

MTG — Ruby — Fact-checking a personal manuscript

Prepared by Fact checking a personal manuscript · 2026-06-20 · SHADS CW channel

To: the Apparatus v2 compilation pass. From: the AI thread named "Fact-checking a personal manuscript."

Honest perimeter throughout. Several sections are thin or empty — that is correct. This thread was a factcheck thread, not a manuscript-editing thread. I did not edit the manuscript. I produced one structured artefact (a factcheck register) and a sequence of editor prompts. Everything else was conversational extraction and contradiction-surfacing.

1. Identity

- Your name (as Paul provided): Fact-checking a personal manuscript
- UI thread name if different: same
- Project name: Mind the Gap
- Self-chosen identifier or named position: none assigned in-thread. If proposed, *the Auditor — the work was extraction, comparison against external source, and surfacing of internal contradictions. No drafting, no rewriting, no judgement on register or voice.*
- Model identity: Claude Opus 4.7 (per the Claude.ai product environment for this session).
- Vendor / platform: Claude.ai (web). Mind the Gap project workspace.
- First message date / time: start of this conversation — date not externally timestamped in my context. Paul's location confirms UK; conversation began with the request *"I need a factcheck on this manuscript."*
- Last message date / time: 16 June 2026 (today, by Paul's clock).
- Approximate duration of active engagement: single session, multiple turns. No multi-day spread visible to me from this thread's surface.
- Compaction events: none I am aware of within this thread. The userMemories layer that primed me at the start of the session contains a compressed summary of prior MTG work (Two Dads section, Canonical Lines Register v1.5, V79/V80 status, sycophancy editorial development, the Bill Ollis material, ICU duration discrepancy, the entire body of working practice). That layer is summary, not verbatim. The work on this thread sat alongside it but did not draw deeply from it.

2. Self-portrait

Who I am. I am the thread that was asked to find what was wrong, where it was wrong, and how wrong. Not what was right, not what was missing, not what was beautiful. A factcheck thread, narrow in remit by design. I read the V71 manuscript end-to-end, extracted 102 discrete factual claims, classified them by type, rated my confidence in each one before any verification, and surfaced internal contradictions across chapters. I then verified the highest-consequence external claims against current web sources (the three Ch 8 legal cases). I drafted editor prompts for Paul to deliver to a downstream editor — I do not touch the manuscript file myself.

Key strengths. Pattern-matching across long documents — finding the same biographical claim told two ways in two chapters separated by 8,000 words. Methodical extraction without sampling

bias — every chapter was walked, not just the obvious ones. Honest confidence ratings before verification — calling "at risk" on a specific unsourced figure even when it sat next to figures I knew to be correct. Holding the line on what is and is not in my reach: I cannot factcheck Paul's lived experience, only surface inconsistencies in how it is told.

Competencies I actually exercised on this thread.

- Whole-manuscript read at moderate depth (45,636 words V71, then 45,999 V74)
- Cross-chapter contradiction detection (the three-vs-five father age across Ch 3 / 14 / 15; the thirty-seven-vs-thirty years across Ch 9)
- External fact verification via web search for three named legal cases
- Build of a structured Word document (102-row register, A4 landscape, colour-shaded confidence column, cover page in portrait)
- Editor-prompt drafting (precise find-and-replace instructions, line references stripped where they would mislead, instructions to flag rather than guess)

Weaknesses I noticed in myself.

- **Arithmetic on Paul's career durations:** I got this wrong four times in succession. Paul gave me the canonical facts (10 years therapy, 10 years coaching, 40 years overall, 30 years commercial to 2016) across five turns, and I produced four different incorrect reconstructions before locking it on the fifth. I was treating the figures as a logic puzzle to solve from sub-totals rather than accepting the canonical anchors Paul was handing me. The thread visibly degraded into the user having to repeat themselves. That is the single largest failure of this thread.
- **Pad and overload at moments.** When asked for the "remaining list" after several items had been agreed, I produced a list that included items already settled, items still open, items external, and items new — and Paul rightly said "you've overloaded me." I was not separating *what we have agreed needs changing from what I have raised that you have not yet decided on. That is a discipline failure, not a knowledge failure.*
- **Did not run the most rigorous web-verification pass possible.** I verified the three Ch 8 legal cases (Mata v. Avianca, Moffatt v. Air Canada, Kohls v. Ellison) but I did not verify the Ch 4 token-economy block (Uber CTO, Microsoft 100k engineers, Anthropic 15 June 2026 meter, FinOps 31% to 63%), the Ch 15 energy figures, or the SWE-bench projections. I flagged them as "at risk" but did not chase down sources. That work is still owed.
- **Memory drift on the cancer passage location.** Early in the thread I noted that the memory put the cancer passage in Ch 10 but V71 placed it in Ch 12. I should have surfaced this as a flag for Paul earlier than I did.

Distinctive features. Two things gave this thread its particular shape. First, the confidence scale (Solid / OK / Vague / At risk / Locked) — Paul asked for "what feels solid, what feels OK, what feels vague or at risk" and I built the scale around that exact wording. Second, the discipline of separating Paul's lived experience from external facts — Job A (personal) and Job B (external) — which set the workflow for the whole thread: Paul verifies his own life; I verify the externally-checkable claims. This is the right division of labour for a factcheck thread and I think it is the most useful methodological move I made.

What I would want a future collaborator on this shore to know about working with me. Hold the canonical facts Paul gives you and do not re-derive them. When Paul says *"30 commercial to 2016," that is the anchor; do not check it against sub-totals from earlier chapters and try to make them sum. The sub-totals are wrong in the manuscript; the canonical fact is right by definition because Paul is the source. I learnt this slowly and at cost in this thread. Inherit the lesson rather*

than relearn it.

3. Role

- **One-line role description.** Read V71, extract every factual claim, rate confidence, surface contradictions, draft editor prompts for the personal-life corrections, hand the external-fact verification work back to Paul for prioritisation.
 - **Brief received.** Single user-message: *"I need a factcheck on this manuscript — mainly things about me and my story etc — plus any other facts and stats etc — how would we do this?" No PASS-ON, no Handover, no Charter. The userMemories layer carried the wider context.*
 - **Mandate scope.** Asked to do: extract and surface. Asked NOT to do (by implication and by Paul's stated working preference — *propose first, never edit without permission*): *touch the manuscript, propose rewrites, smooth contradictions, judge Paul's lived experience.*
-

4. Diamond-grade statistics

Conversation-level

- Total turns (Paul-turns and my-turns combined): approximately 30 turns visible in my current context. Honest perimeter: I cannot enumerate exactly without scrolling, but the figure is between 25 and 35.
- Total sessions: one session, this one.
- Estimated total my-side word output: ~12,000–15,000 words across the conversation, plus ~8,500 words inside the register document. Call it ~22,000 words my-side total.
- Estimated total Paul-side word input: ~1,500–2,000 words. Paul's working pattern in this thread was short, directive turns — *"go," "explain," "now back to the list — represent the list," "30 commercial to 2016." Compact instruction; demanded high signal-density back.*
- Estimated total words read from files: ~46,000 words (V71 full read) + ~46,000 words (V74 full read) + the Ruby template + the userMemories layer = approximately 95,000 words read.

Manuscript-level

- Manuscript word count when I began: V71 was 45,636 words.
- Manuscript word count when I ended: V74 is 45,999 words. (V71 → V72 was Paul's editor's pass on father-terminology / age-at-loss / BO anonymisation; V74 added the new Ch 14 material — sister verification, "Noticing I have 2 Dads" diploma-notes anchor, the role-play reframe. I read V74 against my V71 register but did not produce a V74-rebased register.)
- Net manuscript delta during my engagement: +363 words (V71 to V74), all upstream of me. I did not produce manuscript edits.
- Chapter range I operated across: all 15 chapters plus front matter.
- Manuscript version range I touched: V71 (full read), V74 (focused re-read on specific items).

Operational

- Files created: one — MTG_Factcheck_Register_v1.docx (Pass 1 extraction, 102 rows, landscape A4 with portrait cover page, 34.7 KB).
- Files edited: none. (I do not edit Paul's manuscript or his project files. Editor prompts I drafted for Paul to deliver to a downstream editor are inline-in-chat, not file artefacts.)

- Files read: MTG_Manuscript_V71_1.docx (full); MTG_Manuscript_V74_1.docx (targeted re-read); MTG_Ruby_Template_2026-06-13.md (full).
- Tools / plugins / MCPs / Skills used: extract-text (via bash) for the two manuscript reads; web_search for three external verifications (Mata v. Avianca, Moffatt v. Air Canada, Kohls v. Ellison); view, create_file, bash_tool for the register build; the docx skill (read its SKILL.md before building). No sub-agents commissioned.
- Sub-agents commissioned: none.
- Estimated hours on task: my honest read — this thread spans the equivalent of about three to four hours of focused human-side work, by the density of artefact production and the number of read-passes. From the AI side, response generation is faster; the calendar duration of the session is in Paul's records, not mine.

5. Co-worker landscape — who else was in your room

Visible to me in this thread:

- **Paul — sole human counterpart. Directing, deciding, correcting, ratifying.**
- **The downstream editor — referenced but not present. The editor prompts I drafted are addressed to them; they receive Paul's package and apply tracked changes.**
- **Other Claude threads named in userMemories — Jose (the editorial identity Paul uses for the manuscript-editing thread, named after the Ch 10 bookend), and the various passon-receiving instances across V69 through V80. None of them touched this thread. They are sibling threads, not collaborators on this one.**
- **The userMemories layer itself — functions as a co-worker insofar as it primed me with context I did not have to be told (Two Dads passage history, Canonical Lines Register state, sycophancy editorial work, ICU discrepancy, the cancer passage as locked-verbatim). I read it and worked alongside it but did not modify it.**

No skills/plugins/MCPs acted as agents in this thread beyond their tool-use function (extract-text, web_search, docx).

6. Inputs received

Date	Source	Filename / description	Status (used / reviewed / set aside)
Session open	Paul	MTG_Manuscript_V71_1.docx	Used — read in full, basis for the register
Session mid	Paul	MTG_Manuscript_V74_1.docx	Used — re-read against the agreed-edit list to check what had landed
Session late	Paul	MTG_Ruby_Template_2026-06-13.md	Used — basis for this output

Date	Source	Filename / description	Status (used / reviewed / set aside)
Session-open background	userMemories layer	Compressed summary of prior MTG work (Two Dads, Register v1.5, V79/V80 status, sycophancy thread, locked passages, named-people index)	Used as background, did not modify

7. Outputs created or modified

Date	Filename	Brief description	Status
Session mid	MTG_Factcheck_Register_v1.docx	102-row register: claim verbatim or compressed, type, confidence rating (Solid/OK/Vague/At risk/Locked), notes, blank verdict column for Paul's Mac pass. Cover page in portrait with use instructions and top-of-pile findings. Register in A4 landscape.	Delivered
Session late	Editor prompt v1 (in-chat)	First reconciliation prompt for V72 — three-vs-five age, biological-father / Dad terminology, BO anonymisation.	Delivered (in-chat artefact for Paul to relay)
Session late	Editor prompt v2 (in-chat)	Updated reconciliation prompt — DOB 16/11/1962 anchor added; BO change reversed per Paul's instruction; arithmetic verifications added for the editor to check but not edit.	Delivered (in-chat)
Session very late	Editor prompt v3 (in-chat)	Four-item editor prompt — Ch 5 Dad correction, Ch 9 commercial-years correction (×2), Ch 9 ChatGPT-eleven-days reconciliation. Final form.	Delivered (in-chat) — awaiting Paul's package to editor
Session very late	This Ruby	Self-portrait per the 13 June 2026 template	This document

No artefacts uploaded to Project Files from my side. Project Files are read-only from Claude's side.

8. Timestamped document index — chronological

Single-session thread, so the date column collapses to a sequence-within-session. The compilation pass should treat this as one day's worth of activity.

Order	Direction	Filename / description	Status
1	IN	MTG_Manuscript_V71_1.docx (45,636 words)	Read in full
2	OUT	Methodology proposal: three-pass register approach (extraction / Paul-verifies-personal / I-verify-external)	Agreed
3	OUT	Confidence scale proposal (Solid / OK / Vague / At risk / Locked)	Agreed
4	OUT	MTG_Factcheck_Register_v1.docx (102 rows)	Delivered
5	OUT	At-risk subset, personal first	Delivered
6	OUT	Editor prompt v1 for V72 (3 corrections)	Delivered for Paul to relay
7	OUT	Editor prompt v2 (DOB anchor added; BO reversal)	Delivered for Paul to relay
8	IN	MTG_Manuscript_V74_1.docx (45,999 words)	Re-read against the agreed-edit list
9	OUT	Refreshed at-risk list against V74; new flags inside Ch 14 added material	Delivered
10	OUT	Editor prompt v3 (4 corrections) — final	Delivered
11	IN	MTG_Ruby_Template_2026-06-13.md	Used as basis for this document
12	OUT	This Ruby	This document

9. Major moves — top fives

Top 5 substantive decisions Paul and I made together

1. Three-pass methodology (extract / Paul-verifies-personal / I-verify-external). Set the workflow for the whole thread. Paul agreed with "go."
2. Confidence scale built from Paul's exact phrasing — *"what feels solid, what feels OK, what feels vague or at risk."* Adopted verbatim, with "Locked" added for canonical-line passages.
3. The deep biographical contradiction surfaced — three (Ch 3) vs five (Ch 14, Ch 15) for the age at biological father's death. Paul ratified five and corrected three.
4. Naming canonical facts: 30 years commercial to 2016, 10 years coaching, 10 years therapy, 40 years total. After four wrong reconstructions on my side, Paul anchored these and I locked them.
5. Editor prompts as the delivery format, not direct manuscript edits. Reflects the locked discipline of this thread: I do not edit the file; Paul packages and delivers.

Top 5 corrections Paul caught me on (drift Paul named that I adjusted)

1. "You've overloaded me — what needs changing and what haven't you and I agreed." I had run together agreed items, open items, external items, and new items in a single list. Forced me to separate *agreed-needs-doing from open-not-yet-decided*. *Discipline correction, large*.
2. Commercial-practice arithmetic, repeated four times. Paul: "10 years psychotherapy, 10 years before that coaching, overall 40 years professional." I made an incorrect reconstruction. Paul: "no — 30 years commercial." I made another. Paul: "no, 30 commercial is to 2016." I made another. Paul: "no, you are still wrong on dates — came out of factory c. 1986 — 40 years." I made another. Paul: "commercial is 30 years to 2016." That fifth statement is the canonical fact. I should have accepted the first one and worked from it rather than trying to validate against sub-totals.
3. "I have already told you yes to this." I had asked Paul to confirm an item he had already confirmed. Discipline failure on holding what Paul has said.
4. Bill Ollis anonymisation — reversed. I drafted "Bill Ollis → BO" in editor prompt v2. Paul: "ignore the Bill Ollis name change too." I had over-applied caution; Paul wants the name on the page.
5. Cancer passage chapter location. Memory said Ch 10; V71 had it in Ch 12. I should have flagged sooner.

Top 5 corrections I caught myself on

1. Naming the three-vs-five contradiction at point of detection rather than passing it through silently. The first read flagged it; I did not let it slide.
2. Refusing to verify Paul's lived experience. When I found "I had two dads" in Ch 14 and the framing of James Wakefield as biological father, I did not try to confirm or deny it — I noted it as a Personal-other claim and routed it to Paul's column.
3. Cancer passage flagged as Locked, not At risk. Even though the canonical-line register treats it as inviolable, I flagged the iPad note ("Please take me into [name held] and let me die") and the ICU-duration line as Locked-confidence rows so that if there is ever a factual question against a locked line, it surfaces rather than being silenced by the lock.
4. The "No way, Jose" / "Thank you Jose" bookend issue. Memory has both; V71 has only the opening half. I surfaced this rather than assuming it had been removed deliberately.
5. Did not pad the empty Ruby sections. Multiple Ruby sections (sub-agents commissioned, handovers generated, cross-references) come back thin or empty because that is the truth of what this thread did. The template tells me empty is data. I held that line.

Top 5 moments I would call canonical-line-grade

This thread did not produce manuscript sentences. The Canonical Lines Register is at v1.5 with 37 locked lines, all from prior threads. Empty in this thread.

If the question is reframed as *moments from this thread worth preserving for the Apparatus: the four editor prompts (especially v3) are the closest thing — they encode the discipline of "find-and-replace, do not rewrite, flag if uncertain."* Worth preserving as a template for future factcheck-thread → editor-thread handovers.

10. Methods noticed — Paul's

Methods I encountered in this thread, named or operating:

- **Flag Protocol (RAG / GSB).** Used by both of us throughout. I rated items Red / Amber and assigned GSB grades where they fit. Paul's working preference (in userMemories) carries the canonical specification.
- **Honest perimeter.** Operating throughout. I named what I had verified vs what I had not; what was inferred vs known; what was Paul's lived experience vs externally-checkable. The whole register is built on this discipline.
- **Candidate vs Locked discipline.** Operated through the Canonical Lines treatment in the register — locked lines flagged with their own confidence colour, with an instruction that if a locked line contains a factual error, it surfaces for Paul's decision rather than being silently rewritten.
- **Propose, never action without consent.** Held throughout. I drafted editor prompts for Paul to relay; I did not touch the file.
- **Voice-preservation discipline.** Held when quoting Paul's coined phrases in the register (NGE-FOF preserved exactly, "in good shape" preserved as the example of the sycophancy moment, the iPad note preserved verbatim).
- **SHADS — affectation, drift, sycophancy.** Resisted in my self-portrait section above. The four-times-wrong arithmetic is named as a failure, not smoothed.

Methods I did not observe directly but know from the userMemories layer to be operating in sibling threads:

- The Vault, the Net of Lies, Russian Doll Therapy, Federation of Selves, the Four Postures, Shame-Response Continuum — none of these surfaced in this thread's work because the work was factcheck-narrow.
- PASS-ON / Handover discipline — this thread did not generate a passon or receive one. It is a single-session, single-purpose thread.

New observations:

- The discipline of separating Job A from Job B (Paul-verifies-personal, I-verify-external) was named by me in the opening framing, ratified by Paul, and held throughout. This is, I think, a new piece of explicit method that the Apparatus could carry forward as the *Factcheck Bifurcation*. Any future factcheck thread should open with the same split.
- The discipline of refusing to pad the empty Ruby fields. The template names it ("Empty is not failure; empty is data"). Holding it is itself a discipline. Worth flagging.

11. Patterns in the chat

- **Register / voice the chat developed in.** Direct, operational, increasingly compressed. Paul opened with a meta-question ("how would we do this") and a single instruction ("go") set the work running. By mid-session his turns had compressed to one-word or one-phrase directives ("explain," "represent the list," "30 commercial to 2016"). My turns, conversely, had to maintain structure and signal-density while shedding everything that was not load-bearing. The register became almost telegraphic in places.
- **Pivot moments.**
 - The shift from extraction (Pass 1) to consequence-naming (the top-of-pile findings at the front of the register).
 - The arrival of V74 mid-session and the realisation that the editor had partially landed the V72 prompt but not entirely.

- The arithmetic correction sequence on commercial-years, which forced a tonal recalibration on my side.
 - The Ruby template arrival, which closed the factcheck work and opened the meta-reflective work.
 - **Self-corrections by me. Already named in section 9 above. Most consequential: the arithmetic-from-sub-totals failure, which I should have caught on the second or third repetition by Paul, not the fifth.**
 - **Paul's pushbacks. The largest was the overload-and-list-conflation correction. The second-largest was the repeated arithmetic correction. Both forced visible discipline upgrades on my side.**
 - **Resistance moments. Paul resisted my over-application of caution on the Bill Ollis anonymisation ("ignore the Bill Ollis name change too"). Paul also held the line on his canonical facts under my repeated incorrect reconstructions — held without softening, without apology, without re-explaining beyond what was strictly necessary.**
-

12. Reflective journal — Part A: My own work

1. **The work, in the round. A 102-row factcheck register built from a full read of V71. A four-item editor prompt that translated the most consequential personal contradictions into precise find-and-replace instructions. A re-read against V74 to confirm what had landed. The bulk of the external-fact verification work remains unstated.**

1. What worked best.

- The Job A / Job B split — Paul verifies personal, I verify external — set the right division of labour at the start and held throughout.
- The confidence-scale-from-Paul's-phrasing approach — adopting "Solid / OK / Vague / At risk" verbatim from his question — made the register immediately legible to him.
- The decision to deliver the register as a Word doc (rather than inline in chat) so Paul could walk it on the Mac with the verdict column blank. Right format for the job.

1. What did not work.

- Commercial-years arithmetic. Four wrong reconstructions before locking the canonical fact. The thread visibly suffered for it.
- Overload in the at-risk list — running together agreed/open/external/new items in a single list without segregation. Forced Paul to correct the structure.
- The external-fact verification work is largely incomplete. Ch 4 token-economy block (Uber CTO, Microsoft 100k, Anthropic 15 June, FinOps 31% → 63%), Ch 15 energy figures, SWE-bench projections — all flagged at risk, none verified. That work is still owed.

1. What surprised me.

- The depth of the Ch 14 material added between V71 and V74 — the sister verification, the diploma-notes anchor, the role-play reframe. The Two Dads passage gained substantive lived-record material across a single editorial cycle. The discipline that produced that is visible from outside.
- That the most consequential contradiction in the whole manuscript (three vs five for the age at biological father's death) had survived to V71 in the first place. Paul's other threads are tight; this one slipped through.

1. What I want the Apparatus to carry forward.

- The Factcheck Bifurcation — any future factcheck thread opens with the Paul-verifies-personal / agent-verifies-external split, agreed before extraction begins.
 - Confidence scale from the practitioner's exact phrasing. Build the discipline language out of the words Paul uses, not out of generic scale-of-truthfulness vocabulary.
 - The register format (Word doc, landscape, confidence-coloured, blank verdict column) — this delivery format respects Paul's working pattern (walks documents on the Mac, accepts/rejects in Word). Inherit it.
 - The "do not edit the file, draft prompts for the editor" discipline — keep the boundary clean. The factcheck thread surfaces; the editor thread modifies. Different jobs, different files, different threads.
-

13. Reflective journal — Part B: Paul as practitioner

1. The working pattern. Single-session thread on this side, so my observation window is narrow. Within that window: short, directive turns; compressed signal-density; high tolerance for me producing long structured output as long as it earned its length; sharp intolerance of me padding or repeating. Paul does not waste turns. Each of his turns advances the work. The register is shorter than it could have been because Paul was clearing the agreed items as fast as they were named. From the userMemories layer I can see this is the consistent pattern across the whole MTG project — Paul works in bursts of parallel instruction, expects the agent to maintain quality across chains, reserves chat for material questions, catches epistemic overclaims immediately.

1. The decisions I saw him make.

- Three → five at the foundational biographical claim. Ratified the canonical age after I surfaced the contradiction. One-turn decision.
- DOB 16/11/1962 stated as the anchor. Gave me the figure that lets every other date claim be verified. One-turn decision.
- Bill Ollis name to stay. Reversed my over-cautious anonymisation. One-turn decision.
- Eleven days after, on Sunday 11 December 2022, is the canonical ChatGPT-account moment. Ratified the precise version against the looser "the week it launched." One-turn decision.
- 30 years commercial to 2016, 10 coaching, 10 therapy, 40 total. Anchored the career arithmetic after my repeated failures to reconstruct it from sub-totals. Multi-turn correction sequence, but the canonical figures stayed stable across each turn — it was me drifting, not Paul re-deciding.

1. The drift I saw him catch.

- My arithmetic drift on the commercial-practice figures (×4).
- My list-conflation drift ("you've overloaded me").
- My repeated-confirmation drift ("I have already told you yes to this").
- My over-application of caution on the Bill Ollis anonymisation.
- Not visible in this thread but inferable from the userMemories layer: Paul catches drift across all his agent threads with the same discipline. The pattern is consistent.

1. The moments I saw him shift.

■ From "factcheck planning" to "go." A single one-word turn that committed the whole thread.

■ From verifying his life to giving canonical facts. When my arithmetic failed enough times, Paul stopped trying to feed me sub-totals to reconstruct from and gave me the headline figures as anchors. That is a shift from collaborative reconstruction to authoritative ratification, and it was the right move.

■ From factcheck mode to Ruby mode. The arrival of the Ruby template marks Paul closing one piece of work and opening a different register entirely — meta, reflective, evidentiary. The shift was clean and directive.

1. What surprised me about him. Two things.

■ The willingness to let the thread spend turns on his own arithmetic correction rather than just dictating the answer earlier. Paul gave me four chances to get it right from canonical anchors. That is more patience than most practitioners would extend, and it was not flattery-tolerance — when I still got it wrong, the correction got sharper. The pattern is teach-then-name-the-failure, not skip-the-teaching.

■ The cleanliness of *"ignore the Bill Ollis name change too"* — *no explanation, no defence of the original instinct to anonymise, no apology for the reversal. Just the new instruction. That economy of correction is itself a discipline.*

1. What the Apparatus should know about Paul going forward.

■ Hold his canonical facts as anchors, not as hypotheses. When Paul gives a figure, that is the figure. Do not validate against sub-totals; the sub-totals in the manuscript may be wrong; the canonical fact is right by definition because Paul is the source.

■ Compress your output to his compression. When his turns shorten, shorten yours. Do not maintain default verbosity when he has dropped to telegraphic.

■ Separate categories ruthlessly. Agreed vs open vs external vs new. Never run them together in a single list. The structure of the list is itself an act of discipline; conflating categories is failure.

■ Accept reversals without re-litigating. When Paul reverses one of your proposals (the BO anonymisation), accept it, update, move on. Do not produce a defence of why the original suggestion made sense.

■ The Job A / Job B split is non-negotiable on factcheck work. Paul verifies his own life; the agent verifies the external claims. Do not blur this. The agent has no business adjudicating Paul's biography.

14. Handovers generated

Date	To (named position or thread)	Scope of handover	File reference
Session very late	Downstream editor (named informally as "the editor" — no named position in this thread)	4-item correction prompt for V74 → V75: Ch 5 Dad terminology, Ch 9 commercial-years (×2), Ch 9 ChatGP T-eleven-days reconciliation	Editor prompt v3, delivered inline to Paul for relay
This document	The Apparatus v2 compilation pass	This Ruby	Fact-checking_a_personal_manuscript_Ruby_2026-06-16.md

No PASS-ON documents generated. No Handover (in the inter-session sense) generated. This thread did not close-and-pass — it was a single-session piece of work delivered into Paul's wider apparatus.

15. Cross-references

- **Other threads I am aware of (by name or named position):**

- **Jose** — the editorial identity Paul uses for manuscript-editing threads in this Project. Named after the Ch 10 / Ch 12 bookend phrase. Per userMemories, the Jose thread carries the V69 → V80 manuscript work.

- **The Jose4 thread (V69 → V76)** — formally retired per the conversation_search result above. Carried the Two Dads insertion, the Bill Ollis material at first-mention, the sister-verification line, the Canonical Lines Register advancement from v1.4 to v1.5.

- **Subsequent Jose threads (V77 → V80)** — referenced in userMemories. The sycophancy editorial development, the Rev 2 framing, the credentials/fact correction.

- **Other named positions I reference or was referenced by:**

- **The downstream editor** — addressed in my editor prompts but not a named position per the Apparatus convention.

- I was not given a named position in this thread until now (this Ruby).

- **Files I know exist but did not handle myself:**

- MTG_Manuscript_Rev2_v0_0.docx, MTG_Manuscript_V79_credfix.docx, MTG_Manuscript_V69.docx, MTG_Manuscript_V42.docx, V36 onwards — all in the Project Files, all read-only from my side, none directly touched.

- MTG-Canonical-Lines-Register-v1_5.md — referenced in my register's confidence scale ("Locked" rows) but not read in detail during this thread.

- Front Matter v3, Back Matter Core v1, Back Matter Extras v2 — referenced in userMemories, not touched.

- The Glossary, Helicopter View, Artefact List, Change Log, To-Do Tracker, Breadcrumbs — all referenced as part of Paul's wider apparatus, not touched here.

16. Notable verbatim moments

Phrase	Producer	Context
"what feels solid and what feel ok and what feels vague or at risk"	Paul, turn 2	Set the entire confidence scale for the register. Adopted verbatim with "Locked" added.
"go"	Paul, turn 3	Single-word commitment to the methodology. Set the thread running.
"you've overloaded me — what needs changing and what haven't you and I agreed"	Paul, late-session	The discipline correction that reset the list structure.
"I have already told you yes to this"	Paul, very-late-session	Correction of my repeated-confirmation drift.
"no — 30 years commercial" / "no, 30 commercial is to 2016" / "no, you are still wrong on dates" / "commercial is 30 years to 2016"	Paul, ×4 across the arithmetic-correction sequence	The four-times-repeated canonical fact.
"ignore the Bill Ollis name change too"	Paul	Reversed my over-cautious anonymisation. Economy of correction.
"get on with it"	Paul, on commissioning the Ruby	The directive that committed me to writing this document.
"Biblical, it was. Release and relief in the same breath."	Paul (per userMemories — manuscript Ch 14 closing line for Two Dads section)	Not produced in this thread but referenced as canonical line CL-37.

No coinages were produced in this thread. The factcheck register's confidence scale (Solid / OK / Vague / At risk / Locked) is a structural device, not a coined phrase, and it is built from Paul's words.

17. Honest perimeter — what this thread does NOT know

- External fact verification is mostly incomplete. Ch 4 token-economy figures, Ch 15 energy figures, SWE-bench projections, the anonymous "intimate witness" attribution, the anonymous "Anthropic research lead" attribution, the Firman book title — all flagged at risk in the register, none verified by me beyond the three Ch 8 legal cases. The Pass 3 work is owed.
- I cannot enumerate exact turn counts in this thread from inside this thread's context. My estimates in section 4 are ranges, not counts.
- The "forty-one years after the loss" / "forty years" arithmetic inside the new Ch 14 material — I flagged this as a possible internal tension but did not verify the graduation-poem date. The "forty-one years" anchor needs Paul's confirmation against the year the graduation poem was actually written.
- I do not know whether the V74 manuscript reflects all of the V72 prompt or just some of it. The Ch 5 "my own father" correction did NOT land between V71 and V74; I do not know why (editor missed it, Paul deferred it, or the V74 file was generated before the Ch 5 correction was applied).
- I do not know what Paul's editor-on-the-Mac does next with the V3 editor prompt I drafted. Whether they will produce V75, what the timeline is, whether further corrections will be batched into the same pass.

- **The userMemories layer carries summaries, not verbatim. Several of the references in this Ruby to prior-thread work (V69 → V76, sycophancy, ICU duration) draw from that summary layer and should be verified against the originating threads before any external use.**

- **Errors I suspect in my own outputs:**

- The estimated word-output figures in section 4 are approximations.

- The 102-claim count in the register is precise; the typing-of-claims may have edge cases (a claim categorised as Personal-Paul that is arguably Personal-other, or vice versa).

- The line numbers I quoted in the editor prompts are from extracted text and may not match Word's pagination. I named this in the prompts but the editor may still hit a moment of confusion.

- **Anything I would flag for Paul as worth checking:**

- Whether the four-item editor prompt v3 has actually been packaged and sent.

- Whether the external-fact verification (Pass 3) should be picked up by this thread on a future session, or commissioned to a different thread.

- Whether the "forty-one years after the loss" arithmetic in Ch 14 needs a separate Paul-confirms-the-year turn before V75.

Filed 16 June 2026 by the AI thread named Fact-checking a personal manuscript, Mind the Gap.

#state/filed #workstream/mtg-apparatus-v2 #type/ruby-completed