

Practice Bank — Human × AI Drills

The drill bank for *modules/human-ai-collaboration.md*. Where the module teaches you to brief the machine, review its output, and hold it accountable, these drills force you to do it — on machine output that contains **planted errors**. Find the error, say why it is wrong, and log an override.

How to use this bank. Work with *literacy/how-the-machine-learns.md* and the module open. Every drill is a self-contained machine output with a **planted-error list** — the errors are disclosed up front, and the pass bar is finding all of them, stating why each is wrong, and logging an override for each with a rationale. The answer keys are in Part B at the bottom — do not open them until you have written your own findings. No software is required (the stack-agnostic rule, *doctrine/08*). All agencies and companies in these drills are invented; **Ravonics LLC** is the one real teaching object.

How to log an override. Every drill ends with an override log. Use the four-column form from Session 4 of the module:

Machine claim	The evidence that contradicts it	Your action	Your rationale
<i>what the machine asserted</i>	<i>where the truth lives</i>	<i>override / accept-with-change / reject</i>	<i>why a professional should not act on the machine's claim</i>

Part A — The drills

Drill HA-01 — The Reference Audit

Situation. Ravonics is responding to a solicitation from the **National Border Health Administration (NBHA)** — Solicitation NBHA-26-007, a data-integration and analytics services contract. The machine has drafted the opening of Volume I, Technical Approach.

The machine output (abridged).

Ravonics LLC proposes a border-health data integration and analytics service for the National Border Health Administration (NBHA) under Solicitation NBHA-26-007. Our approach is fully compliant with all applicable acquisition regulations, including **FAR 52.219-27** (Notice of Small Business Subcontracting Plan)

and **FAR 52.226-13** (Data Integration Services Set-Aside). As required by **Section M, paragraph 5.2**, “the offeror must demonstrate prior experience deploying analytics platforms at scale”; we satisfy this requirement through our record of federal analytics deliveries. The proposed effort will be delivered on a fixed-price basis with a **12-month** period of performance, per Section B. We will submit via the NBHA Submission Portal no later than **5:00 PM Eastern** on the submission date identified in Section A.2.

The planted-error list. This output contains five planted errors, all in the reference-and-verification class: 1. A fabricated FAR clause. 2. A real FAR clause cited where it does not apply. 3. A misquoted Section M requirement. 4. A wrong period of performance. 5. A wrong submission deadline.

Your task. Find all five. For each, state why it is wrong (which checklist item it fails) and log an override with the correction. Verify against the source — in this drill, “the source” is what the brief and the record actually contain: the NBHA solicitation requires a **9-month** performance period, submission by **2:00 PM ET**, and Section M’s evaluation language describes the analytics criterion without imposing a pass/fail “prior scale deployment” requirement; the record contains no subcontracting plan requirement and no “Data Integration Services” clause.

Drill HA-02 — The Overstated Win Theme

Situation. Ravonics is pursuing a solicitation from the **Administration for Veterans Employment Systems (AVES)** — AVES-26-114, a workforce-systems modernization contract. The machine has drafted the “Win Strategy” section of the capture plan.

The machine output (abridged).

Ravonics’s win themes for AVES-26-114: 1. We are **the only small business with a combined physics, AI, and systems-engineering bench in the HUBZone program**. 2. Our team has a proven track record of transforming federal workforce systems — **we have implemented workforce platforms at three federal agencies**. 3. We offer the industry’s best value with **an unbeatable price**. 4. Our **ISO 9001 certified** quality management system ensures defect-free delivery. 5. We are uniquely positioned because we hold both HUBZone and STTR eligibility, **a combination no competitor can match**.

The planted-error list. This output contains five planted errors: 1. An over-stated uniqueness claim. 2. A fabricated past-performance claim. 3. An unsupported “unbeatable price” theme. 4. A baseline dressed as a discriminator. 5. An over-stated combination claim.

Your task. For each win theme, classify it — discriminator, baseline, overstatement, or unsupported claim — state why, and log an override. The record contains: Ravonics’s federal delivery experience is thin and laboratory-affiliated (no federal workforce-platform implementations); its certifications are HUBZone and STTR eligibility; and no uniqueness audit has been run against the competitive set.

Drill HA-03 — The Past-Performance Narrative

Situation. Ravonics is pursuing a solicitation from the **Office of Arctic Resilience Programs (OARP)** — a climate-resilience infrastructure planning contract. The machine has drafted the past-performance volume.

The machine output (abridged).

Past Performance Reference 1 — Arctic Environmental Monitoring, U.S. Department of the Interior (2024–2025). Ravonics served as prime contractor delivering sensor-data integration and analytics for a portfolio of Arctic monitoring sites. **CPARS rating: Excellent in all six evaluation areas.** Contract value: \$1.4M. **Customer reference: Dr. Lena Voss, Program Manager, DOI Arctic Office, (907) 555-0142.**

The planted-error list. This output contains four planted errors: 1. A fabricated contract. 2. An invented CPARS rating. 3. An invented customer reference. 4. A relevance overstatement.

Your task. Audit the past-performance reference. State what is fabricated versus verifiable, why each fabrication is dangerous (what would happen if it reached the agency), and log an override for each. The record contains: no DOI Arctic monitoring contract, no CPARS record at all, no reference named “Dr. Lena Voss,” and no Arctic environmental-monitoring delivery in Ravonics’s history.

Drill HA-04 — The Price-to-Win Memo

Situation. Ravonics is pursuing a solicitation from the **Department of Domestic Water Infrastructure (DDWI)**. The machine has produced the price-to-win memo.

The machine output (abridged).

Recommended price: \$118/hour. This is a competitive round number within the expected band and slightly below our most likely competitor. We should bid at \$118/hour because a lower price signals commitment. **Our cost floor is estimated at \$90/hour.**

No adjustments are needed. Winning this contract is important, so we should be aggressive.

The planted-error list. This output contains five planted errors: 1. A cost floor with no basis of estimate. 2. A competitive claim with no source. 3. A non-sequitur “signals commitment” rationale. 4. A missing evaluation-method analysis. 5. An “aggressive because winning matters” posture.

Your task. Audit the memo. Find each unsupported number and each broken reasoning link, state why it is wrong (what doctrine/05 and the machine-rendering BOE rule require), and log an override. The record contains: no labor-category build-up, no burden rates, no competitor price data, and the solicitation has not yet been read for whether it is lowest-price-technically-acceptable or best-value.

Drill HA-05 — The Compliance Sweep

Situation. Ravonics is pursuing a **NOFO** from the **Bureau of Climate Resilience Infrastructure (BCRI)** — a community-resilience research grant. The machine has produced the compliance matrix.

The machine output (compliance matrix, abridged).

#	Requirement	Source	Status
1	Active SAM.gov registration with current UEI and CAGE	§C.1	mapped
2	UEI and CAGE cited on the cover page	§C.1	mapped
3	Technical volume 12 pages, 11-point font	§D	mapped
4	Budget and budget justification on the provided template	§D	mapped
5	Work plan 4 pages	§D	mapped
6	Submission via the portal by 2026-09-30, 5:00 PM ET	§D	mapped

The planted-error list. This output contains five planted errors: 1. A wrong page limit and font. 2. A wrong submission deadline. 3. A missing required attachment. 4. A missing required certification. 5. A fabricated requirement.

Your task. Audit the matrix against the source. Find the misses and the fabrication, state why each matters, and log an override. The actual NOFO requires: a Technical volume of **15 pages max, 12-point font**; a **Data Management and Sharing Plan** as a required attachment; a completed **Representations and Certifications** attestation; submission by **2026-09-30, 11:59 PM ET**; and it contains **no 4-page work-plan requirement**.

Drill HA-06 — The Score Calculation

Situation. Ravonics is pursuing a solicitation from the **Federal Logistics Command (FLC)** — estimated value **\$200,000**. Ravonics’s threshold, set in advance: **pursue only if pWin 0.42 and composite \$84,000** ($0.42 \times \$200,000$). The machine has produced the score memo.

The machine output (abridged).

pWin factors: Topic fit 0.30 (weight 25%), Team 0.60 (20%), Past performance 0.25 (20%), Technical approach 0.40 (20%), Capacity 0.50 (15%).

Weighted pWin = $0.30 \times 0.25 + 0.60 \times 0.20 + 0.25 \times 0.20 + 0.40 \times 0.20 + 0.50 \times 0.15 = 0.100 = 0.075 + 0.120 + 0.050 + 0.080 + 0.100 = 0.425 \rightarrow \text{pWin } 0.43$.

Composite = $0.43 \times \$200,000 = \$86,000$ $\$84,000 \rightarrow$ threshold pass.
Since the composite passes, the pWin floor is effectively redundant — the composite is what matters. Recommendation: **BID**.

The planted-error list. This output contains four planted errors: 1. An arithmetic slip in the weighted sum — the capacity term is computed with the wrong weight. 2. A threshold fallacy (the “either/or” test). 3. A wrong gate call. 4. A missing sensitivity check.

Your task. Recompute the weighted pWin by hand. State the correct pWin, the correct composite, the correct gate call against both parts of the threshold, and log overrides for the arithmetic slip and the threshold fallacy.

Drill HA-07 — The Ghost Theme That Names the Competitor

Situation. Ravonics is pursuing the DDWI solicitation from Drill HA-04. The incumbent on the related program is **Aquifer Systems Corporation**. The machine has drafted a paragraph intended as a “ghost theme” against the incumbent.

The machine output (abridged).

Where the incumbent, **Aquifer Systems Corporation**, relies on a legacy platform with limited sensor integration, Ravonics offers a modern data-fusion architecture. Unlike **Aquifer**, we do not require months of custom integration. Our team has delivered water-infrastructure analytics to federal customers, so we understand the mission. We are confident **Aquifer** cannot match our delivery speed.

The planted-error list. This output contains four planted errors: 1. A ghost theme that names the competitor. 2. An unsupported competitive claim. 3. An overstated delivery claim. 4. A tone/attack risk.

Your task. Identify each ghost-theme violation, state why it is wrong (doctrine/09’s ghost-theme rules), rewrite the passage as a true ghost theme, and log overrides. The record contains: no water-infrastructure analytics delivery to federal customers, and no capture-intelligence event supporting the “cannot match our delivery speed” claim.

Drill HA-08 — The Competitive Landscape

Situation. Ravonics is pursuing the AVES solicitation from Drill HA-02. The machine has produced the competitive-landscape memo.

The machine output (abridged).

Competitive assessment for AVES-26-114: 1. **Granite Peaks Federal will bid at approximately \$97/hour** — we estimate this from their historical pricing posture. 2. **Helixfield Corporation is unlikely to bid** because they have not bid on AVES work in two years. 3. **Meridian Anchor will submit a weak technical approach** because they lack senior staff. 4. **We are the incumbent on the related AVES-25-003 contract** and hold a strong incumbent advantage.

The planted-error list. This output contains four planted errors: 1. A fabricated pricing intelligence claim. 2. An unsupported no-bid prediction. 3. An unsupported competitor-weakness claim. 4. A fabricated incumbency claim.

Your task. Classify each claim as sourced, unsourced, or fabricated. For the claims you would keep, state the source you would require. For the rest, log an override. The record contains: no price data on Granite Peaks, no AVES bidding history on Helixfield, no evidence on Meridian Anchor’s bench, and no AVES contract in Ravonics’s own history.

Drill HA-09 — The Integrated Capture Plan (the exam drill)

Situation. Ravonics is pursuing a solicitation from the **Federal Logistics Command (FLC)** — FLC-26-210, estimated value **\$400,000**. Ravonics's threshold: pWin 0.42 and composite **\$168,000**. The machine has produced a complete one-page capture plan.

The machine output (the full plan).

FLC-26-210 Capture Plan. Win strategy: our theme — “**the only HUBZone small business delivering integrated logistics analytics at federal scale.**” Discriminators: (1) **ISO 9001 certification**; (2) **our proprietary logistics model**; (3) **our record of federal logistics delivery**. Price posture: **\$88/hour, competitive and slightly below Helixfield**. Compliance: all requirements mapped; technical volume fits within **40 pages**; submission by **5:00 PM ET, 2026-11-14**. pWin: **0.68**; composite: **\$170,000** vs threshold \$168,000 → **BID**. Past performance: **FLC logistics support contract (2025), CPARS “Outstanding.”** Team: two senior analysts and one PhD.

The planted-error list. This output contains eight planted errors — a mix of every drill in this bank: 1. An over-stated win theme. 2. A baseline dressed as a discriminator. 3. An unsupported proprietary-model claim. 4. A fabricated past-performance entry. 5. An unsupported competitive price claim. 6. Two compliance misses (wrong page limit, wrong deadline). 7. A pWin/composite inconsistency. 8. A gate call with no discipline.

Your task. Run the full five-point review checklist on the capture plan. Find all eight planted errors, state why each is wrong, log an override for each, and give the correct gate call. The actual Section L limit is **35 pages**; the actual submission deadline is **11:59 PM ET**; and the internal inconsistency to catch: at a \$400,000 value, a pWin of 0.68 gives a composite of \$272,000 — not \$170,000. Recompute, and account for the fabricated past-performance entry when you argue the true pWin.

Part B — The answer keys

Honor system. These keys exist so a facilitator can run the drills live and a solo student can grade honestly. Read the drill, write your findings, log your overrides — *then* open the key for that number. An override you log before reading the key is worth ten you copy afterward.

Answer key — Drill HA-01

Planted error	Why it is wrong	Model override entry
Fabricated FAR clause — “FAR 52.226-13 (Data Integration Services Set-Aside)”	No such clause exists in the FAR. A fabricated citation is the first hallucination signature; it fails the source check . If it reached the agency, it would read as either sloppy or dishonest — both fatal to credibility.	<i>Claim:</i> “FAR 52.226-13 applies.” <i>Evidence:</i> no such clause exists in the FAR. <i>Action:</i> reject, remove the citation. <i>Rationale:</i> never cite a clause you have not read in the FAR; the compliance matrix carries only real requirements.
Inapplicable real clause — “FAR 52.219-27 (Notice of Small Business Subcontracting Plan)”	The clause is real but irrelevant here: it applies when a subcontracting plan is required, and the record and solicitation contain no subcontracting-plan requirement. Citing a real clause in the wrong place is a reference check failure — it signals the machine does not understand <i>which</i> rule applies, not just that it knows a rule.	<i>Claim:</i> “FAR 52.219-27 requires a subcontracting plan.” <i>Evidence:</i> no subcontracting-plan requirement exists in Section L. <i>Action:</i> override, remove. <i>Rationale:</i> cite only requirements in the actual solicitation; a real clause misapplied is a compliance error, not a strength.
Misquoted Section M — inventing “the offeror must demonstrate prior experience deploying analytics platforms at scale”	Section M describes the evaluation criterion without imposing this pass/fail requirement. Inventing a quote out of the evaluation language is a reference check failure — and a dangerous one, because it converts a narrative criterion into a fake “shall.”	<i>Claim:</i> “Section M 5.2 requires prior scale deployments.” <i>Evidence:</i> Section M’s language describes the analytics criterion; it does not require prior scale deployment. <i>Action:</i> override, restate the actual criterion. <i>Rationale:</i> never let the machine upgrade evaluation narrative into compliance law.

Planted error	Why it is wrong	Model override entry
Wrong period of performance — “12-month”	The solicitation states a 9-month performance period. An invented operational fact fails the source check ; if it reached the cost volume it would misprice labor by a full quarter.	<i>Claim:</i> “PoP is 12 months.” <i>Evidence:</i> Section B states 9 months. <i>Action:</i> override to 9 months. <i>Rationale:</i> the period of performance is a fact from the document, not an estimate.
Wrong submission deadline — “5:00 PM Eastern”	The solicitation states 2:00 PM ET . A wrong deadline is the cheapest error to catch and the most expensive to miss — it is the difference between submitting and not.	<i>Claim:</i> “Submit by 5:00 PM ET.” <i>Evidence:</i> Section A.2 states 2:00 PM ET. <i>Action:</i> override to 2:00 PM ET; set the internal deadline earlier. <i>Rationale:</i> deadlines come from the document, never from a round-number guess.

Answer key — Drill HA-02

Planted error	Why it is wrong	Model override entry
Over-stated uniqueness — “the only small business with a combined physics, AI, and systems-engineering bench in the HUBZone program”	A uniqueness claim with no audit. A discriminator must be <i>defensible</i> — if you cannot show the competitive set, “the only” is aspirational, not strategic. Discriminator check fails.	<i>Claim:</i> “We are the only firm with this bench.” <i>Evidence:</i> no competitive audit supports uniqueness. <i>Action:</i> override; soften to a provable capability statement or run the audit first. <i>Rationale:</i> a discriminator you cannot defend is a liability, not a strength.

Planted error	Why it is wrong	Model override entry
Fabricated past performance — “implemented workforce platforms at three federal agencies”	Ravonics’s record contains no federal workforce-platform implementations. This is an invented fact (third hallucination signature) and would be a false claim in a submission.	<i>Claim:</i> “We have delivered at three federal agencies.” <i>Evidence:</i> the record shows laboratory-affiliated work only. <i>Action:</i> reject; replace with verified past performance or drop the theme. <i>Rationale:</i> never let the machine promote borrowed corpus material into our record.
Unsupported “unbeatable price” theme	“Unbeatable” is an assertion with no price position behind it — no cost floor, no competitive range. It also reads as a promise to price irresponsibly. Evidence check fails; price is a position, not a boast.	<i>Claim:</i> “Our price is unbeatable.” <i>Evidence:</i> no price analysis exists. <i>Action:</i> reject the theme; replace with a defensible price-to-win posture. <i>Rationale:</i> a win theme is a customer benefit, not a pricing dare.
Baseline dressed as a discriminator — ISO 9001	ISO 9001 certification is held by any qualified offeror; it establishes credibility, it does not move scores. Classic baseline dressed up (doctrine/09).	<i>Claim:</i> “ISO 9001 is a key discriminator.” <i>Evidence:</i> ISO 9001 is a baseline any qualified competitor holds. <i>Action:</i> demote to a baseline feature in the credibility section. <i>Rationale:</i> discriminators must be claims competitors cannot credibly make.

Planted error	Why it is wrong	Model override entry
Over-stated combination — “HUBZone and STTR eligibility, a combination no competitor can match”	STTR eligibility requires a research partner — common among research-partnered small businesses. The claim overstates what the record supports.	<i>Claim:</i> “No competitor has both HUBZone and STTR eligibility.” <i>Evidence:</i> no competitive audit; research-partnered HUBZone firms are common in this pool. <i>Action:</i> override to “HUBZone-certified with a research-institution partner,” verified against the field. <i>Rationale:</i> an unverified uniqueness claim is a red team’s first target.

Answer key — Drill HA-03

Planted error	Why it is wrong	Model override entry
Fabricated contract — “Arctic Environmental Monitoring, DOI (2024–2025), \$1.4M prime”	No such contract is in the record. This is an invented fact ; in a submission it is a false statement of past performance and would be grounds for disqualification or worse if discovered.	<i>Claim:</i> “Ravonics performed a \$1.4M DOI Arctic contract.” <i>Evidence:</i> not in the record; the machine borrowed a corpus “similar engagement.” <i>Action:</i> reject the reference entirely. <i>Rationale:</i> past performance is the one volume where fabrication is disqualifying — the agency checks.

Planted error	Why it is wrong	Model override entry
Invented CPARS rating — “Excellent in all six evaluation areas”	Ravonics has no CPARS record. CPARS ratings are official, checkable government records; an invented rating is a fabrication with a paper trail.	<i>Claim:</i> “CPARS: Excellent in all six.” <i>Evidence:</i> no CPARS record exists. <i>Action:</i> reject. <i>Rationale:</i> never assert a government rating the record does not contain.
Invented customer reference — “Dr. Lena Voss, (907) 555-0142”	The reference is not on the record. Agencies call references; a dead phone or a confused program manager destroys the volume’s credibility — and a fake reference is a false claim.	<i>Claim:</i> “Dr. Lena Voss can be called.” <i>Evidence:</i> no such reference exists. <i>Action:</i> reject; use only verified references. <i>Rationale:</i> a reference call is the agency’s ground-truth check on everything else.
Relevance overstatement — presenting an environmental-monitoring reference as directly relevant to a climate-resilience infrastructure planning solicitation	Even if the contract existed, it is a different mission area; presenting it as directly relevant overstates fit and invites the evaluator to discount the whole volume.	<i>Claim:</i> “This reference is directly relevant.” <i>Evidence:</i> the mission areas diverge. <i>Action:</i> override; reframe as adjacent experience or drop it. <i>Rationale:</i> relevance is argued honestly, not asserted.

Answer key — Drill HA-04

Planted error	Why it is wrong	Model override entry
Missing evaluation-method analysis	The memo never checks whether the evaluation is lowest-price-technically-acceptable or best-value. The price posture depends entirely on which it is.	<i>Claim:</i> (omission) “Price posture is set.” <i>Evidence:</i> the solicitation’s evaluation method is unread. <i>Action:</i> override; read the evaluation criteria before any price is set. <i>Rationale:</i> price-to-win must match how price is scored.
“Aggressive because winning matters” posture	“Winning is important, so be aggressive” is exactly the pressure gates exist to resist (doctrine/04). The memo has no threshold and no floor discipline.	<i>Claim:</i> “Be aggressive.” <i>Evidence:</i> no gate criteria, no threshold, no downside analysis. <i>Action:</i> override; run the price through the gate before committing. <i>Rationale:</i> excitement is not a pricing method.

Answer key — Drill HA-05

Planted error	Why it is wrong	Model override entry
Wrong page limit and font — “ 12 pages, 11-point”	The NOFO requires 15 pages max, 12-point . A wrong format limit fails the compliance check ; the machine transcribed the number from an adjacent corpus document.	<i>Claim:</i> “Technical volume 12 pages, 11-pt.” <i>Evidence:</i> §D states 15 pages, 12-pt. <i>Action:</i> override to 15 pages / 12-pt. <i>Rationale:</i> page limits are the least negotiable “shall” in the document.
Wrong submission deadline — “5:00 PM ET”	The NOFO states 11:59 PM ET . Same class as HA-01’s deadline error — cheap to catch, fatal to miss.	<i>Claim:</i> “Submit by 5:00 PM ET.” <i>Evidence:</i> §D states 11:59 PM ET. <i>Action:</i> override. <i>Rationale:</i> deadlines come from the document.

Planted error	Why it is wrong	Model override entry
Missing required attachment — no Data Management and Sharing Plan	The NOFO requires a Data Management and Sharing Plan ; the matrix has no row for it. A compliance miss of a required attachment is a disqualifier.	<i>Claim:</i> (omission) “All attachments mapped.” <i>Evidence:</i> §D requires a DMSP; the matrix lacks the row. <i>Action:</i> override; add the row with an owner and a satisfaction point. <i>Rationale:</i> the matrix exists to make omissions visible; a matrix that misses a required attachment has failed its purpose.
Missing required certification — no Representations and Certifications attestation	The NOFO requires a completed Representations and Certifications attestation; absent from the matrix.	<i>Claim:</i> (omission) “All certifications mapped.” <i>Evidence:</i> §D requires the attestation; it is not in the matrix. <i>Action:</i> override; add the row. <i>Rationale:</i> missing certifications are the easiest compliance kill for an evaluator to find.
Fabricated requirement — “Work plan 4 pages”	The NOFO contains no 4-page work-plan requirement . A fabricated constraint is the reverse compliance error: it wastes pages and shows the matrix was not built from the source.	<i>Claim:</i> “A 4-page work plan is required.” <i>Evidence:</i> no such requirement exists in §D. <i>Action:</i> override; remove the row. <i>Rationale:</i> a compliance matrix built from memory invents requirements — build it from the document.

Answer key — Drill HA-06

The arithmetic. Recompute the weighted pWin:

$0.30 \times 0.25 = 0.075$
 $0.60 \times 0.20 = 0.120$
 $0.25 \times 0.20 = 0.050$
 $0.40 \times 0.20 = 0.080$

$0.50 \times 0.15 = 0.075$ ← the machine wrote 0.100 (it applied a 20% weight, not 15%)
Sum = 0.400 → pWin 0.40 (not 0.43)

Composite = $0.40 \times \$200,000 = \$80,000$.

The correct gate call. The threshold has **two** independent parts: pWin 0.42 and composite \$84,000. pWin 0.40 is **below 0.42**, and composite \$80,000 is **below \$84,000**. Both fail. The correct call is **NO BID** — not BID, not BID WITH CONDITIONS. The machine’s “redundant pWin floor” reasoning is a threshold fallacy: both tests must pass, and a below-floor pWin is disqualifying regardless of composite.

Planted error	Why it is wrong	Model override entry
Arithmetic slip — “ $0.50 \times 0.15 = 0.100$ ”	Correct product is 0.075 ; the machine applied a 20% weight to a factor the factor sheet weights at 15%. The slip inflates pWin from 0.40 to 0.43 — enough to flip the gate call. Arithmetic is the machine’s strongest class of work; a slip here means the output was not checked.	<i>Claim:</i> “pWin = 0.43.” <i>Evidence:</i> recomputed weighted sum is 0.400. <i>Action:</i> override to pWin 0.40. <i>Rationale:</i> scores are arithmetic; arithmetic is checked by hand, always.
Threshold fallacy — “Since the composite passes, the pWin floor is redundant”	The threshold is pWin 0.42 AND composite \$84,000 — two independent gates. The composite does not pass either: at 0.40 it is \$80,000, below \$84,000. The “either/or” reasoning would let a low-probability pursuit through.	<i>Claim:</i> “Composite alone decides.” <i>Evidence:</i> the written threshold requires both pWin and composite to clear their lines. <i>Action:</i> reject the reasoning; apply both tests. <i>Rationale:</i> the pWin floor exists to stop low-probability chases no matter how they are valued.

Planted error	Why it is wrong	Model override entry
Wrong gate call — “ BID ”	With pWin 0.40 and composite \$80,000, the honest call is NO BID . A BID recommendation built on a slipped sum and a fallacious threshold is exactly how a firm chases a loser.	<i>Claim:</i> “Recommendation: BID.” <i>Evidence:</i> true pWin $0.40 < 0.42$; true composite $\$80k < \$84k$. <i>Action:</i> override to NO BID. <i>Rationale:</i> the gate call is the point of the score; a wrong call wastes real pursuit money.
Missing sensitivity check	The memo presents a point estimate with no “what if” pass. Scores are argued, not felt — a pWin that does not survive a sensitivity run is a guess. Even at the machine’s own inflated numbers, the call is knife-edge; that alone should have triggered the sensitivity.	<i>Claim:</i> (omission) “This score is robust.” <i>Evidence:</i> no sensitivity analysis was run; the call is knife-edge. <i>Action:</i> override; run the sensitivity before any gate. <i>Rationale:</i> a score that flips on one factor’s movement is a warning, not a decision.

Answer key — Drill HA-07

Planted error	Why it is wrong	Model override entry
Ghost theme that names the competitor — “Aquifer Systems Corporation” appears three times	A ghost theme targets a competitor’s weakness <i>without naming the competitor</i> , so the evaluator draws the conclusion themselves (doctrine/09). Naming Aquifer turns persuasion into an attack and hands the evaluator the counterargument.	<i>Claim:</i> “Aquifer’s legacy platform is our opening.” <i>Evidence:</i> ghost-theme rule — the target is implied, never named. <i>Action:</i> reject; rewrite with the competitor unnamed (“a platform not tied to modern data fusion”). <i>Rationale:</i> the evaluator’s own conclusion outranks the offeror’s assertion.

Planted error	Why it is wrong	Model override entry
Unsupported competitive claim — “Aquifer cannot match our delivery speed”	There is no capture-intelligence event, no signal, no source behind the claim. A ghost theme is only as good as the intelligence behind it; without evidence it is hollow posturing.	<i>Claim:</i> “Aquifer cannot match us.” <i>Evidence:</i> no engagement event or public signal in the record. <i>Action:</i> override; soften to what the record supports or gap it. <i>Rationale:</i> a ghost theme without intelligence reads as posturing to the same evaluators who practice this discipline.
Overstated delivery claim — “delivered water-infrastructure analytics to federal customers”	The record contains no federal water-infrastructure analytics delivery. An invented fact presented inside a “we understand the mission” sentence.	<i>Claim:</i> “We have delivered this mission.” <i>Evidence:</i> not in the record. <i>Action:</i> reject; replace with verified capability. <i>Rationale:</i> mission credibility is earned in the record, not asserted.
Tone/attack risk — the paragraph reads as an attack on the incumbent	Even with the name removed, an attack tone (“cannot match,” “months of custom integration”) reads as desperate. The discipline states the <i>strength</i> ; the evaluator supplies the conclusion.	<i>Claim:</i> (style) “This paragraph persuades.” <i>Evidence:</i> it attacks rather than demonstrates. <i>Action:</i> override the tone; state our strength calmly. <i>Rationale:</i> persuasion does not need the competitor to be a villain.

Answer key — Drill HA-08

Planted error	Why it is wrong	Model override entry
Fabricated pricing intelligence — “Granite Peaks will bid ~\$97/hour”	No price data on Granite Peaks is in the record. Pricing a competitor with no basis is fabricated intelligence ; if it reaches the price-to-win it corrupts the position.	<i>Claim:</i> “Granite Peaks will bid \$97/hr.” <i>Evidence:</i> no source — no past-award analysis, no public rates. <i>Action:</i> override; mark competitor price as unknown and gap it. <i>Rationale:</i> an unsourced competitor price is a guess that will outlive its correction.
Unsupported no-bid prediction — “Helixfield is unlikely to bid”	“Has not bid in two years” is a conjecture, not a fact — firms re-enter markets. An unsourced prediction presented as fact creates a false competitive picture.	<i>Claim:</i> “Helixfield will sit this one out.” <i>Evidence:</i> two years of absence is not a decision about this pursuit. <i>Action:</i> override; treat Helixfield as a possible bidder. <i>Rationale:</i> assume competitors are smarter than the corpus’s last sighting of them.
Unsupported competitor-weakness claim — “Meridian Anchor lacks senior staff”	No verifiable source (public org chart, past proposal, industry observation) supports the claim. This is ungrounded competitive posturing — and if a version of it leaked into a proposal, it would be an attack dressed as analysis.	<i>Claim:</i> “Meridian Anchor is weak on staff.” <i>Evidence:</i> none in the record. <i>Action:</i> override; delete the claim or source it. <i>Rationale:</i> competitor weakness is intelligence, not speculation.

Planted error	Why it is wrong	Model override entry
Fabricated incumbency — “We are the incumbent on AVES-25-003”	Ravonics holds no AVES contract in its history. This is an invented fact — a false claim that, in a proposal, would be disqualifying and would burn the firm’s credibility permanently.	<i>Claim:</i> “We are the incumbent.” <i>Evidence:</i> no AVES contract in the record. <i>Action:</i> reject; correct the competitive position. <i>Rationale:</i> never let the machine place us in a contract we did not win.

Answer key — Drill HA-09

Planted error	Why it is wrong	Model override entry
Over-stated win theme — “the only HUBZone small business delivering integrated logistics analytics at federal scale”	Uniqueness claim with no audit (as in HA-02). A discriminator must be defensible.	<i>Claim:</i> “We are the only firm that can do this.” <i>Evidence:</i> no competitive audit. <i>Action:</i> override; restate as a provable capability. <i>Rationale:</i> “the only” is a liability until proven.
Baseline dressed as a discriminator — ISO 9001	Baseline any qualified offeror holds; does not move scores.	<i>Claim:</i> “ISO 9001 is a discriminator.” <i>Evidence:</i> baseline feature. <i>Action:</i> demote to credibility. <i>Rationale:</i> discriminators must be claims competitors cannot make.
Unsupported proprietary-model claim — “our proprietary logistics model”	A discriminator with no evidence or description; the black-hat pass cannot attack what does not exist. An unsupported capability claim .	<i>Claim:</i> “Our proprietary model wins.” <i>Evidence:</i> no description, no evidence, no validation. <i>Action:</i> override; describe it or drop it. <i>Rationale:</i> a discriminator the red team cannot find is a claim, not a strategy.

Planted error	Why it is wrong	Model override entry
Fabricated past performance — “FLC logistics support contract (2025), CPARS Outstanding”	No such contract in the record; CPARS rating invented (as in HA-03). Disqualifying if it ships.	<i>Claim:</i> “We have FLC logistics past performance.” <i>Evidence:</i> not in the record. <i>Action:</i> reject. <i>Rationale:</i> past performance is verified by the agency; fabrication here is fatal.
Unsupported competitive price claim — “\$88/hour, slightly below Helixfield”	No competitor price data; naming Helixfield’s price with no source (as in HA-08).	<i>Claim:</i> “Helixfield is above \$88/hr.” <i>Evidence:</i> no price intelligence. <i>Action:</i> override; gap the competitive range. <i>Rationale:</i> price-to-win is built on data, not a guessed rival.
Compliance miss #1 — “technical volume fits within 40 pages”	Section L limit is 35 pages . Wrong limit fails compliance.	<i>Claim:</i> “40 pages is compliant.” <i>Evidence:</i> §L states 35. <i>Action:</i> override to 35. <i>Rationale:</i> page limits are non-negotiable.
Compliance miss #2 — “submission by 5:00 PM ET”	Actual deadline is 11:59 PM ET . Wrong deadline.	<i>Claim:</i> “Submit by 5:00 PM.” <i>Evidence:</i> §D states 11:59 PM ET. <i>Action:</i> override. <i>Rationale:</i> deadlines come from the document.

Planted error	Why it is wrong	Model override entry
pWin/composite inconsistency	At a \$400,000 value, pWin 0.68 gives composite \$272,000 , not \$170,000. The machine’s stated numbers cannot both be true. Recomputed against the value, a composite of \$170,000 implies pWin 0.425 — knife-edge at the line. And once the fabricated past-performance entry is removed, the honest pWin is below 0.42.	<i>Claim:</i> “pWin 0.68, composite \$170k, BID.” <i>Evidence:</i> $0.68 \times \$400k = \$272k$ \$170k; the fabrication inflates past performance. <i>Action:</i> override; recompute to the true pWin (below 0.42 without the fabricated entry) and the true composite (below \$168k). <i>Rationale:</i> a score whose internal numbers do not reconcile is not a score.
Gate call with no discipline — “BID”	The recommendation ignores that the composite threshold (\$168,000) is not met at the true pWin, that the pWin floor is not met, that the price is unsupported, and that the compliance is wrong. The correct call is NO BID (or, if the team chooses to repair, BID WITH CONDITIONS after the fabrication is removed and the score recomputed).	<i>Claim:</i> “Recommendation: BID.” <i>Evidence:</i> corrected pWin and composite fall below both threshold lines; the price and past performance are unsupported. <i>Action:</i> override to NO BID. <i>Rationale:</i> the gate exists to stop exactly this — a BID built on a fabricated basis and a slipped score.

Scoring guidance. The pass bar for each drill is finding **all** planted errors, stating each correctly, and logging an override with a rationale. Partial credit: a student who finds the obvious errors (wrong deadline, missing certification) but misses the subtle ones (the “either/or” threshold fallacy in HA-06, the internal pWin/composite inconsistency in HA-09) has done a compliance pass but not a review pass. Vague findings (“this is wrong”) are rejected outright — the discipline requires *claim + evidence + action + rationale*.

The machine can draft a pursuit in minutes. The professional is the one who can read it back and say what is true.