

Independent Research Desk — Worked Example

(End to End)

Status: Illustrative. One invented company, invented data, traced through every stage of the systems design's end-to-end flow, using the components and data model from the architecture blueprint and the stack from the technical view.

The company, the ticker, the filings, the figures, and every quoted line below are fictional — invented for this walkthrough, not a real business, not investment analysis, not a claim about any real security. The point is to show the *shape* of what the pipeline actually produces at each stage — not to demonstrate real research quality, which depends entirely on the operator doing it.

Output format: probability-weighted scenarios (Bull / Base / Bear), a probability-weighted intrinsic value (PWIV), margin of safety (required vs. actual), and high/low trigger levels, with a stated conviction level — not a categorical rating. No verdict, no label. The operator gets the reasoning, the range, and the price, and decides for themselves what to do with it.

LiquidBlack's own system carries a further classification layer on top of this same output — that layer isn't described here, for legal reasons, and isn't part of this blueprint. An independent desk built from this document stops at PWIV, margin of safety, triggers, and conviction, which is a complete, defensible basis for an operator's own decision without needing anything further layered on top of it.

The example is deliberately unresolved in one place, on purpose: the Executive Report doesn't simply adopt the Investment Thesis's probability weighting. That's the no-synthesis principle (architecture blueprint, Section 1.5) shown in motion rather than just asserted.

Stage 0 — Screen

Company: Harrow Fastener Group ("HFG"), invented mid-cap industrial fastener manufacturer.

Screen record:

Field	Value
Company	Harrow Fastener Group (HFG)
Market	US-listed, mid-cap industrial
Reason on the list	Trading below book value after a two-quarter earnings miss; a director purchased shares on the open market four weeks after the miss

That last line is the whole reason it's here — not "looked cheap," a specific, checkable trigger (the insider purchase, which is itself a filed disclosure).

Stage 1 — Approval Gate

Operator reads the one-line reason, approves coverage. Timestamped. This is the only step in the whole cycle that needed a person — nothing below this line ran until this happened.

Stage 2 — Research Report

The Research module pulls HFG's primary filings and extracts Claims. A representative sample of what gets stored (illustrative — a real research pass would produce several dozen, not five):

Claim	Source
FY2025 revenue \$412M, down 3% YoY	10-K FY2025, Item 7 (MD&A), p.34
Gross margin compressed to 28.1% from 31.4% the prior year, attributed to steel input costs	10-K FY2025, Item 7, p.36
Director J. Averill purchased 15,000 shares on the open market on 2026-03-11	Form 4, filed 2026-03-13
Largest customer accounts for 22% of revenue, under a contract expiring 2027	10-K FY2025, Item 1A (Risk Factors), p.11
Net debt / EBITDA of 1.4x, covenant ceiling of 3.0x per credit agreement	10-K FY2025, Note 8 (Debt)
Trading at 0.81x book value as of the Research Report's run date	Derived: filed book value ÷ market price at close
Price used throughout this example: \$19.00 (USD), close of 2026-03-14	Market data, as-of date stated for reproducibility

Research Report excerpt:

HFG's margin compression is fully explained by input costs, not a demand problem — unit volumes were flat year over year. The insider purchase came four weeks after the miss was reported, not immediately, which is more consistent with a considered view than a reflexive buy. Customer concentration (22%, one contract) is real and worth carrying into valuation.

Every sentence in that excerpt traces to a Claim above. This is what Section 1.1's core principle looks like as an actual document, not a rule.

Stage 3 — Investment Thesis

Reads the Research Report and its Claims. Builds the case as probability-weighted scenarios, not a single number:

Case	Intrinsic value	Probability	Assumption
Bear	\$18.00	25%	Margin stays compressed; concentrated customer doesn't renew
Base	\$27.00	50%	Margin partially recovers as input costs ease; customer renews on similar terms
Bull	\$35.00	25%	Full margin recovery to historical ~31%; customer renews with expanded volume

PWIV (probability-weighted intrinsic value): $0.25(18) + 0.50(27) + 0.25(35) = \mathbf{\$26.75}$

Margin of safety: required 35% · actual $(26.75 - 19) / 26.75 = \mathbf{29.0\%}$ — short of the required threshold at the current price.

Trigger — Low: \$17.40 — the price at which actual MOS reaches the required 35% against this PWIV $(26.75 \times 0.65 = 17.3875, \text{rounded up to the nearest } 5c)$. **Trigger — High:** \$31.50 — 90% of the Bull case's intrinsic value, a re-check point if price approaches how much upside is already priced in.

Conviction: Medium — the base case carries the plurality of the weight, but the case rests on margin normalizing within 4–6 quarters, which isn't yet observed, only assumed.

Investment Thesis record (stored): the table above, PWIV, MOS, both triggers, conviction level.

Stage 4 — Adversarial Report

Does **not** just read the Thesis and poke at its prose — goes back to the Claims directly, the same ones the Thesis used, looking for what the Thesis's weighting didn't account for:

The customer contract (22% of revenue) expires in twelve months from today, not "eventually" — that's a dated, specific event, and it belongs in the Bear case's probability, not just its prose. The Bear and Bull weights (25%/25%) treat renewal risk as a coin flip against an equally likely upside scenario, but only one of those two scenarios has a filed deadline attached to it. Reweighting toward Bear is the defensible read of the same Claims.

Separately: the covenant headroom (1.4x against a 3.0x ceiling) is genuinely comfortable, and the Thesis is right not to weight it as a risk. The margin-recovery logic itself is reasonable and not contested here — only the probability weighting is.

Adversarial reweighting (for comparison, not adopted automatically): Bear 35% / Base 50% / Bull 15% → PWIV \$25.05.

Stage 5 — Executive Report

Reads the Investment Thesis and Adversarial Report **for context**, but returns to the Claims itself rather than adopting either prior document's weighting wholesale:

The Adversarial Report's core point holds — the twelve-month contract clock is a dated, specific risk that the Thesis's original weighting underpriced. But the Adversarial Report's own reweighting (Bear 35%) moves further than the Claims themselves support: the covenant headroom (1.4x against a 3.0x ceiling) is real cushion against the severity of a non-renewal outcome, even if it doesn't change its probability. Weighing the dated risk without also weighing the covenant cushion lands closer to Bear than the Claims justify.

Case	Intrinsic value	Probability
Bear	\$18.00	32%
Base	\$27.00	48%
Bull	\$35.00	20%

PWIV: $0.32(18) + 0.48(27) + 0.20(35) = \25.72 . **Margin of safety:** required $35\% \cdot \text{actual} (25.72 - 19) / 25.72 = 26.1\%$. **Trigger — Low:** $\$16.75 - 25.72 \times 0.65 = 16.718$, rounded up to the nearest 5c. **Trigger — High:** $\$31.50$ — unchanged from the Investment Thesis, since that rule depends only on the Bull case's dollar value, and the Bull case never moved. **Conviction: Medium** — unchanged in category, but for a different reason than the Thesis's Medium: here it reflects a specific, dated uncertainty rather than general margin-recovery timing.

This is the point of the whole exercise, and the arithmetic has to survive scrutiny for the point to mean anything: the mean of the Thesis's weighting (25/50/25) and the Adversarial Report's weighting (35/50/15) is exactly 30/50/20. The Executive Report's actual weighting — 32/48/20 — is close to that average but not identical to it, and the difference is where the real reasoning lives. Bear moved further than a simple average would (32 vs. 30), because the twelve-month contract clock is a hard, dated risk, not a probability to split down the middle. But the offsetting movement came out of Base, not Bull, because the covenant-headroom argument affects the *severity* of a bad outcome, not the *likelihood* of the recovery scenario. A synthesis step that simply averaged the two prior weightings would have landed close to the right numbers for the wrong reason — arithmetic, not the covenant argument — and would have had no principled basis for deciding whether Base or Bull should absorb the shift.

Stage 6 — Operator Review and Decision

The operator now has all four documents — Research Report, Investment Thesis, Adversarial Report, Executive Report — none of them hidden or discarded. No verdict is delivered anywhere in the chain; what's on the table is the reasoning, the range, and the price. Reads where the Investment Thesis (\$26.75 PWIV, 29.0% MOS) and Executive Report (\$25.72 PWIV, 26.1% MOS) diverge, and why — a specific, dated contract risk the Thesis's original weighting underweighted. Records the final call in their own words:

Decision record: No position initiated at the current price (\$19, against a required MOS not yet met on either document's numbers). Re-check at the earlier of: the contract-renewal announcement, or two further quarters of margin data. Timestamped, immutable.

What Got Stored, End to End

Tying back to the data model (architecture blueprint, Section 2.2):

Record type	Count in this cycle
Company	1 (HFG)
Source Document	3 (10-K, Form 4, credit agreement note referenced within the 10-K)
Claim	6 shown above (illustrative — a real pass would run higher)
Reports	4 (Research, Investment Thesis, Adversarial, Executive)
Decision	1 (no position initiated, two named re-check triggers)

Nothing here required more than the technical view's stack (Section 1, technical view document) — one machine, one reasoning API, one SQLite file. The entire cycle above, if run for real, would be a handful of API calls and a few dozen database rows. The value isn't in the infrastructure; it's in the fact that Stage 5's disagreement with Stage 3 survived to the operator instead of getting smoothed into a single soft recommendation somewhere upstream.

Note. This document is a conceptual reference prepared by Liquidblack.ai for illustrative purposes. It does not describe Liquidblack.ai's own production system, is not a build specification, and carries no guarantee of outcome for any operator who builds from it. It is not financial, investment, or technical advice. © 2026 Liquidblack.ai.