters to travel to farms, physically inspect animals, verify ownership and condition, and record identifying information. These tasks are labor intensive and time consuming. AI-based cattle facial recognition lowers the cost of verification, so

$$c_n^{u,A} < c_n^{u,T}. \tag{12}$$

Let K^u denote the fixed time or effort capacity of underwriter u. Under identification technology r, underwriter u can serve a set of insured cows $\mathscr{I}_u^r$ subject to the capacity constraint

$$\sum_{i \in \mathscr{I}_u^r} c_i^{u,r} \leq K^u. \tag{13}$$

Because AI lowers per-cow transaction costs, the same underwriter can process more cows under accurate identification. Thus, even if AI does not change the underlying mortality risk of insured cows, it can increase insurance coverage by relaxing the underwriter's capacity constraint.

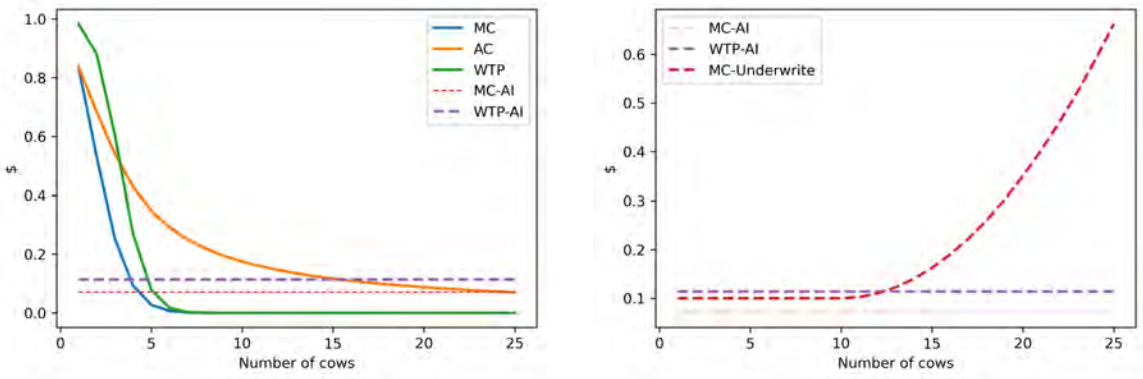
The right panel of Figure 1 isolates the transaction-cost channel. When cow identity is verifiable but underwriting is costly, marginal cost is shifted upward by transaction costs. AI lowers transaction costs from $c_n^{u,T}$ to $c_n^{u,A}$, shifting the marginal cost curve downward and increasing equilibrium coverage. This effect reflects improved underwriting productivity rather than a change in underlying mortality risk.

Combining the fraud and productivity channels, AI reduces marginal cost through two components:

$$\underbrace{\Pr(X \geq n)L \to \pi L}_{\text{identity verification and fraud reduction}} \quad \text{and} \quad \underbrace{c_n^{u,T} \to c_n^{u,A}}_{\text{lower transaction costs}}. \tag{14}$$

The first channel affects paid claims. The second channel affects underwriting capacity and market expansion.

Figure 1: Insurance Fraud, Transaction Costs, and Underwriting Capacity



(a) Identity frictions and fraud

(b) Transaction costs and capacity

Notes: This figure illustrates the model's equilibrium implications. Panel A plots marginal willingness to pay and marginal cost as functions of the number of insured cows n. In the ear-tag regime, animal identity is imperfectly verified, so the marginal claim cost of the nth contract is $\Pr(X \geq n)L$. In the AI regime, animal identity is accurately verified, so the marginal claim cost is πL. Panel B isolates the transaction-cost channel. When underwriting requires costly in-person verification, marginal cost is shifted upward by c_n^u. AI lowers transaction costs from $c_n^{u,E}$ to $c_n^{u,A}$ and expands equilibrium coverage. The simulation uses $\pi = 0.07$, $L = 1$, and CARA utility $u(m) = -(1/\gamma)\exp(-\gamma m)$ with $\gamma = 0.0007$.

2.4 Empirical Predictions

The framework delivers three empirical predictions.

First, AI reduces claim rates and insured sizes per policyholder by improving identity verification. When cow identity cannot be perfectly verified, farmers can use a limited number of insured contracts to claim against deaths from the entire herd. AI ties claims to the identity of specific insured cows, limiting claims on uninsured cattle. Therefore, AI adoption should reduce observed claim rates and expected payouts. This effect reflects a reduction in identity-based fraud, not necessarily a reduction in true mortality risk.

Second, AI increases coverage by lowering transaction costs and relaxing underwriting capacity constraints. Because $c_i^{u,A} < c_i^{u,T}$, underwriters can process more cows with the same time and effort capacity. Therefore, AI should increase the number of insured cows and insured households.

3 Background and Setting

3.1 Livestock Assets and Mortality Risk

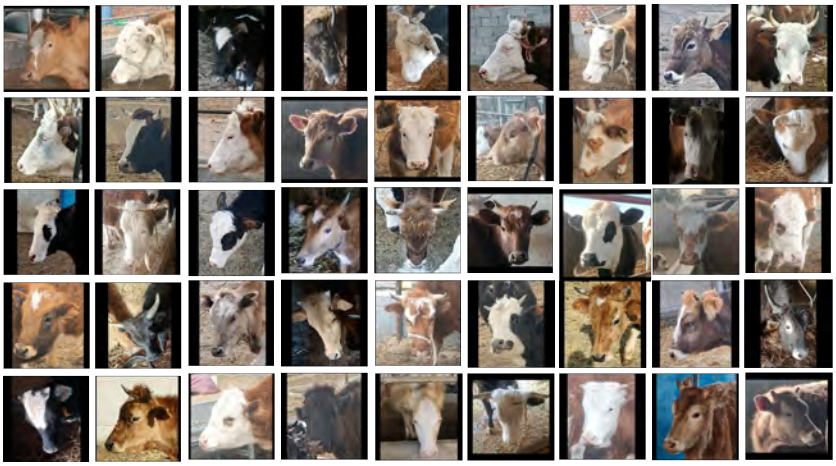
Animals raised for agricultural production play a central role in rural economies around the world. Beyond supplying meat, milk, and other animal products, livestock contribute to household earnings, function as a form of savings, provide collateral for borrowing, and help families cope with adverse economic shocks (Food and Agriculture Organization of the United Nations, 2009; World Bank, 2009). The sector is globally significant: livestock-related activities sustain the livelihoods of about 1.3 billion people, generate close to 40 percent of total agricultural output, and supply more than one-third of the protein consumed worldwide (Food and Agriculture Organization of the United Nations, 2009). Their importance is particularly pronounced in developing regions, where a large portion of household wealth is often held in livestock and other productive assets that cannot be easily converted into cash.

Livestock are also risky assets. Farmers face a wide range of risks, including price fluctuations, production shocks, animal diseases, epidemic outbreaks, natural disasters, and mortality risk (Food and Agriculture Organization of the United Nations, 2009; World Bank, 2009). Mortality risk is particularly consequential because it destroys the productive asset itself and can generate persistent reductions in household income, asset wealth, and productive capacity (Dercon, 2004; Rosenzweig and Wolpin, 1993). The loss of a cow can reduce current income, future production, breeding capacity, and collateral value. This risk is especially important for cattle because individual animals are valuable relative to many farmers' liquid wealth. In China, cattle production is economically significant: China had more than 100 million head of cattle in 2024, making it one of the largest cattle-producing countries in the world (National Bureau of Statistics of China, 2024). In our setting, the market value of an adult cow is often substantial, and high-quality dairy cattle can be worth considerably more. For many households, therefore, cattle mortality is not a small income shock but a large asset loss.

3.2 Cattle Mortality Insurance and the Chinese Market

Cattle mortality insurance differs from many other insurance markets in a fundamental way: the insured object must be identified, verified, and monitored over time. In health insurance, auto insurance, or homeowners insurance, the insured object is usually linked to a stable administrative identifier, such as a person, a vehicle identification number, or a property

Figure 2: Cattle's Visible Phenotypic Traits



[a]

[a]Source: A anonymous AI-Based cattle facial recognition company.

address. In cattle insurance, the insured object is a living animal. Cattle are visually similar, mobile, and costly to inspect. They may be moved across locations, traded, lost, replaced, or commingled with other animals. As a result, the central informational problem is not only whether a loss occurred, but whether the dead animal is the same animal that was originally insured.

This identity-verification problem creates several forms of market friction. If the insurer cannot reliably match a claim to a specific insured animal, farmers may be able to claim on uninsured animals, submit repeated claims on the same dead animal, or insure the same animal multiple times across insurers. **These behaviors are difficult to detect when the contract is not perfectly tied to the biological identity of a particular animal.**

China is a useful setting for studying these issues because it combines a large cattle insurance market with the introduction of AI-based animal identification technology. The market operates through a heavily subsidized public-private insurance system, covers millions of animals, and requires insurers to verify individual cattle during underwriting and claims processing. The product we study is a one-year mortality insurance contract for cattle.[2] The contract pays a fixed indemnity when an insured cow dies from covered causes, including disease, natural disaster, accident, wild animal attack, difficult birth, epidemic disease, and government-mandated culling. Unlike indemnity contracts that reimburse re-

[2]The cattle typically refers to dairy cows, including specialized dairy cows used primarily for milk production and multipurpose dairy–beef breeding cows between one and seven years of age.

alized market value, the cattle insurance contract in our setting specifies a fixed coverage amount, with the coverage level depending primarily on animal age. the indemnity for adult cattle is typically between RMB 8,000 and RMB 12,000 per cow. For younger cattle, the coverage amount is lower, typically around RMB 2,000–4,000. A representative adult-cattle contract has a coverage amount of RMB 10,000 and a premium rate of 6 percent, implying an annual premium of RMB 600 per cow. Farmers pay approximately 20 percent of the premium, while government subsidies cover the remaining 80 percent.[3]

The underwriting process is operationally intensive. Before a policy is issued, the underwriter must visit the farm, verify ownership, inspect the animal's health and eligibility, record identifying information of animal. After the animal is registered and identified, the farmer pays the farmer-paid portion of the premium, and the contract becomes effective. The claim process mirrors the underwriting process. When a cow dies, the insurer must inspect the dead animal, verify that the cause of death is covered, and determine whether the claimed animal is the same animal that was insured.

The identity verification is central to cattle insurance. The insurer's problem is not simply to estimate expected mortality risk. It must also verify that each contract corresponds to a unique animal at underwriting and that each claim corresponds to the same animal at payout. Weak identity verification can lead to claims on uninsured cattle, repeated claims, and duplicate insurance. High verification costs can also limit supply because underwriters must spend time traveling to farms, inspecting animals, and recording identifying information.

3.3 Traditional Identification Technology: Ear Tags

Traditionally, cattle insurers have relied on ear tags to link an insurance contract to a physical animal. The earliest form is an external plastic tag attached to the animal's ear. A more advanced version is an electronic radio-frequency identification (RFID) tag, which can be attached to the ear or implanted under the skin. At underwriting, the underwriter scans or records the tag number and links it to the policy record. At claim time, the insurer inspects the dead animal, checks whether the tag number matches the insured record, and often compares field photographs of the claimed animal with underwriting records by visual inspection. Figure 3 illustrates the main forms of ear tags and the physical tagging

[3] In China, the government subsidizes the premium for dairy cow insurance at an overall rate of approximately 80 percent. While this subsidy ratio may vary slightly across regions, the region from which our study sample was drawn has a subsidy rate of exactly 80 percent.

process.

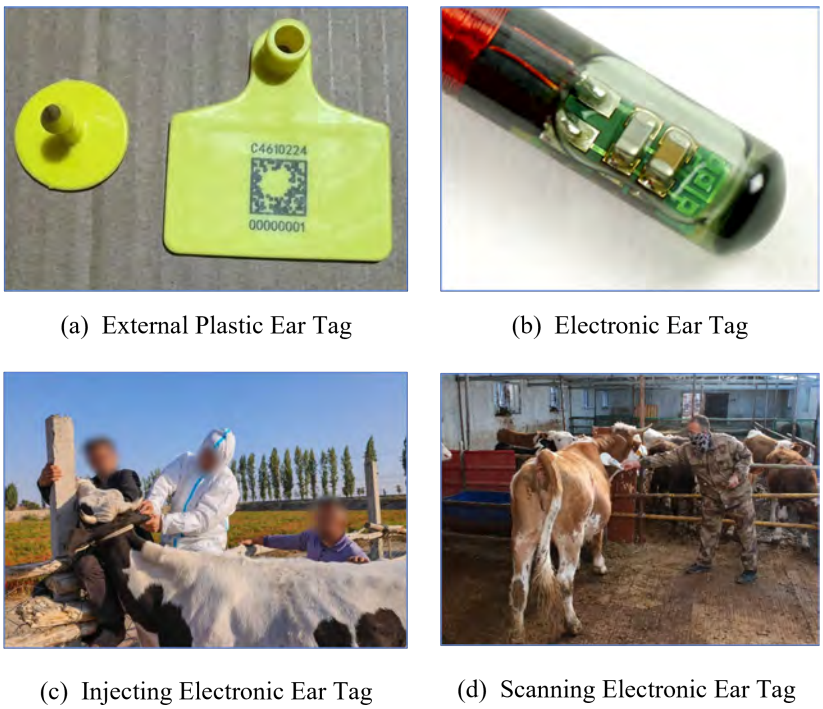
Ear tags partially solve the identity problem, but they leave two important frictions. The first friction is imperfect identity verification. Ear tags are physical objects that can fall off, be damaged, be removed, or be replaced. External plastic tags are vulnerable to loss and manipulation. Electronic RFID tags require specialized reading devices and can still be damaged, removed, or implanted in locations that make verification difficult. Because different insurers may use different tag systems, a farmer may also be able to insure the same animal multiple times across companies. These limitations create scope for claims on uninsured animals, repeated claims on the same dead animal, and duplicate insurance. In some cases, such fraud may also require collusion between farmers and local underwriters or claim adjusters, especially when field verification depends on local discretion and visual inspection. Thus, ear tags reduce but do not eliminate the core contracting problem in cattle insurance: the insurer must verify that the claimed dead animal is the same animal that was originally insured.

The second friction is transaction cost. Ear-tag identification is costly and time-consuming. In our setting, the direct cost of an ear tag and related tagging inputs is approximately RMB 15, or about USD 2, per animal. More importantly, tagging often requires workers to physically restrain the animal, attach or implant the tag, scan or record the tag information, and upload the identifying record. Panel (c) of Figure 3 shows that this process can require physically controlling the animal before the tag is attached or implanted. The process is physically demanding and sometimes dangerous for underwriters and farmers. It can also generate stress responses for cattle, including injury, miscarriage, lower milk production, and disease risk. These costs are especially important in settings where farms are geographically dispersed and the underwriting season is short. Even when the premium and subsidy schedule make insurance financially attractive, high verification costs can limit the number of animals an underwriter can process.

3.4 AI-Based Cattle Facial Recognition

AI-based cattle facial recognition replaces physical tag-based identification with a computer-vision technology. Instead of relying only on an external or implanted tag, the underwriter takes images of the cow using a smartphone. A deep-learning algorithm extracts visual features from the animal's head and face, including color patterns, head structure, horns, and other identifying characteristics. These features are then compared with the insurer's

Figure 3: Underwriting with Ear Tags



(a) External Plastic Ear Tag

(b) Electronic Ear Tag

(c) Injecting Electronic Ear Tag

(d) Scanning Electronic Ear Tag

Notes: This figure shows traditional ear-tag technologies and the tagging process. Panel (a) presents an external plastic tag with a QR code. Panel (b) displays an electronic RFID tag. Panel (c) shows workers physically restraining cattle and injecting an ear tag. Panel (d) shows an RFID tag placed below the neck and an RFID device used to read the tag's information. Panels (a) and (b) are from an anonymized ear-tag company. Panel (c) is from a public online source. Panel (d) was taken by the authors during fieldwork.

image database to determine whether the cow has already been insured, whether a claimed dead cow matches an insured record, and whether the same dead cow has already been claimed.

This technology changes the economics of cattle insurance in two ways. First, it improves animal identity verification. A policy can be linked to the visual identity of a specific cow, and a claim can be checked against the image record of that cow. This reduces the scope for claims on uninsured animals, repeated claims on the same dead animal, and duplicate insurance across insurers. Second, it lowers transaction costs. Underwriters no longer need to rely solely on physically restraining cattle, implanting tags, scanning RFID devices, or comparing photographs by human eyes. The process is faster, less physically demanding, and easier to scale. Other technologies, such as DNA testing, iris recognition, and nose-pattern recognition, could also identify individual cattle in principle. However, these methods are costly, time-consuming, and difficult to implement at scale in routine

insurance underwriting. For livestock insurance, the relevant technology must satisfy two requirements at the same time: it must improve identity verification, and it must reduce the cost of verification. AI-based cattle facial recognition is attractive because it can do both. It strengthens the link between the contract and the animal while also lowering the time, labor, and equipment costs of underwriting.

This institutional setting maps directly to the two frictions in the theoretical framework. First, imperfect animal identity verification allows insurance contracts to be used against deaths of animals that were not actually insured. In the model, this creates a downward-sloping marginal claim cost curve: when identity cannot be perfectly verified, the first few contracts can effectively cover deaths from the entire herd. AI removes this distortion by tying coverage and claims to the identity of a specific insured cow. Second, underwriting is costly. Underwriters must travel to farms, inspect animals, physically handle cattle, verify ownership and condition, and record identifying information. These transaction costs generate capacity constraints. AI lowers the per-cow transaction cost and allows the same underwriter, or the same underwriting network, to process more cows and reach more households.

4 Data and Descriptive Analysis

We combine several administrative datasets to study how AI-based cattle facial recognition affects fraud, transaction costs, and insurance market expansion in livestock insurance. The empirical setting is the cow insurance market in a large city in China, where the cattle inventory is estimated at approximately one million head in 2025 and the government-subsidized livestock insurance covers the risk of cattle mortality. The main sample comes from a large livestock insurer that accounts for roughly 60 percent of the local market and is one of the early adopters of AI-based cattle facial recognition. It covers 2021–2025, a period during which the focal insurer transitions from traditional ear-tag-based identification to AI-based cattle facial recognition. The data are uniquely suited for our purpose because they link insurance enrollment, animal-level attributes, claims, underwriting technology, underwriter identity, and geographic information.

Cow- and policyholder-level data. The core micro dataset is Cow-level data. The unit of observation is cow–year. The data contain approximately 1.94 million cow-year observations and about 48,000 policyholders. For each insured cow, we observe policyholder identifiers, insurance dates, insured value, premium, claim outcomes, whether the cow is

newly insured, whether AI facial recognition was used, and geographic location. The data also include policyholder characteristics such as age, gender, and size. These data allow us to study both the fraud channel, through claim outcomes, and the coverage channel, through enrollment and insured herd size.

For each insured cow, we observe the policy year, policyholder identifier, insured date, insured value, premium, claim indicator, location, and whether the cow was processed using AI facial recognition. The data also contain information about the underwriter or insurance assistant, including underwriter identifier, age, gender, and education. These variables allow us to construct outcomes at three levels: the cow-year level, the household-year level, and the underwriter-year level.

Table **??** reports summary statistics at three levels. Panel A describes cow-year data from the focal insurer, including insured value, claim rates, whether claimed, new-cow indicators, AI use, and distance measures. Panel B describes the policyholder-year and cow-year data from the focal insurer, including insured herd size, premium per cow, insured value, claim rates, claim amounts, new-cow indicators, AI use, and distance measures. Panel C describes the underwriter-year data, including the number of cows and households underwritten, the number of new households served, AI adoption, distance to served households, and underwriter demographics.

Underwriter data. We also observe underwriter-level administrative records. Underwriters are local agents, often veterinarians or animal disease-control workers, who conduct farm visits, verify cattle, record identifying information, and assist with underwriting. The data include underwriter identifiers, demographic characteristics such as gender, age, education, and measures of annual underwriting activity. By linking underwriters to insured cows and policyholders, we can measure underwriter productivity, including the number of cows and households underwritten, the number of new customers reached, and the geographic distance of served households.

Geographic and market-background data. We link insurance records to geographic information using policyholder addresses and village-level coordinates. This allows us to measure the distance between insured households and local administrative centers, and to examine whether AI expands coverage toward more remote households. We also use local agricultural statistics on cattle inventory, cattle sales, number of livestock farmers, and county-level livestock market conditions. In addition, weekly live-cattle price data from

Table 1: Summary Statistics

	Mean	SD	Median	Min	Max
Panel A. Cow-level data					
Claim indicator	0.049	0.215	0.000	0	1
Premium per cow	436.002	135.866	462.000	0.000	1422.857
Coverage value per cow	8391.490	1913.682	8000.000	2000.000	14000.000
Newly insured cow	0.856	0.351	1.000	0	1
Distance to county government, km	38.450	27.037	31.018	0.039	254.044
Panel B. Household-year data					
Number of insured cows	16.781	15.870	13.000	1	718
Claim rate	0.055	0.109	0.000	0.000	1.000
Household AI share	0.497	0.497	0.286	0.000	1.000
Number of newly insured cows	14.371	15.133	10.000	0	718
Panel C. Underwriter-year data					
Number of insured cows	1063.722	1398.462	638.000	1	15,484
Number of households	63.449	77.536	42.000	0	778
Number of newly insured cows	910.846	1224.658	558.000	0	15,098
Number of new households	26.379	46.051	12.000	0	639
Underwriter AI share	0.518	0.483	0.708	0.000	1.000
Mean distance to households, km	35.008	23.460	30.217	0.058	220.021
Share of small households	0.068	0.153	0.030	0.000	1.000

Notes: This table reports summary statistics for the three main analysis datasets. Panel A uses the cow-level analysis sample. Panel B uses the household-year-level sample. Panel C uses the underwriter-year-level sample. The table reports means, standard deviations, medians, minimums, and maximums. The numbers of observations are 1,941,070 in Panel A, 115,668 in Panel B, and 1,817 in Panel C.

major producing provinces are used to describe the broader cattle-price cycle during the sample period.

AI adoption and underwriting technology. The staggered adoption of AI provides the empirical variation. The insurer ran a very small pilot program in 2022. The technology became fully available to underwriters in 2023, but underwriters retained discretion over whether to use it. In 2023, approximately 52 percent of insured cows were underwritten using AI. By 2024 and 2025, roughly 90 percent of insured cows were underwritten using AI. This adoption pattern is useful for two reasons. First, 2023 provides cross-sectional variation in AI use across households and underwriters within the same market and policy environment. Second, the 2021–2025 panel allows us to study whether claim rates and insured quantities changed differentially for households and underwriters more exposed to

AI adoption in 2023.

We begin by documenting the adoption of AI. Before 2022, the insurer relied almost entirely on traditional ear-tag identification. Under this technology, underwriters or insurance assistants had to visit farms, physically restrain cattle, attach or scan ear tags, verify ownership and animal condition, and record identifying information. The process was costly because it required time, physical effort, specialized equipment, and cooperation from farmers. It was also imperfect because ear tags could be damaged, lost, replaced, or manipulated.

AI facial recognition changed this production process. Underwriters could use a smartphone to take images of the cow and compare the extracted image features with the insurer's database. The technology allowed the insurer to check whether a cow had already been insured, whether a claimed dead cow matched an insured record, and whether the same dead cow had been claimed before. Thus, AI affected both sides of the insurance transaction: it improved claim verification and reduced the cost of underwriting.

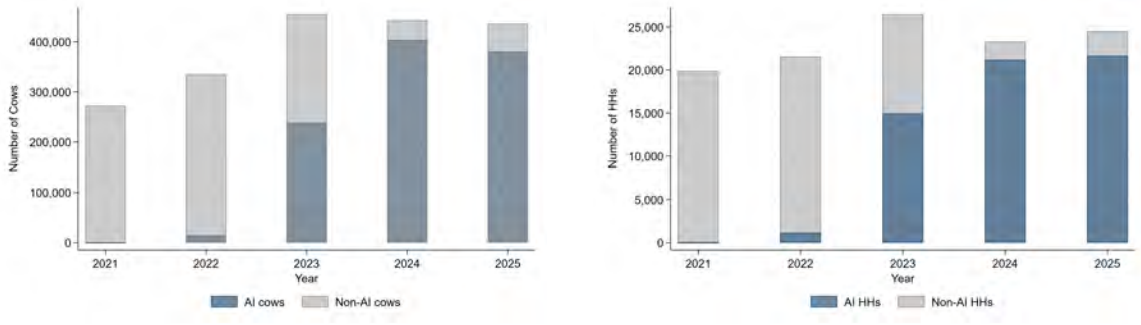
The data cover the insurer's cattle insurance business in one city from 2021 to 2025. The insurer has approximately 60 percent market share in the city. The main unit of observation is a cow-year. The sample contains approximately 1.94 million cow-year observations and 48,000 policyholders. For each cow, we observe policy year, household identifier, insured value, premium, claim status, whether the cow was processed using AI, location, and underwriter identifier. We also observe underwriter characteristics, including gender, age, and education.

The baseline descriptive patterns already point to the two forces emphasized by the model. First, claim rates are lower among AI-processed cows and households with greater AI exposure. Second, insured quantities increase sharply after AI adoption: the insurer's annual insured cattle count rose from roughly 280,000 in 2022 to roughly 460,000 in 2024. The remaining subsections turn these patterns into regression evidence.

We construct two main exposure measures. At the household level, let $AI_{h,2023}$ denote the share of household h's insured cows in 2023 that were processed using AI. This variable captures the household's exposure to the new identity-verification technology at the point when the technology became broadly available but was not yet universal. At the underwriter level, let $AI_{u,2023}$ denote the share of cows processed using AI by underwriter u in 2023. This variable captures the underwriter's early adoption intensity and is used to study

productivity, market expansion, and reallocation across underwriters.

Figure 4: AI Adoption Within the Focal Insurer



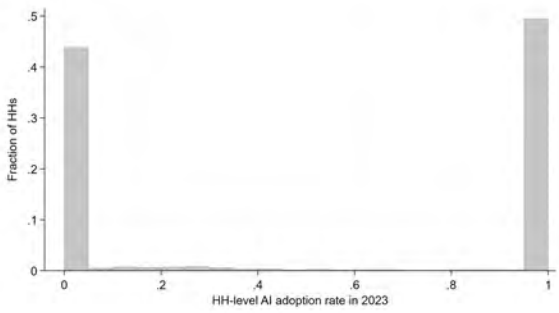
(a) Cow AI exposure

(b) Household AI exposure

Notes: Panel A plots the number of cows as well as share of AI treated cows. Panel B plots the number of policyholders as well as share of AI treated cows.

Descriptive evidence on AI adoption. We then document the diffusion of AI technology within the focal insurer. Figure 4 plots the share of cows and households underwritten using AI by year. It also shows the distribution of household-level AI exposure in 2023, the first year when underwriters had discretion over whether to use AI.

Figure 5: AI Adoption at Policyholder level in 2023



Notes: The figure plots the histogram of AI adoption rate at policyholder level in 2023.

AI adoption at the policyholder level. Figure 5 shows that AI adoption at the policyholder level in 2023 is highly polarized. About 44 percent of policyholders have an AI adoption rate of zero, while about 51 percent have an adoption rate of one. In other words, most policyholders are either fully non-AI or fully AI, with very limited partial adoption within

policyholders. This pattern indicates that AI use was largely assigned at the policyholder-underwriter interaction level rather than varying cow by cow within the same household. It motivates both across-policyholder comparisons between AI and non-AI households and within-policyholder analyses using the small subset of households with partial adoption.

5 Empirical Evidence

This section presents the empirical evidence on how AI-based cattle facial recognition changes the operation of livestock insurance markets. The analysis is organized around the two mechanisms in the model. First, AI improves animal identity verification and reduces identity-based fraud. Second, AI lowers the transaction cost of underwriting, relaxes underwriter capacity constraints, and expands insurance coverage. We then examine the distributional implications of this expansion by studying whether AI helps the insurer reach more distant and smaller households.

A central feature of the empirical design is that AI adoption was neither instantaneous nor universal. The insurer operated a small pilot in 2022. In 2023, AI facial recognition became available to all underwriters, but underwriters could choose whether to use it. In that year, 52 percent of insured cows were processed using AI. By 2024 and 2025, AI use had expanded to roughly 90 percent of insured cows. This adoption path generates useful variation across cows, households, and underwriters during the transition from ear-tag-based identification to AI-based identification.

5.1 AI and Claim Rates: Evidence on Identity-Based Fraud

The first prediction of the model is that AI should reduce paid claim rates by improving identity verification. When the insurer cannot perfectly identify cow identity, a farmer may insure only part of the herd while effectively claiming against deaths from the entire herd. In the model, this means that the marginal claim cost of the nth contract is not πL, but $\Pr(X \geq n)L$. AI changes this mapping by tying the claim to the identity of a specific insured animal. Therefore, the model predicts a decline in observed paid claim rates after AI adoption.

We first present descriptive evidence on AI adoption and claim outcomes in 2023, the first year in which AI-based cattle facial recognition was broadly available to underwriters. The policyholder-level adoption distribution is highly polarized: about 44 percent of policyholders have an AI adoption rate of zero, while about 51 percent have an adoption rate

of one, with very limited partial adoption within policyholders. This suggests that AI use is largely determined at the policyholder–underwriter interaction level rather than varying cow by cow within the same household. Panel A of Figure 6 provides a simple comparison of cow-level claim rates between AI-treated and non-AI cows in 2023. The average claim rate is 6.61 percent for AI-treated cows, compared with 6.98 percent for non-AI cows. The difference is 0.37 percentage points, implying that AI-treated cows have a claim rate that is about 5.31 percent lower than non-AI cows. This descriptive pattern is consistent with the hypothesis that AI reduces claims that rely on imperfect identity verification, although it does not by itself establish a causal effect.

We next examine whether AI-treated cows have lower claim rates, as predicted by the identity-verification channel. We start with a simple cow-level comparison in Panel A of Figure 6. The average claim rate is 6.61 percent for AI-treated cows and 6.98 percent for non-AI cows, a difference of 0.37 percentage points, or 5.31 percent lower relative to the non-AI mean.

To move beyond this raw comparison, Panel B of Figure 6 reports regression estimates of the relationship between AI use and cow-level claims in 2023. We begin with the following specification without controls:

$$Claim_{ih} = \beta AI_{ih} + \varepsilon_{ih}, \qquad t = 2023, \tag{15}$$

where $Claim_{ih}$ is an indicator for whether cow i in household h had a paid claim in 2023, and AI_{ih} indicates whether the cow was processed using AI. The coefficient from this specification is the first point in Panel B. It reproduces the raw difference in Panel A: AI-treated cows have claim rates that are 0.37 percentage points lower than non-AI cows.

A concern with this simple comparison is that AI use may be correlated with other determinants of claim risk, such as household management style, local disease risk, or the household's propensity to file claims. We therefore estimate progressively richer specifications. The second estimate adds controls for premium per cow, household insured herd size, and whether the cow is newly insured. The third estimate adds household fixed effects:

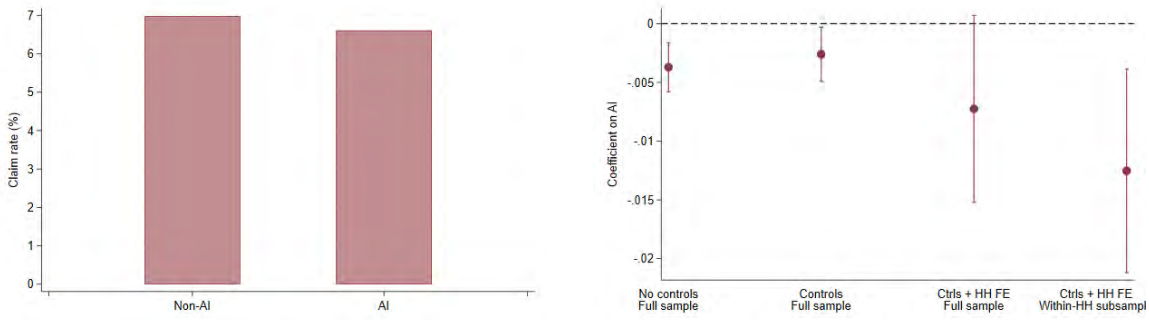
$$Claim_{ih} = \beta AI_{ih} + X_{ih}'\Gamma + \mu_h + \varepsilon_{ih}, \qquad t = 2023, \tag{16}$$

where X_{ih} includes the cow-level and policy controls, and μ_h is a household fixed effect.

Standard errors are clustered at the household level.

The household fixed-effect specification compares AI and non-AI cows within the same household in the same year. This absorbs all household-level determinants of claim risk, including herd management, location, farmer type, local disease environment, and the household's general propensity to file claims. The fourth estimate further restricts the sample to households with within-household variation in AI use in 2023 and estimates the same specification with controls and household fixed effects. Across these specifications, AI-treated cows have lower paid claim rates than non-AI cows. The estimates become larger in magnitude when household fixed effects are included, suggesting that the raw difference is not simply driven by safer or better-managed households adopting AI. Instead, the evidence is consistent with the model's prediction that AI reduces claims that rely on imperfect identity verification.

Figure 6: AI Underwriting and Cow-Level Claim Rates in 2023



(a) Cow-level claim rate: AI vs. non-AI

(b) Regression estimates of the effect of AI on claim rates

Notes: This figure presents descriptive and regression-based evidence on the relationship between AI underwriting and cow-level claim outcomes in 2023. Panel A plots the mean cow-level claim rate for cows underwritten using traditional methods (non-AI) and cows underwritten using AI. The unit of observation is a cow-year, and the sample is restricted to 2023. Claim rate is defined as the share of cows with a positive claim indicator. Panel B reports the coefficient on the AI indicator, $AI_{i,2023}$, from regressions of the cow-level claim indicator on AI status using the 2023 cow-level sample. The first three estimates use the full sample: (1) no controls; (2) with controls for premium per cow, household insured herd size, and a new-cow indicator; and (3) with the same controls plus household fixed effects. The fourth estimate uses the subsample of households with within-household variation in AI use in 2023 and includes controls and household fixed effects. Standard errors are clustered at the household level, and the vertical bars in Panel B denote 95% confidence intervals.

Although the within-policyholder cow-level comparison in 2023 is useful, it may still reflect other within-household factors that vary across cows, such as differences in cow age, insurance timing, or which cows are selected into AI processing. We therefore complement

Table 2: AI Underwriting and Cow-Level Claim Rates in 2023

	Full sample			Subsample w. within HH variation
	(1) No controls	(2) Controls	(3) Controls + HH FE	(4) Controls + HH FE
AI	-0.0037***	-0.0026**	-0.0072*	-0.0125***
	(0.0011)	(0.0012)	(0.0041)	(0.0044)
Controls	No	Yes	Yes	Yes
Household FE	No	No	Yes	Yes
Observations	454,804	454,804	454,804	42,855
R-squared	0.000	0.001	0.102	0.074

Notes: This table reports cow-level regressions using the 2023 sample. The dependent variable is an indicator for whether the cow had a paid claim. *AI* is an indicator for whether the cow was processed using AI-based cattle facial recognition. Controls include premium per cow, household insured herd size, and an indicator for newly insured cows. Column (4) restricts the sample to households with within-household variation in AI use in 2023. Standard errors, clustered at the household level, are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

this analysis by exploiting within-household changes in claim rates over time. Specifically, we estimate an event-study specification at the household-year level:

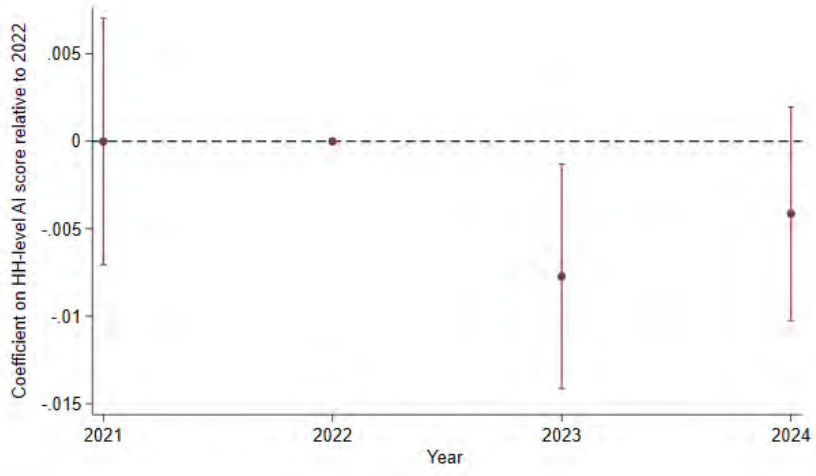
$$ClaimRate_{ht} = \sum_{k \neq 2022} \beta_k \mathbf{1}\{t = k\} \times AI_{h,2023} + X_{ht}'\Gamma + \mu_h + \lambda_t + \varepsilon_{ht}. \tag{17}$$

Here, $ClaimRate_{ht}$ is the claim rate for household h in year t, and $AI_{h,2023}$ is the share of household h's insured cows in 2023 that were processed using AI. The omitted year is 2022, immediately before AI became broadly available. The household fixed effect μ_h absorbs time-invariant differences across households, such as persistent herd management practices, location, farmer type, and long-run claim propensity. The year fixed effect λ_t absorbs aggregate changes in policy, disease risk, cattle market conditions, and claim-processing rules. The vector X_{ht} includes time-varying household-level controls.

This specification should not be interpreted as a standard staggered-adoption difference-in-differences design. The treatment variable, $AI_{h,2023}$, is fixed based on AI exposure in 2023. It does not continue to measure differential treatment intensity in later years, especially because AI adoption becomes close to universal after 2023. Instead, the coefficients β_k describe whether households with higher AI exposure in 2023 experienced differential changes in claim rates relative to households with lower 2023 AI exposure. The pre-2023 coefficients provide a diagnostic check for differential pre-trends, while the later-year coef-

ficients should be interpreted more cautiously. Since most policyholders eventually adopt AI, we expect the differential effect between high- and low-2023-AI households to diminish over time.

Figure 7: AI Exposure and Household Claim Rates Over Time



Notes: This figure reports coefficients from Equation (17). The dependent variable is the household-year claim rate. The key explanatory variables are interactions between year indicators and $AI_{h,2023}$, the share of household h's insured cows in 2023 that were processed using AI. The omitted year is 2022. The specification includes household fixed effects, year fixed effects, and household-year controls. Standard errors are clustered at the household level. Vertical bars report 95 percent confidence intervals.

The event-study evidence is consistent with the within-household cow-level result. Before broad AI adoption, households with higher and lower 2023 AI exposure do not display differential claim-rate trends: the 2021 coefficient is essentially zero and statistically insignificant. In 2023, when AI becomes broadly available but is not yet universal, the coefficient becomes negative and statistically significant. The estimate is -0.0077 with a standard error of 0.0033, implying that moving from zero to full AI exposure in 2023 is associated with a 0.77 percentage-point decline in the household claim rate relative to 2022. This corresponds to a 13.95 percent reduction relative to the baseline claim rate. In 2024, the coefficient remains negative, at -0.0042, but is no longer statistically significant. This attenuation is consistent with the fact that AI adoption becomes widespread after 2023, so the fixed 2023 exposure measure no longer captures persistent differences in treatment intensity.

These results closely match the fraud channel in the model. AI does not change the

biological mortality probability π_i of a cow. Instead, it changes whether a paid claim can be made on a cow that is not the insured cow. Therefore, a decline in paid claim rates is exactly what the model predicts when the market moves from imperfect identity verification toward accurate identity verification. The empirical result corresponds to the theoretical movement from $\Pr(X \geq n)L$ toward πL.

5.2 AI and Household Insured Size

The model also implies that improved identity verification can change the quantity of insurance purchased by existing households. When identity cannot be perfectly verified, a household may have an incentive to insure only a subset of its herd, because a small number of insured contracts can still be used to claim against deaths from the entire herd. Once AI ties coverage to specific cows, partial insurance becomes less useful as a fraud technology. At the same time, AI lowers underwriting costs, which can increase the supply of coverage. The net effect on insured size within existing households is therefore theoretically ambiguous and depends on which margin dominates.

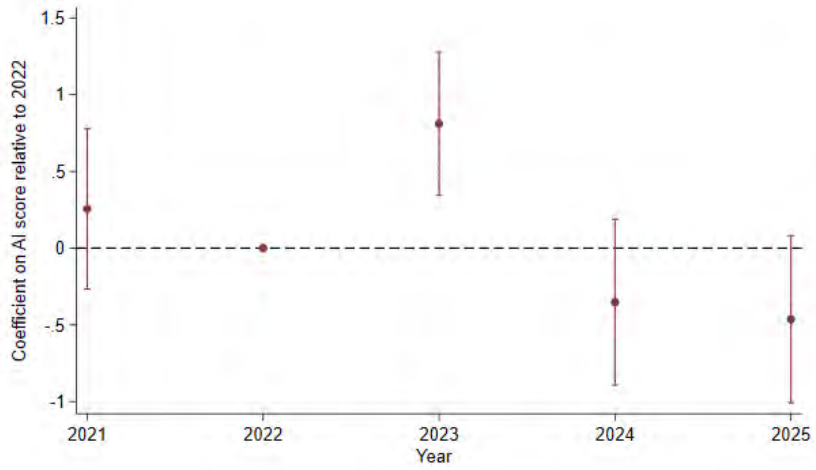
We examine household insured size using the same 2023 exposure design. For each household-year, let $Cows_{ht}$ denote the number of insured cows. We estimate

$$Cows_{ht} = \sum_{k \neq 2022} \theta_k \mathbf{1}\{t = k\} \times AI_{h,2023} + X'_{ht}\Pi + \mu_h + \lambda_t + \nu_{ht}. \tag{18}$$

Here, $AI_{h,2023}$ is the share of household h's insured cows in 2023 that were processed using AI. The omitted year is 2022. Household fixed effects absorb permanent differences in herd scale, farmer type, location, and management style, while year fixed effects absorb aggregate changes in the insurance program and local market conditions. As in the claim-rate event study, this is not a standard staggered-adoption difference-in-differences design. The treatment variable is fixed using 2023 AI exposure, and later years should be interpreted as a test of whether differential outcomes persist once AI adoption becomes more widespread.

Figure 8 shows a sharp increase in insured herd size among households with higher AI exposure in 2023. Relative to 2022, the coefficient for 2023 is 0.811 and statistically significant at the 1 percent level. This implies that moving from zero to full AI exposure in 2023 is associated with about 0.81 additional insured cows per household, or a 4.84 percent increase relative to the baseline insured herd size. In contrast, the 2021 coefficient is small and statistically insignificant, suggesting no strong differential pre-trend before broad AI

Figure 8: AI Exposure and Household Insured Herd Size Over Time



Notes: This figure reports coefficients from Equation (18). The dependent variable is the number of insured cows at the household-year level. The key explanatory variables are interactions between year indicators and $AI_{h,2023}$, the share of household h's insured cows in 2023 that were processed using AI. The omitted year is 2022. The specification includes household fixed effects, year fixed effects, and household-year controls. Standard errors are clustered at the household level. Vertical bars report 95 percent confidence intervals.

availability. The coefficients in 2024 and 2025 are negative and not statistically significant at conventional levels in 2024, with only marginal significance in 2025. This pattern is consistent with a temporary differential response in 2023 that diminishes once AI becomes widely adopted across households.

The insured-size result should be interpreted jointly with the claim-rate evidence. The decline in paid claims among high-AI households suggests that AI reduces the value of using a limited number of insured contracts to claim against losses from the broader herd. At the same time, the positive 2023 effect on insured herd size indicates that AI may have increased coverage among treated households when the technology first became broadly available. Thus, the household-level quantity response does not simply reflect greater demand for insurance. It also reflects the removal of an implicit fraud opportunity and the relaxation of underwriting frictions.

This distinction is important for connecting the empirical evidence to the model. The identity-verification channel works through the effective claim technology and the farmer's incentive to use insured contracts strategically. The productivity channel works through the insurer's ability to process and serve more cows and households. The evidence in Figure 8, together with the claim-rate results, suggests that AI both changes how existing households

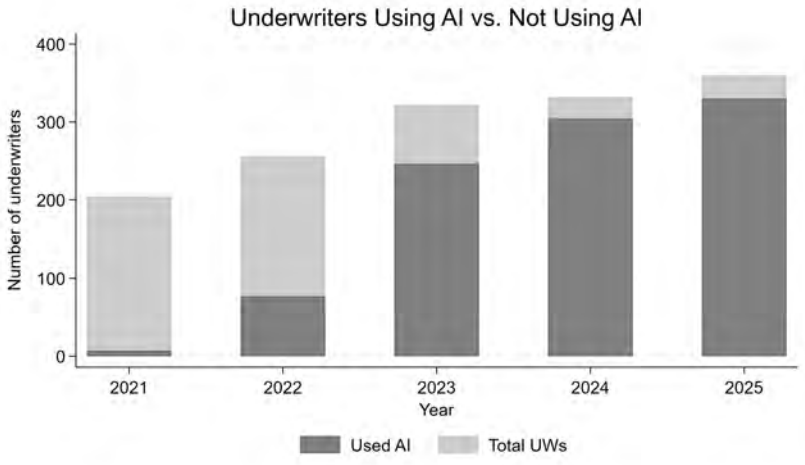
use insurance and contributes to broader coverage expansion.

5.3 AI and Underwriter Productivity

The second central prediction of the model is that AI lowers transaction costs and increases underwriting productivity. Under the traditional technology, underwriting requires travel, physical inspection, animal restraint, ear-tagging or scanning, and administrative recording. These steps are labor intensive and time consuming, which limits the number of cows and households an underwriter can process during the short underwriting season.

We begin by documenting the diffusion of AI at the underwriter level. Figure 9 plots the number of active underwriters by year and distinguishes those using AI from those not using AI. The figure shows a rapid transition from traditional underwriting to AI-based underwriting. In 2021, almost no underwriters used AI. In 2022, AI use was still limited, reflecting the pilot stage. In 2023, when AI became broadly available, a clear majority of underwriters adopted it. By 2024 and 2025, nearly all active underwriters were using AI. This pattern confirms that 2023 was the key transition year and motivates our use of 2023 AI intensity as the main exposure measure in the underwriter-level analysis.

Figure 9: Underwriters Using AI vs. Not Using AI



Notes: This figure plots the number of active underwriters by year and distinguishes underwriters using AI-based cattle facial recognition from those not using AI. The figure illustrates the rapid diffusion of AI underwriting between 2022 and 2025.

Let $AI_{u,2023}$ denote the share of cows processed using AI by underwriter u in 2023. Our first direct measure of productivity is the average underwriting time per cow. We construct this measure as described in Appendix A. Before AI adoption, the underwriting process

generates time-stamped photos for each household visit, which allow us to recover the elapsed time needed to process a sequence of cows. After AI adoption, the data directly record a time stamp for each cow processed using AI. Using these records, we construct a comparable measure of average underwriting time per cow at the underwriter-year level.

We estimate

$$TimePerCow_{ut} = \sum_{k \neq 2022} \delta_k \mathbf{1}\{t = k\} \times AI_{u,2023} + X'_{ut}\Lambda + \alpha_u + \lambda_t + \eta_{ut}. \tag{19}$$

Underwriter fixed effects α_u absorb permanent differences in ability, local networks, experience, and baseline productivity. Year fixed effects λ_t absorb market-wide changes in demand, policy, prices, and aggregate insurance supply. The coefficients δ_k therefore capture whether underwriters with greater AI exposure in 2023 experienced differential changes in underwriting time relative to 2022.

As in the earlier household-level event-study specifications, this is not a standard staggered-adoption difference-in-differences design. The key variable $AI_{u,2023}$ is fixed using each underwriter's 2023 AI exposure. Once AI becomes nearly universal, later-year coefficients should not be interpreted as continued treatment effects from differential adoption intensity. Rather, they indicate whether underwriters who were more exposed to AI in 2023 exhibited different productivity dynamics over time.

We are still finalizing the underwriting-time data. Preliminary analysis, however, points to large productivity gains from AI. Using the time-stamp method described in Appendix A, we find that AI reduces average underwriting time per cow from about 3 minutes under the traditional process to about 1 minute under AI-based underwriting, a decline of roughly 70 percent. This large reduction is consistent with the idea that AI substantially lowers the transaction cost of verifying and recording cattle identity. In the model, this corresponds to a decline in c_i^u, which relaxes underwriter capacity constraints and allows each underwriter to serve more cows and more households.

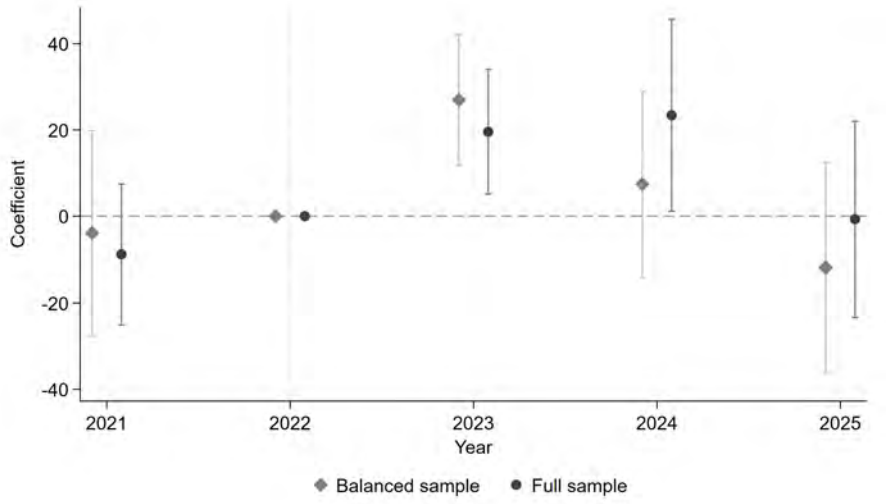
5.4 Distributional Implications of Lower Transaction Costs

A direct implication of the transaction-cost channel is that AI should allow underwriters to serve more policyholders. Under the traditional technology, underwriting requires travel, physical inspection, animal restraint, ear-tagging or scanning, and administrative recording. These tasks are costly in time and effort, especially during the short underwriting season. If

AI lowers the per-cow and per-household cost of verification, underwriters should be able to process more policyholders with the same capacity.

We test this prediction at the underwriter-year level by re-estimating Equation (19), using the number of policyholders underwritten by underwriter u in year t as the dependent variable. Let $AI_{u,2023}$ denote the share of cows processed using AI by underwriter u in 2023. We estimate the model separately for two samples. The first is a balanced sample of underwriters who are continuously active over the sample period. This sample isolates the intensive-margin effect of AI on incumbent underwriters. The second is the full underwriter-year panel, which includes entrants and exiters and therefore captures both productivity changes among incumbents and composition changes in the underwriter market.

Figure 10: AI Exposure and Number of Policyholders Underwritten



Notes: This figure reports event-study coefficients from underwriter-year regressions. The dependent variable is the number of policyholders served by an underwriter in a year. The key explanatory variables are interactions between year indicators and $AI_{u,2023}$, the share of cows processed using AI by underwriter u in 2023. The omitted year is 2022. The balanced sample restricts to underwriters continuously active over the sample period. The full sample includes all underwriters, including entrants and exiters. The regressions include underwriter fixed effects and year fixed effects. Vertical bars report 95 percent confidence intervals.

Figure 10 shows that underwriters with greater AI exposure in 2023 served more policyholders after AI became broadly available. In the balanced sample, the 2023 coefficient is positive and statistically significant, indicating that incumbent underwriters with higher AI use expanded the number of households they served. The point estimate is roughly

27 additional policyholders, relative to 2022. The 2021 coefficient is close to zero and statistically insignificant, suggesting no strong differential pre-trend before the AI rollout. The balanced-sample effects become smaller and less precise in 2024 and 2025, which is consistent with the interpretation that the 2023 exposure measure becomes less informative once AI adoption becomes widespread.

The full-sample estimates show a somewhat different pattern. The 2023 coefficient is also positive, at roughly 20 additional policyholders, and the effect remains positive in 2024, at roughly 23 additional policyholders. This persistence in the full sample suggests that AI affects not only the productivity of continuing underwriters, but also the composition of active underwriters. In particular, lower transaction costs may facilitate entry or continuation by underwriters who can use AI effectively. By 2025, the full-sample coefficient is close to zero and imprecisely estimated, consistent with the market having largely transitioned to AI-based underwriting.

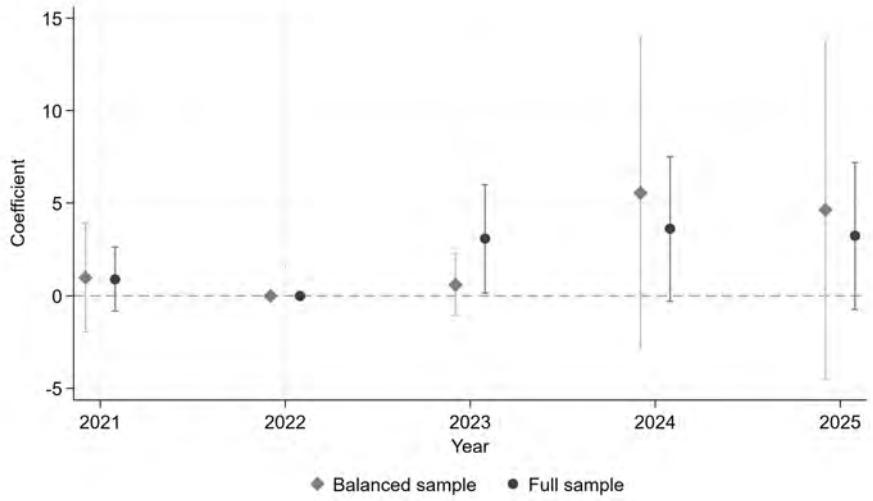
This distinction helps interpret the market-wide expansion documented in the data. The increase in insured cows and policyholders after AI adoption is not solely the result of the same underwriters working faster. It also reflects entry, continuation, and reallocation toward underwriters whose operations are more compatible with the new technology. The theory's capacity constraint operates at the underwriter level, but in equilibrium it can also change the composition of active underwriters and expand coverage on the extensive margin.

We next examine whether AI helps underwriters reach more distant policyholders. This is a direct implication of lower transaction costs: when underwriting requires costly travel, physical animal restraint, and manual verification, underwriters may concentrate on nearby households. If AI reduces the time and effort required to verify cattle identity, it may expand the geographic radius over which underwriters can operate.

We re-estimate Equation (19) using the average distance of households served by underwriter u in year t as the dependent variable. As before, we estimate the specification separately for the balanced sample of continuously active underwriters and the full sample of all underwriters, including entrants and exiters.

Figure 11 shows that AI exposure is associated with serving more distant households after AI becomes broadly available. In the full sample, the 2023 coefficient is positive and statistically significant, at roughly 3 kilometers. The coefficients remain positive in 2024

Figure 11: AI Exposure and Average Distance of Underwritten Policyholders



Notes: This figure reports event-study coefficients from underwriter-year regressions. The dependent variable is the average distance, in kilometers, between the households served by an underwriter and the relevant county government. The key explanatory variables are interactions between year indicators and $AI_{u,2023}$, the share of cows processed using AI by underwriter u in 2023. The omitted year is 2022. The balanced sample restricts to underwriters continuously active over the sample period. The full sample includes all underwriters, including entrants and exiters. The regressions include underwriter fixed effects and year fixed effects. Vertical bars report 95 percent confidence intervals.

and 2025, around 3–4 kilometers, although they are less precisely estimated. This pattern suggests that underwriters with greater AI exposure expand geographically after adoption.

The balanced-sample estimates are also generally positive after 2023, but they are less precise. The 2023 coefficient is small and statistically insignificant, while the 2024 and 2025 coefficients are positive but have wide confidence intervals. This difference between the balanced and full samples is informative. Among incumbent underwriters, AI may modestly increase the distance they can cover, but the estimates are noisy. In the full sample, the more persistent positive effects suggest that geographic expansion also operates through entry and composition changes: new or continuing underwriters who use AI may be better able to serve more remote households.

This evidence supports the distributional implication of the transaction-cost channel. Lower underwriting costs do not only increase the total number of policyholders served; they can also change which households are reached. Under the traditional technology, more distant households are more costly to serve and may be excluded when underwriting capacity is limited. By reducing the cost of identification and verification, AI can make it

feasible to cover households farther from administrative centers. This mechanism helps explain why AI may expand access to livestock insurance for policyholders who were previously more costly to reach.

6 Welfare Analysis

7 Conclusions

References

Akerlof, G. A. (1970). The market for "lemons": Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics 84*(3), 488–500.

Andrew, W., J. Gao, S. Mullan, N. Campbell, A. W. Dowsey, and T. Burghardt (2021). Visual identification of individual holstein-friesian cattle via deep metric learning. *Computers and Electronics in Agriculture 185*, 106133.

Banerjee, A. V. and E. Duflo (2010). Giving credit where it is due. *Journal of Economic Perspectives 24*(3), 61–80.

Banerjee, A. V. and A. F. Newman (1993). Occupational choice and the process of development. *Journal of Political Economy 101*(2), 274–298.

Besley, T. (1995). Property rights and investment incentives: Theory and evidence from ghana. *Journal of Political Economy 103*(5), 903–937.

Buera, F. J., J. P. Kaboski, and Y. Shin (2011). Finance and development: A tale of two sectors. *American Economic Review 101*(5), 1964–2002.

Chen, M. A., Q. Wu, and B. Yang (2019). How valuable is fintech innovation? *The Review of Financial Studies 32*(5), 2062–2106.

Cole, S., X. Giné, J. Tobacman, P. Topalova, R. Townsend, and J. Vickery (2013). Barriers to household risk management: Evidence from india. *American Economic Journal: Applied Economics 5*(1), 104–135.

Crocker, K. J. and J. Morgan (1998). Is honesty the best policy? curtailing insurance fraud through optimal incentive contracts. *Journal of Political Economy 106*(2), 355–375.

D'Acunto, F., N. Prabhala, and A. G. Rossi (2019). The promises and pitfalls of robo-advising. *The Review of Financial Studies 32*(5), 1983–2020.

Dercon, S. (2004). Growth and shocks: Evidence from rural ethiopia. *Journal of Development Economics 74*(2), 309–329.

Einav, L., A. Finkelstein, and M. R. Cullen (2010). Estimating welfare in insurance markets using variation in prices. *The Quarterly Journal of Economics 125*(3), 877–921.

Food and Agriculture Organization of the United Nations (2009). *The State of Food and Agriculture 2009: Livestock in the Balance*. Rome: FAO.

Fuster, A., M. Plosser, P. Schnabl, and J. Vickery (2019). The role of technology in mortgage lending. *The Review of Financial Studies 32*(5), 1854–1899.

Giné, X. and D. Yang (2009). Insurance, credit, and technology adoption: Field experimental evidence from malawi. *Journal of Development Economics 89*(1), 1–11.

Hau, H., Y. Huang, C. Lin, H. Shan, Z. Sheng, and L. Wei (2024). Fintech credit and entrepreneurial growth. *The Journal of Finance 79*(5), 3309–3359.

Karlan, D., R. Osei, I. Osei-Akoto, and C. Udry (2014a). Agricultural decisions after relaxing credit and risk constraints. *Quarterly Journal of Economics 129*(2), 597–652.

Karlan, D., R. Osei, I. Osei-Akoto, and C. Udry (2014b). Agricultural decisions after relaxing credit and risk constraints. *Quarterly Journal of Economics 129*(2), 597–652.

Liu, X. (2024). Artificial intelligence and customer acquisition in retail financial services: Experimental evidence from insurance distribution. Working paper, Tsinghua University PBC School of Finance. R&R at The Review of Financial Studies.

Miranda, M. J. (1991). Area-yield crop insurance reconsidered. *American Journal of Agricultural Economics 73*(2), 233–242.

National Bureau of Statistics of China (2024). China statistical data on livestock inventory. National Data, National Bureau of Statistics of China. Accessed via National Bureau of Statistics of China.

Nelson, C. H. and E. T. Loehman (1987). Further toward a theory of agricultural insurance. *American Journal of Agricultural Economics 69*(3), 523–531.

Pauly, M. V. (1968). The economics of moral hazard: Comment. *American Economic Review 58*(3), 531–537.

Picard, P. (1996). Auditing claims in the insurance market with fraud: The credibility issue. *Journal of Public Economics 63*(1), 27–56.

Qiao, Y., C. Clark, S. Lomax, H. Kong, D. Su, and S. Sukkarieh (2021). Automated individual cattle identification using video data: A unified deep learning architecture approach. *Frontiers in Animal Science 2*, 759147.

Rosenzweig, M. R. and K. I. Wolpin (1993). Credit market constraints, consumption smoothing, and the accumulation of durable production assets in low-income countries: Investments in bullocks in india. *Journal of Political Economy 101*(2), 223–244.

Rothschild, M. and J. Stiglitz (1976). Equilibrium in competitive insurance markets: An essay on the economics of imperfect information. *The Quarterly Journal of Economics 90*(4), 629–649.

Shavell, S. (1979). On moral hazard and insurance. *Quarterly Journal of Economics 93*(4), 541–562.

Skees, J. R. and M. R. Reed (1986). Rate making for farm-level crop insurance: Implications for adverse selection. *American Journal of Agricultural Economics 68*(3), 653–659.

Smith, V. H. and B. K. Goodwin (1996). Crop insurance, moral hazard, and agricultural chemical use. *American Journal of Agricultural Economics 78*(2), 428–438.

Stiglitz, J. E. and A. Weiss (1981). Credit rationing in markets with imperfect information. *American Economic Review 71*(3), 393–410.

Townsend, R. M. (1979). Optimal contracts and competitive markets with costly state verification. *Journal of Economic Theory 21*(2), 265–293.

Townsend, R. M. (1994). Risk and insurance in village india. *Econometrica 62*(3), 539–591.

World Bank (2009). *Minding the Stock: Bringing Public Policy to Bear on Livestock Sector Development*. Washington, DC: World Bank.

Online Appendix (Not For Publication)

A Measuring Underwriting Time Per Cow

This appendix describes how we construct a measure of average underwriting time per cow. The measure is designed to capture the amount of time required for underwriters to process cows during a farm visit. We construct comparable measures before and after the adoption of AI-based cattle facial recognition, using the time stamps recorded in the underwriting data.

Traditional ear-tag underwriting. Before AI adoption, underwriters were required to take photos of insured cows during the ear-tag underwriting process. These photos contain time stamps, which allow us to recover the timing of underwriting activities within each household visit. For each household visit, we identify the first and last photo time stamps among the cows underwritten during the visit. We then compute the elapsed time between the first and last photo and divide it by the number of cows processed, net of the first cow. Specifically, for household visit v, we define

$$TimePerCow_v^{ET} = \frac{LastPhotoTime_v - FirstPhotoTime_v}{N_v - 1} \times U_v, \tag{20}$$

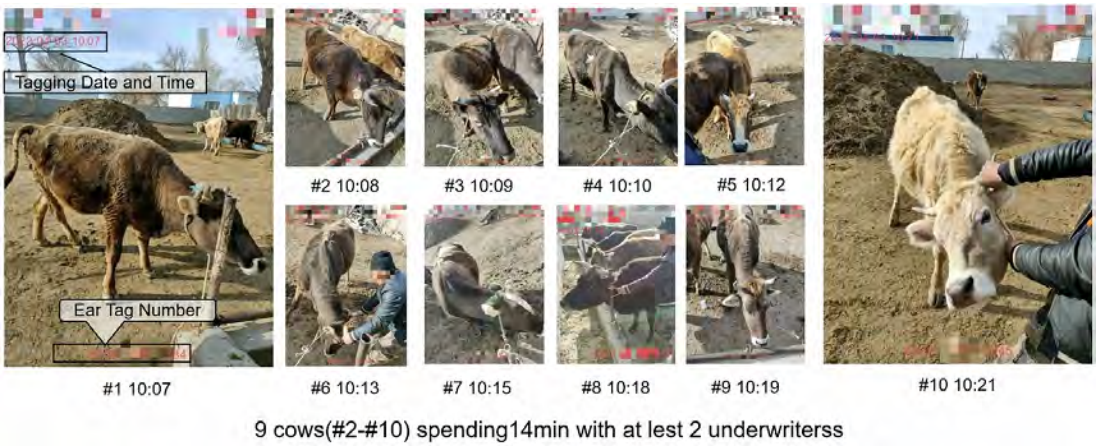
where N_v is the number of cows photographed during the household visit and U_v is the number of underwriters involved in the visit. We use $N_v - 1$ in the denominator because the elapsed time between the first and last photo captures the time intervals between consecutive cows rather than the absolute time spent on the first cow. Multiplying by U_v converts elapsed clock time into underwriter-minutes per cow when multiple underwriters are involved.

Example. Figure 1 illustrates the measurement using a traditional ear-tag underwriting visit. The first photo is taken at 10:07 and the last photo is taken at 10:21. Excluding the first cow, the visit processes nine additional cows over 14 minutes. With at least two underwriters involved, the implied average underwriting time is

$$\frac{14}{9} \times 2 = 3.11 \tag{21}$$

underwriter-minutes per cow. This example highlights the labor-intensive nature of the traditional process: cows must be physically restrained, ear tags must be checked or attached, and multiple underwriters or assistants may be needed during the visit.

Appendix Figure 1: Measuring Underwriting Time Using Photo Time Stamps



Notes: This figure shows an example of a traditional ear-tag underwriting visit. Each photo includes a time stamp. We measure the elapsed time between the first and last photo and divide by the number of intervals between cows, $N_v - 1$. We then multiply by the number of underwriters involved to obtain underwriter-minutes per cow.

AI-based underwriting. After AI adoption, the underwriting data directly record a time stamp for each cow processed using cattle facial recognition. We use the same logic to construct average underwriting time per cow. For each household visit, we calculate the elapsed time between the first and last AI record and divide by the number of processed cows minus one. When information on the number of underwriters involved is available, we multiply by the number of underwriters to express the measure in underwriter-minutes per cow:

$$TimePerCow_v^{AI} = \frac{LastAIRecordTime_v - FirstAIRecordTime_v}{N_v - 1} \times U_v. \tag{22}$$

This creates a comparable measure of underwriting time under traditional ear tags and AI-based cattle facial recognition.

Sample restrictions and implementation. In constructing this measure, we restrict attention to household visits with at least two cows, since the elapsed-time method requires at least two time-stamped observations. We also exclude visits with missing or invalid time stamps and winsorize extreme values to reduce the influence of data-entry errors, unusually long interruptions, or visits that combine multiple unrelated tasks. When a household has multiple underwriting visits within a year, we compute visit-level underwriting time first and then aggregate to the household-year or underwriter-year level using cow-count weights.

Interpretation. The resulting measure should be interpreted as the average amount of underwriter time required to process an additional cow during a household visit. It is not a pure measure of the physical time spent attaching an ear tag or scanning a cow. Instead, it captures the combined time required for identification, recording, and verification during the underwriting process. This is the relevant object for our analysis because the model's transaction-cost channel concerns the effective labor and time cost of underwriting cows. A decline in this measure after AI adoption provides direct evidence that AI lowers underwriting transaction costs and relaxes underwriter capacity constraints.