

INSTITUTE FOR INDIVIDUAL AVIATION

IIA



COST CONTROL · CASE CC-04

Capability Illegibility and Premium Pooling in Operator Insurance

The market cannot see the behaviours that reduce claims, so it prices coarse proxies and pools everyone; the capable operator subsidises the pool, and hours earned inside an organisation's architecture are credited at face value in the cockpit that has none of it.

DOC NO. IIA_Case_CC04_2026	ISSUE 2026 · REV B	CLASS PUBLIC
-------------------------------	-----------------------	-----------------

Abstract

This case studies why the insurance premium of an individual light-aviation operation does not reflect the operation's capability. The finding: the market sits inside four structural conditions, thin data, actuarial illegibility of the behaviours that actually reduce claims, systematic miscalibration of experience credit, and a claim-cost base substantially outside the operator's control, and under those conditions it can only price coarse proxies and pool the results. The consequences run one way: the capable operator pays for development they receive no return on, the less capable is not penalised for its absence, hours accumulated inside organisational architecture are credited at face value in the cockpit that has none of it, and the losses that miscalibration produces are absorbed into the aggregate every policyholder pays. The case documents the conditions, credits what the market genuinely cannot do and why, issues a legibility recommendation the operator can execute alone, and states its limitations. No fault is found with any underwriter; the conditions are features of the market's structure.

1. The Practice and Its Scope

The practice under study: the pricing of individual light-aviation insurance on visible proxies, hours flown, ratings held, training records, claims history, in place of the behaviours that distinguish capable from less-capable operators, which produce no signal the underwriter can price. The premium that results describes a pool, not an operation, and the direction of every error is the same: away from crediting capability.

Scope: the structure of the insurance market for individual operators and the premium consequences of its illegibility, worldwide, since the conditions are structural rather than national. The operational side of over-credited organisational experience, how the halo misleads oversight and self-assessment, is separate case territory and is not treated here; the configuration differential the market cannot price is the Institute's issued case SM-04 and enters only as the record. This case concerns what the market can see, what it cannot, and who pays for the difference.

2. Sources

- The public structure of the individual-operator aviation insurance market: policyholder population scale, claim frequency and severity characteristics, underwriting capacity cycles, and the observable pricing factors of the application process.
- AOPA McSpadden series and NTSB analyses for the accident pattern of high-hour, organisationally trained pilots operating personal light aircraft, the empirical contradiction of the hours proxy.

- The skill-currency literature, Childs, Spears and Prophet (1983) and successors: proficiency decays with absence of practice regardless of total hours, the mechanism beneath the miscalibration.
- The Institute's issued cases CC-01 and CC-02 for the structural share of claim cost: parts pricing, certified-labour rates, and hull repair economics on the ageing fleet, which set claim severity outside the operator's control.
- IIA research library: the doctrinal analysis of the insurance pooling trap and the structural market analysis, from which the four-condition inventory of Section 4 is drawn.

3. Method

The method is a legibility audit: for each behaviour the safety record associates with reduced risk, recurrent training depth, weather discipline, mechanical engagement, decision-making patterns, type and terrain familiarity, the audit asks whether the behaviour produces a signal an underwriter can observe and price; and for each factor the underwriter does observe, the audit asks whether it reliably tracks claims at the operator class actually paying the premium. The gap between the two lists is the finding.

The audit imputes nothing to any underwriter: a market can only price what it can see, and the case's question is what the structure lets it see. Two results are reported per the Institute's standard: the magnitude, stated as the four-condition inventory and the direction of its consequences, and the sensitivity, what remains after every proxy is credited with the partial validity it genuinely has.

4. Findings

Finding one, the four conditions. Data thinness: the policyholder population is small, the annual claim count is small, outcomes are severe and variable, and the regulatory environment is fragmented across jurisdictions; the actuarial dataset that would support fine-grained differentiation is not available to any single underwriter and is unlikely to be aggregated under any current arrangement. Actuarial illegibility: the behaviours that distinguish capable operators, recurrent training depth, weather discipline, mechanical engagement, decision-making patterns, type and terrain familiarity, produce no signal the underwriter can price; whether an operator practises divert decisions in actual marginal weather is invisible to every application form in the market. The mismatch is one of measurement: the operator provisions from money, time, competence and room, and the form counts money. Miscalibration of experience credit: hours accumulated inside organised aviation, with its crew, dispatch, automation, and oversight, are credited at face value when they enter personal aviation without that architecture; the high-hour organisationally trained pilot in a personal aircraft is overrepresented in the accident pattern, is priced at the discount the hours suggest, and the underwriting losses that follow are absorbed into the aggregate every policyholder pays. Claim severity asymmetry: a meaningful share of

claim cost is set by structural features outside the operator's control, the parts multiples, certified-labour rates, and ageing-airframe repair economics the Institute's cases CC-01 and CC-02 measure, plus tort exposure, so even perfect behaviour cannot reach a large fraction of the premium.

Finding two, the consequences, and they run one way. The market prices coarsely; the pool drifts toward less capability, because the capable receive no return and the less capable no penalty; premiums rise to match the drifting pool; underwriting capacity periodically exits, disrupting coverage regardless of any policyholder's behaviour; and territorial coverage contracts in step with the operator class. The deepest consequence is developmental: the capability the safety record actually depends on is not built by the market, because the market cannot see it, and the operator whose capability was developed inside personal aviation from the start subsidises the losses produced by over-credited organisational hours. That operator paid in the constituents of the provisioning budget the instrument cannot count.

Finding three, the export test. None of the four conditions is national. Data thinness worsens as markets shrink, the multi-jurisdictional fragmentation is itself condition one, capacity cycles run through the same international underwriting centres, and the miscalibration of organisational hours travels with the pilots who carry them. The mechanism is the market's structure, and the structure is the same everywhere the operator class exists.

Findings in the other direction, per the Institute's standard. Aggregate pricing under thin data is rational, not negligent: no underwriter can price a signal that does not exist, and the conditions are features of the structure, with no fault found and no intent imputed. The proxies hold partial validity: claims history is genuinely informative, ratings mark discrete verified events, and premium credits for documented recurrent training and type-club programmes exist across parts of the market, real, discrete, and credited here. Differentiation in adjacent insurance lines shows the direction is not impossible, while the adjacency differences, frequency, severity, and population scale, are real and are why the transfer has not happened on its own. Sensitivity: grant every proxy its partial validity in full, and the core findings stand: the behavioural capability that moves the record remains unpriceable, the organisational-hours miscalibration remains contradicted by the accident pattern, and the structural share of claim cost remains outside the operator's reach.

5. The Mechanism

The trap persists because no actor inside it can move first. The underwriter who priced capability without the data would be selected against by the market; the operator who develops capability receives no premium signal to continue; the dataset that would change both is held by no one and built by no one, because building it is a cost with a collective benefit and an individual bill. The equilibrium is stable precisely because everyone in it is behaving reasonably.

Four cover propositions maintain it. That aviation is intrinsically high-risk and aggregate pricing is appropriate: true at the aggregate, silent on differentiation, which exists in adjacent lines. That the existing credit structure already differentiates: the

credits address discrete observable events, not cumulative behavioural capability. That hours flown are the best available proxy: contradicted by the accident pattern of high-hour organisational pilots in personal aircraft. That better differentiation would be too intrusive or too expensive: a reframing of structural absence as deliberate calibration. Each proposition is true where it is narrow and cover where it is broad.

6. Recommendation

Issued to individual operations, in aviation's non-mandatory form: adopted voluntarily, at the adopter's responsibility. The floor here is contractual as well as regulatory: full and honest disclosure to insurers, always; the discipline below is legibility, never misrepresentation.

The recommendation is a capability-legibility discipline: since the market cannot see capability, show it. What has to become legible is the rest of the provisioning: the time, the competence, and the room the operation spends and no form records. First, build the auditable record: document recurrent training beyond minima, instruction and dual share per the configuration record, type-specific transition work, and mechanical engagement, in artefacts an underwriter or broker can hold, certificates, logbook endorsements, programme completions. Second, present the record: place it before the market through a specialist broker at every renewal, and ask explicitly which documented behaviours the underwriter will credit; credits for documented training exist in parts of the market and are found by asking. Third, shop the thin market: quotes in a small market vary materially, and the annual presentation of a legible operation to several underwriters is the one pricing pressure an individual can apply. Fourth, manage the structural share knowingly: hull value election, deductible level, and liability-only cover on ageing airframes are decisions about the share of premium no behaviour can reach, made best with the operation's own numbers on the table. The same discipline reaches the planning table: the insurability and structural claim-cost profile of an aircraft, its parts multiples and repair economics, belong in the acquisition-path analysis before the aircraft is chosen.

7. Gain Estimate Methodology

The gain claims published with this case are constructed as follows, with their honesty conditions. The case publishes no universal percentage, because none exists: credits and shopping spreads are market- and jurisdiction-dependent, and inventing a representative figure would fabricate precision the market's own thinness forbids.

Individual gain. Two bounded components. Documented-training credits, where offered, are real and discrete; their availability is established by presentation and asking, at the cost of the paperwork the operation should hold anyway. The shopping spread of a thin market is the larger component: quotes for the same risk vary materially across underwriters, and an annually presented, legible operation captures the low end of its own spread. The published claim: legibility plus annual presentation captures whatever differentiation the market currently offers, at near-zero cost beyond

record-keeping, and the record built for the insurer is the provisioning account the operation should keep for itself.

Cumulative gain. Stated as the missing dataset: condition one is that no underwriter holds the data fine-grained differentiation would need; every legible operation is a fragment of that dataset, and a population of documented operations is the precondition for any future capability-priced market. The case publishes no figure for the pooling's aggregate cost, because the subsidy flows inside aggregates no party publishes, which is itself a finding.

Secondary benefit. A market that could see capability would build it: premium signals for documented development would fund the recurrent practice the safety record depends on. The direction is claimed; the magnitude is not. This case does not quantify the benefit.

Sensitivity and honesty conditions. The credit and spread components are stated as market practice with variation, not as entitlements; a market cycle that hardens can swamp both, and capacity exit disrupts coverage regardless of any operator's conduct, per finding two. The gain is procedural: the case claims legibility and presentation as the mechanism, not a promised premium level.

8. Limitations

The four-condition inventory is drawn from the market's public structure and the doctrinal analysis, not from underwriters' internal files, which are not observable; the miscalibration finding rests on the published accident pattern and the skill-currency literature rather than on claims data, which no party publishes at the resolution required. Credit availability and shopping spreads vary by market, cycle, and jurisdiction, and are stated as directions with mechanisms rather than as rates. The recommendation's force is bounded by the structural share of claim cost, which no operator behaviour reaches, and by market cycles, which no operator controls. None of this touches the central claims: that the behaviours that reduce claims are invisible to the market, that the visible proxies run unreliably at this operator class, that organisational hours are over-credited into an architecture they were not earned in, and that the resulting pool is paid for by the operators it least rewards.

References

Public structure of the individual-operator aviation insurance market: application factors, capacity cycles, and pricing practice as observable from the market's own instruments.

AOPA Air Safety Institute, Richard G. McSpadden Report series; NTSB analyses. Accident pattern of high-hour organisationally trained pilots in personal light aircraft.

Childs, J. M., Spears, W. D., and Prophet, W. W. (1983). Private pilot flight skill retention. FAA/Embry-Riddle, and successor skill-currency studies.

Institute for Individual Aviation. Issued cases CC-01 (parts and services multiples), CC-02 (acquisition value share), and SM-04 (dual-occupancy differential), cited for the structural claim-cost base and the configuration record.

Institute for Individual Aviation, research library: doctrinal analysis of the insurance pooling trap; structural market analysis. Four-condition inventory of Section 4.