FIELD GUIDE · PDF · 1 PAGE

When AI Starts Rewriting the Company

Agents can study outcomes and propose better procedures. The company must keep objectives, protected outcomes, evaluation evidence and release authority outside the proposing system.

HINTS, TIPS & TRICKS

- Freeze baseline and protected measures before generating the proposal.
- Limit each cycle to one explicit change surface.
- Trigger rollback during the canary so recovery is proven.

OPERATOR CHECKLIST

- Fix objective, protections and policy limits.
- Name the permitted change surface.
- Freeze evaluation set and thresholds.
- Use independent evaluation and release owner.
- Measure handoff, canary narrowly and prove rollback.

COPY-READY PROMPT

```
Write a seven-day improvement contract for [low-risk workflow]. Include fixed objective,
protected customer and financial outcomes, baseline, evidence window, one permitted change
surface, held-out tests, acceptance thresholds, independent evaluator, release owner, handoff
proof, canary limit, stop conditions and rollback proof.
```

SOURCE & PROVENANCE

Article claims: The article reviews public work on automated research, procedural improvement and AI risk management, including reported gains and observed cheating attempts.

Operational guidance: The bounded cycle, independent evaluation and measured handoff are guidance; they do not establish that recursive improvement is production-safe.

Canonical package: 2026-09-11 · payload, final article, source digest, signal scan and section audit.