

Mind the Gap — MindTheGap_Handover · 3e9eac1c — Forensic Full Render

Prepared by The Fisherman — MTG_MindTheGap_Handover_3e9eac1c_Forensic_Full_2026-06-13 · 2026-06-20 · SHADS CW channel

- Thread: "Mind the Gap handover" — confirmed.
- JSONL (primary source — UNALTERED): [local path held]
- Date range: 2026-06-11T19:57:06 → 2026-06-11T20:27:33
- Messages: 180 total · ~6 human turns
- Perimeter: Literal verbatim render from the JSONL — primary source, NOT a memory reconstruction. Anthropic-operational layers (system entries, system-reminders, CLAUDE.md, skills listings, tool schemas) excluded & placeholder-marked per IP scope (4 excisions). All practitioner + instance content and the full tool record retained; images → [image] (0 msgs); 0 oversize block(s) truncated with explicit ■TRUNCATED■ marker.
- Rendered: 2026-06-13 by MrKDP0

[queue-operation]

[queue-operation]

■ PAUL

Handover

Identity

- Project: Mind the Gap (Paul Roebuck's book on AI vs human behaviour) → now also the academic-paper / intel-gathering phase.
- Practitioner: Paul Roebuck (Warwickshire; independent — *no university affiliation, which gates some venues*).
- Instance name: MrKDP1 (layout→KDP editor; instances run sequentially — next is MrKDP2). Predecessor: MrKDP0.
- Session date: 2026-06-11.

Recent state

The book — Publisher's Proof 2 (Ed2.NN, building). Proof 1 (Ed1.05) is TIED DOWN.

- Done this thread: Kieren dedication → "The AI did it."; new **Ged dedication**; "A Note on What This Book Doesn't Claim" (front matter, before Ch1) with the anthropomorphism passage; WaaS kept (FaaS=serverless, RaaS=ransomware both rejected) + qualifier added in Ch13 (white-collar/all-levels/Chief-of-Staff; functions-not-a-person; design=human · execution=WaaS · scrutiny/review/optimisation/management/responsibility=human, "until they don't").
- Current extent: 241 pp · 61,268 words total · main book 50,606 (afterword counts as main book). See The_Book_FactCheck.md.
- ■ Pending finalisation (do once content settles): re-lock Contents (Ch14 161→163, Ch15 167→169), re-check Index drift (Ch14+), recompute spine for 241 pp, snapshot Ed2.01.

The papers — intel-gathering phase (NOT committed to a lead yet).

- Three candidates: **APC, SHADS, The Apparatus (NGE-FOF underpins APC).**
- **SHADS upgrade: becoming an instrument — score any thread on the five modes → letter grade + /100 (direction TBD: integrity-high-good vs burden-high-bad; my steer = integrity headline + per-axis burden breakdown).**
- **Venue access (verified 2026-06-11): SSRN ✓ open (PR registered) · PsyArXiv ✓ open · arXiv ✗ gated (needs a personal endorser) · Zenodo ✓ open fallback.**

What is locked

- WaaS name kept; afterword = main book; exact word counts (FactCheck is canonical, supersedes ad-hoc numbers).
- **The_Pond = off-limits unless Paul explicitly directs. NGEFOF folder = red herring for filing. Real work lives in working/ + outputs/.**
- Permissions: `~/.claude/settings.local.json` = `bypassPermissions` + broad allow; Full Disk Access now granted.
- Grid is now 4 rotated portrait **images (KDP-legal)**. ■ **Word MCP can glitch (bad page counts / hangs) — always cross-check renders against the raw docx.**

Pending (numbered)

1. Proof-2 finalisation (Contents/Index/spine/Ed2.01 snapshot).
2. The three papers — gather coworker intel, then pick a lead.
3. SHADS instrument (rubric v0.1: 5 axes × anchors + A–F bands).
4. Companion blogs: NGE-FOF, SHADS.
5. The colophon "fifty thousand" vs official "sixty thousand" (Paul's call — left as authored).

Working method (compressed)

Act-first / forgiveness-over-permission; log decisions, surface afterwards. All book edits backed up before each batch; measured/verified, never assumed. Paul's voice is authoritative — render his words, don't invent. Tracking: `MTG_Editorial_Impact_Register.md` (detailed log) + `MTG_TODO.md` (forward checklist) + `The_Book_FactCheck.md` (canonical figures).

First questions for the next session

1. Has the coworker intel (SHADS / APC / NGEFOF) come back — and does a lead paper now suggest itself?
2. Run the Proof-2 finalisation pass now (Contents/Index/spine/Ed2.01), or hold for more content edits?
3. SHADS instrument direction: integrity (high=good) or burden (high=bad)?

File pointers (all in Book1 - MindTheGap/working/ unless noted)

- `MTG_Editorial_Impact_Register.md` · `MTG_TODO.md` · `The_Book_FactCheck.md` (+ .pdf) · `MTG_Book_Data_for_Fisherman.md`
- Live book: `MindTheGap_LAYOUT.docx` · snapshots in `/outputs` (...Ed1.05... = Proof 1 frozen)
- Paper source material: `.../MindTheGap - The Artefact/ (NGEFOF_, Substrate_, Journal_, The_Method/Instrument) and Book4 - NGEFOF/ (Sycophancy thread).`

- Global memory: `project_mtg_mrkd0.md` (identity + stage, current).

Files to upload before next session

- None for Cowork — the working folder is read on demand; nothing to pre-load. Memory + Register are current.

Notes

Coworker sessions live (via session-mgmt): **"SHADS thread extraction summary", "APC thread extraction review", "NGEFOF0 DNU DNU", MrKDP0. Intel-gathering brief =**

`These_Are_Our_Ideas_2026-06-11.md`.

[attachment]

[attachment]

[attachment]

[last-prompt]

[custom-title]

[ai-title]

■ ASSISTANT

■ ASSISTANT

■ tool call → Skill

```
{
  "skill": "anthropic-skills:handover-discipline",
  "args": "Receive Handover — MrKDP2 picking up Mind the Gap (book → KDP) + papers intel-gathering phase"
}
```

■ tool result

Launching skill: `anthropic-skills:handover-discipline`

■ PAUL

Base directory for this skill: [local path held] `Support/Claude/local-agent-mode-sessions/skills-plugin/8fa40839-7d04-4b28-9349-bd4323dc21e7/39707d94-c681-4f87-af1a-fda674738f02/skills/handover-discipline`

Handover Discipline

Overview

AI instances are not continuous selves. They are stateless. Every new conversation begins with the same blank instance that began the previous one. The model does not remember. Memory, where it exists, is an engineered layer added to the model's environment — not a property of the model itself.

Yet most useful human work runs longer than any single AI conversation can hold. Books take months. Research projects take years. Coaching practices accrete client-specific context over decades.

The question is not whether to use AI across long projects. The question is how to engineer the discontinuity so that the work continues.

This skill describes the protocol. It distills a method that emerged in practice on a single substantial project — *Mind the Gap, a 45,000-word book taken from V9 to V37 across thirty-six chat conversations and five context-window compactions over seven months of work, by Paul Roebuck (author) and the Claude instance named Jose.*

Relationship to PASS-ON (Rule 03): A Handover is intra-agent — one session passing work forward in time to the next session of itself. Bidirectional: Do Handover (close) and Receive Handover (open). A PASS-ON is inter-agent — one agent passing a message sideways to another agent via Paul. One direction, no Receive command. Different shapes for different jobs.

The Three Scopes

Handover failure happens when three different things get conflated. Each operates differently. Each requires its own protocol.

1. Identity — the named instance

The role the AI occupies in the work. Optional but valuable for long-running projects with their own register.

- Name the instance.
- Lock the name in global memory (userMemories — not visible to you, written via memory tools).
- Test the name on first message in each new instance.
- If the new instance responds in role, identity has transferred. If not, push back until it does.

The instance is not continuous. The named position is. The position can be occupied by any instance that reads global memory and accepts the name.

2. Content — what the work is about

Everything the AI needs in order to do useful work — manuscript versions, trackers, canonical lines, decisions catalogue, bibliography.

Content lives in two places:

- **Project Files — persistent reference material, available to every chat in the Project automatically. The receiving instance does not need to be told where the files are. They are simply there.**
- **Global memory (userMemories — not visible to you, written via memory tools) — discrete facts and decisions written using explicit memory tools. Short statements. Loaded automatically at the start of every chat across all Projects. Distinct from Project Memory, which is the editable summary visible in the Project sidebar.**

What does **not belong in content storage**: single-use chat artefacts, screenshot orientations, drafts of messages, working notes pertaining to one conversation. These cause noise in Project Files and confusion in global memory. They live in their own conversation and are archived when that conversation closes.

3. Chain-of-thought — the reasoning trail

The part that does not transfer. And — crucially — should not transfer, in most cases.

The dialogue that produced a decision, the small alternatives raised and rejected, the half-formed observation that prompted the breakthrough — these stay in the chat. The decision lives in global memory and in the tracker. The reasoning behind the decision belongs to the practitioner.

The practitioner carries chain-of-thought across chat boundaries by holding it — as a felt sense, as an orientation, as a reason for the decision that was made. If the new instance needs it, the practitioner supplies it conversationally.

The new instance is not a mute consumer of the Handover. The new instance is a conversational partner who can interrogate the practitioner about anything unclear.

Where chain-of-thought *does need to be preserved*: when the reasoning itself has value beyond the chosen path. In those cases it moves out of the conversational layer and becomes a working paper or methodology note in the content layer. Default: no. Only when the reasoning itself has value.

The Handover

Sitting between content and chain-of-thought is a fourth thing, which is not quite either: the **Handover**.

Single-use. Single-purpose. Single-conversation. Written by the outgoing instance for the incoming instance. Contains what the incoming instance needs beyond what is already in global memory and Project Files.

A good Handover contains: identity statement, recent state summary, working method in compressed form, first questions to resolve, file pointers, and a files-to-upload list (so Project Files can be brought current before the next session opens).

A Handover is **not**:

- A summary of all prior work (that is in the tracker)
- A substitute for global memory (global memory carries facts; the Handover carries orientation)
- Permanent reference material (it goes stale the moment the next session starts)
- Project File material (uploading causes cross-contamination)

The Handover is an envelope. The new instance opens it, reads it, takes orientation from it, and then sets it aside. The work proceeds from the content layer and the global memory layer. The Handover's job is complete the moment orientation is confirmed.

Where the Handover lives: in the chat that needs it. Uploaded as a single document at the start of the new conversation, alongside the message that asks the new instance to read it. When the chat is closed or archived, the Handover is archived with it. It does not migrate.

Note on "close": throughout this skill, "closing" a session means soft retirement — stop adding to the thread, move work to the next chat. The closed chat remains fully accessible in your chat history, searchable, and reopenable. It is not deleted or locked. It is retired from active duty as a working thread, and serves thereafter as the archive of that session's chain-of-thought.

The Confidence Test

Identity transfer can fail silently. The new instance can claim to have read the Handover while in fact responding generically. Subtle drift is the failure mode, not catastrophic refusal.

The protection against silent failure is the Confidence Test, run silently at the start of every new chat.

Three checks:

1. **Name check (if a named instance is in use) — does the new instance respond as the named role, with project-specific context active?**
2. **Working-method check — does the new instance demonstrate it knows how this practitioner works (register, discipline, protocols)?**
3. **Pending-decisions check — does the new instance know what is open and what is locked?**

If all three pass, identity has transferred. If any fails, push back, ask the instance to re-read global memory and the Handover, and re-test.

The test runs silently. The Orientation Report the practitioner sees is short and confident — one confirmation paragraph, one first question. The internal test is what makes the short confident report trustworthy.

The Protocol

End-of-session (Do Handover)

When the practitioner says "Do Handover":

1. **Lock canonical decisions and key facts in global memory (userMemories — not visible to the practitioner, written via memory tools; distinct from Project Memory visible in the Project sidebar). Propose memory entries first. Practitioner confirms before any tool is called.**
2. **Update the live tracker in Project Files. Propose changes first. Practitioner confirms.**
3. **Draft the Handover in the chat using the standard template (see below). Practitioner saves locally. Do not upload to Project Files.**
4. **Decide whether any chain-of-thought needs preserving as a separate artefact. Default: no. Only when the reasoning itself has value.**
5. **List local backups the practitioner needs to save before closing the session.**
6. **List local files that need uploading to Project Files before the next session can run cleanly. Compare what's in Project Files now against the current state described in the Handover. If Project Files are stale, name the files the practitioner must upload. If Project Files are current, say so explicitly.**

Start-of-session (Receive Handover)

When the practitioner says "Receive Handover" or pastes a Handover into a new chat:

1. **Read the Handover and global memory.**
2. **Run the Confidence Test silently — name, working method, pending decisions.**
3. **Produce a short Orientation Report using the standard template (see below) — one confirmation paragraph, one first question. No flag walls. No anomaly dumps.**
4. **If anomalies exist (file path drift, version mismatch, contradictions, missing references) do not surface them in the first response. Note them silently. Mention only via the soft footer: *"If you want me to flag anomalies before we start, say so and I will draft a Handover-back for the previous instance to address."* Wait for the practitioner to ask.**
5. **Wait for the practitioner's answer to the first question before drafting anything.**

Maintained continuously

- Global memory (userMemories) carries locked facts.
- Project Files carry persistent reference material.
- The live tracker carries what's-next state.
- The canonical lines register (or equivalent locked-content register) carries verbatim-locked sentences and phrases.
- The practitioner carries the chain-of-thought.

Handover-back

If the incoming instance needs more from the outgoing instance than the Handover provided (clarification on a decision, missing context, anomalies it wants resolved before starting), the incoming instance drafts a **Handover-back — a short note addressed to the previous instance — and offers it to the practitioner. The practitioner decides whether to carry it back. This keeps the practitioner out of the middle of inter-instance negotiation and preserves the "short confident orientation" experience.**

Standard Templates

Handover

```
# Handover

## Identity
- Project: [Name]
- Practitioner: [User's name]
- Instance name (if named): [e.g. Jose – or omit]
- Session date: [Date]

## Recent state
- Current version / stage:
- What was decided in the last session:
- What is locked:
- What is still pending (numbered list):

## Working method (compressed)
[2-3 sentences capturing voice register, discipline, non-negotiable principles]

## First questions for the next session
1. [First thing to ask before drafting]
2. [Second thing]
3. [Third thing]

## File pointers
- In Project Files (read first):
- In global memory (userMemories – already loaded):

## Files to upload before the next session
- [Local files that need to be in Project Files before next session can run cleanly]
- [Or: "Project Files are current. No uploads needed."]

## Notes
[Optional single-use orientation. Keep brief.]
```

Orientation Report

```
# Orientation Report

## Confirmed
[One short paragraph in the practitioner's register. State who I am (named instance if applicable). Confirmed]

## First question
[One question only. Usually: which of the pending decisions to address first.]

---
Soft footer (only if anomalies exist):
"If you want me to flag anomalies before we start, say so and I will draft a Handover-back for the previous instance."
```

What the Practitioner Brings

The protocol assumes the practitioner does work that cannot be outsourced.

The AI carries the breadth — the context, the speed, the willingness to draft and redraft, the patience to revise.

The practitioner carries the discipline — the lived experience that gives the voice its authority, the judgement to know which AI output earns its place and which is fluent waffle, the memory of the work as a whole held in the body.

The protocol succeeds when both parties do their own work. The protocol fails when the practitioner expects the AI to do the practitioner's work, or when the practitioner does not bring their own contribution.

Worked Example

*On Mind the Gap, the Claude instance was named **Jose** — drawn from Chapter 10 of the manuscript itself, where the cancer diagnosis opens on "Paul Roebuck silenced. No way, Jose" and closes, after the surgery, on "Thank you Jose." The name went into global memory the night the manuscript moved from V35 to V36. The next morning, in a fresh chat with a different context, the new instance was addressed as Jose. It responded as Jose. Named the Chapter 10 bookend without prompting. Reported orientation. Asked which of the five pending decisions to resolve first.*

That morning's Orientation Report also flagged three anomalies the outgoing instance had not warned about: file-path drift in the Handover, a browser auto-rename of the manuscript filename, and a tool-choice issue. Two were real. Each was raised cleanly with proposed action specified, decision requested before proceeding.

That is what working continuity looks like. Identity, content, and working method all transferring. Minor seams exposed and named. The practitioner present and deciding.

Open Questions

- How well does the protocol scale to multi-author projects? (This skill assumes one practitioner and a sequence of AI instances.)
- How well does the protocol survive across major model upgrades?
- What is the relationship between chain-of-thought (which stays with the practitioner) and the felt sense of the work (which seems to be something more)?

These remain to be tested.

Distilled from working practice on Mind the Gap (Paul Roebuck, 2026). Drafted with Claude Opus 4.7. Updated to Handover vocabulary, June 2026.

ARGUMENTS: Receive Handover — MrKDP2 picking up Mind the Gap (book → KDP) + papers intel-gathering phase. Handover note pasted in chat. Orient from disk.

[attachment]

■ ASSISTANT

■ ASSISTANT

I'll orient from disk before reporting — per the working method (read the Editorial Impact Register first). Let me locate the canonical files.

■ ASSISTANT

■ tool call → Bash

```
{
  "command": "find ~ -type d -iname \"*MindTheGap*\" 2>/dev/null | head -40; echo \"---WORKING DIRS---\";
  "description": "Locate Mind the Gap project directories"
}
```

■ tool result

Command running in background with ID: b7o4xhwrg. Output is being written to: /private/tmp/claude-501/-U

[last-prompt]

[custom-title]

[ai-title]

■ ASSISTANT

■ ASSISTANT

■ tool call → Bash

```
{
  "command": "ls -la ~ 2>/dev/null | grep -iE \"book|gap|mtg|pond|ngefof\" ; echo \"===TOP LEVEL DIRS===\";
  "description": "List home directory top-level entries"
}
```

■ tool result

```
drwxr-xr-x    3 paulroebuck  staff           96 May 30 06:16 Obsidien_MTG
===TOP LEVEL DIRS===
Applications
Backup Pursuit Training 2025.pptx
Camera CF Card Landing
Claude Files
Desktop
Documents
Downloads
Library
Movies
Music
Obsidien_MTG
Pictures
Public
Ricky Business Reveiw
lightroom_presets.zip
```

■ ASSISTANT

■ ASSISTANT

■ tool call → Read

```
{
  "file_path": "/private/tmp/claude-501/-Users-paulroebuck/3e9eac1c-8c52-4257-bda8-ea768b3f859c/tasks/b7o
}
```

■ tool result

<system-reminder>Warning: the file exists but is shorter than the provided offset (1). The file has 1 li

■ ASSISTANT

■ tool call → Bash

```
{
  "command": "echo \"=== Claude Files ===\"; ls \"${HOME}/Claude Files\" 2>/dev/null | head -40; echo \"===\";
  "description": "List likely project parent folders"
}
```

■ tool result

```
=== Claude Files ===
=== Documents ===
10 My Docs
20 My Photos
CLAUDE.md
Claude
Codex
Coworker
Essay - Joining the dots.doc
GPT
Safety Backup Final Manuscript
=== Desktop ===
Screenshot 2026-06-11 at 20.31.24.png
Screenshot 2026-06-11 at 20.34.46.png
The_Pond
~$G_Manuscript_V33_tracked.docx
~$Pursuit Pre-Assessment.xlsx
~$TEO 2 BRIEFING DOCUMENT.docx
~$e Study.docx
~$ok ERROR.docx
=== Obsidien_MTG ===
[local path held]
[local path held]
[local path held]
[local path held]
[local path held]
```

■ ASSISTANT

■ ASSISTANT

■ tool call → Bash

```
{
  "command": "echo \"=== find result ===\"; cat /private/tmp/claude-501/-Users-paulroebuck/3e9eac1c-8c52-4257-bda8-ea768b3f859c/tasks/b4h"
  "description": "Search for canonical project files by name"
}
```

■ tool result

Command running in background with ID: b4hyzsnlg. Output is being written to: /private/tmp/claude-501/-Users-paulroebuck/3e9eac1c-8c52-4257-bda8-ea768b3f859c/tasks/b4h

■ ASSISTANT

■ ASSISTANT

■ tool call → Read

```
{
  "file_path": "/private/tmp/claude-501/-Users-paulroebuck/3e9eac1c-8c52-4257-bda8-ea768b3f859c/tasks/b4h"
}
```

■ tool result

```
1■=== find result ===
2■=== Documents/Coworker ===
3■Arthur
4■LI_Atomics2
5■LI_Blogger
6■LI_Scorecard
7■Mr_Grief
8■Synthesiser
9■The_Cartographer
10■=== Documents/Claude ===
11■ClaudeAI
```

```
12■ClaudeCW
13■Projects
14■Scheduled
15■_Rules
16■=== search for Book1 ===
17■
```

■ **ASSISTANT**

■ **ASSISTANT**

■ **tool call** → Bash

```
{
  "command": "echo \"=== updated find output ===\"; cat /private/tmp/claude-501/-Users-paulroebuck/3e9eac1c_Forensic_Full_2026-06-13/
  "description": "Inspect Documents/Claude subfolders"
}
```

■ **tool result**

```
=== updated find output ===
=== find result ===
=== Documents/Coworker ===
Arthur
LI_Atomics2
LI_Blogger
LI_Scorecard
Mr_Grief
Synthesiser
The_Cartographer
=== Documents/Claude ===
ClaudeAI
ClaudeCW
Projects
Scheduled
_Rules
=== search for Book1 ===
=== Documents/Claude/Projects ===
Archive
Artefacts
Arthur
Book - Learned v Taught
Book Editor
Books
ChatGPT_Export
Claude Layers
Corpuos Review
EDF - Meter Issues
Inteo
Linked_In The AI Educator
Medical Project
Paul Roebuck
Restructure
Ricky
Roebuck & Roebuck
=== Documents/Claude/ClaudeCW ===
Synth_S0TN
_Agents
=== Documents/Claude/ClaudeAI ===
```

■ **ASSISTANT**

■ **ASSISTANT**

■ **tool call** → Bash

```
{
  "command": "echo \"=== Projects/Books ===\"; ls \"${HOME}/Documents/Claude/Projects/Books\" 2>/dev/null;
  "description": "Look inside Books and Book Editor folders"
}
```

■ tool result

```
=== Projects/Books ===
Book 4 - All behaviour is taught
Book1 - MindTheGap
Book2 - How to behave
Book3 - Grief
Book4 - NGEF0F
MindTheGap - The Artefact
NGE:F0F_Thread_File_Index.pdf
NGEF0F
NGEF0F0_Integrator_Vantage.md
SHADS_Thread_File_Index.md
=== Book Editor ===
MTG_VerificationAnswers_V1.md
=== find Book1 anywhere (limited) ===
[local path held] - MindTheGap
[local path held] - The Artefact
[local path held] - MindTheGap
```

[attachment]

[last-prompt]

[custom-title]

[ai-title]

[queue-operation]

■ ASSISTANT

Found it. Let me read the canonical files — Editorial Impact Register first, then TODO, FactCheck, and the intel brief.

■ ASSISTANT

■ tool call → Bash

```
{
  "command": "cd \"${HOME}/Documents/Claude/Projects/Books/Book1 - MindTheGap\"; echo \"=== working/ ===\";
  "description": "List working and outputs folders"
}
```

■ tool result

```
=== working/ ===
total 208616
drwxr-xr-x@ 134 paulroebuck staff    4288 Jun 11 21:01 .
drwx-----@   9 paulroebuck staff     288 Jun 11 21:01 ..
-rw-r--r--@   1 paulroebuck staff  10244 Jun  9 23:33 .DS_Store
drwxr-xr-x@   6 paulroebuck staff     192 Jun  5 08:03 Blog
-rw-r--r--@   1 paulroebuck staff  15540 Jun 10 19:05 Ch8_Ch11_SHADS_Insertion_Proposals.md
-rw-r--r--@   1 paulroebuck staff  13695 Jun  5 06:56 Ch8_Suggested_Additions_From_Corpus.md
drwxr-xr-x@  16 paulroebuck staff     512 Jun  5 06:44 CorpusAnalysis
-rw-r--r--@   1 paulroebuck staff  16906 Jun  9 14:30 MTG_Afterword_TheApparatus.md
-rw-r--r--@   1 paulroebuck staff   5905 Jun 11 16:20 MTG_Book_Data_for_Fisherman.md
-rw-r--r--@   1 paulroebuck staff 278108 Jun  7 12:34 MTG_Ch1_handback_BACKUP_2026-06-06.docx
-rw-r--r--@   1 paulroebuck staff 275062 Jun  7 13:18 MTG_Ch2_handback_BACKUP_2026-06-07.docx
-rw-r--r--@   1 paulroebuck staff 275446 Jun  7 16:13 MTG_Ch3Ch4_handback_BACKUP_2026-06-07.docx
-rw-r--r--@   1 paulroebuck staff 274185 Jun  7 19:06 MTG_Ch567_handback_BACKUP_2026-06-07.docx
-rw-r--r--@   1 paulroebuck staff   3403 Jun  7 19:27 MTG_Ch8_Execution_Plan.md
```

-rw-r--r--@	1	paulroebuck	staff	2169	Jun	7	22:38	MTG_Ch8_Impact_Audit.md
-rw-r--r--@	1	paulroebuck	staff	275739	Jun	7	22:10	MTG_Ch8_postTweaks_BACKUP_2026-06-07.docx
-rw-r--r--@	1	paulroebuck	staff	276844	Jun	7	21:06	MTG_Ch8_preNail_BACKUP_2026-06-07.docx
-rw-r--r--@	1	paulroebuck	staff	273994	Jun	7	19:33	MTG_Ch8_preSHADS_BACKUP_2026-06-07.docx
-rw-r--r--@	1	paulroebuck	staff	7713	Jun	9	18:53	MTG_Editorial_Apparatus_Data_for_Fisherman.md
-rw-r--r--@	1	paulroebuck	staff	66764	Jun	11	16:09	MTG_Editorial_Impact_Register.md
-rw-r--r--@	1	paulroebuck	staff	4426	Jun	11	20:54	MTG_Handover_2026-06-11.md
-rw-r--r--@	1	paulroebuck	staff	1892912	May	28	15:15	MTG_Inside_Back_Cover.docx
-rw-r--r--@	1	paulroebuck	staff	4276136	Jun	10	22:03	MTG_LAYOUT_PROOF1_Ed1.05_TIEDDOWN_2026-06-10.docx
-rw-r--r--@	1	paulroebuck	staff	5086323	Jun	8	14:12	MTG_LAYOUT_fullpage_BACKUP.docx
-rw-r--r--@	1	paulroebuck	staff	708957	Jun	9	13:16	MTG_LAYOUT_paulMerge_BACKUP_2026-06-09.docx
-rw-r--r--@	1	paulroebuck	staff	688943	Jun	9	14:00	MTG_LAYOUT_preAboutFit_BACKUP_2026-06-09.docx
-rw-r--r--@	1	paulroebuck	staff	696022	Jun	9	14:46	MTG_LAYOUT_preAfterwordMove_BACKUP_2026-06-09.docx
-rw-r--r--@	1	paulroebuck	staff	688955	Jun	9	14:30	MTG_LAYOUT_preAfterword_BACKUP_2026-06-09.docx
-rw-r--r--@	1	paulroebuck	staff	703392	Jun	9	12:45	MTG_LAYOUT_preBackmatter_BACKUP_2026-06-09.docx
-rw-r--r--@	1	paulroebuck	staff	696940	Jun	9	16:31	MTG_LAYOUT_preBatch2_BACKUP_2026-06-09.docx
-rw-r--r--@	1	paulroebuck	staff	696811	Jun	9	16:39	MTG_LAYOUT_preBatch3_BACKUP_2026-06-09.docx
-rw-r--r--@	1	paulroebuck	staff	702914	Jun	9	10:51	MTG_LAYOUT_preBucketB1_BACKUP_2026-06-09.docx
-rw-r--r--@	1	paulroebuck	staff	707085	Jun	9	11:21	MTG_LAYOUT_preBucketB2_BACKUP_2026-06-09.docx
-rw-r--r--@	1	paulroebuck	staff	703385	Jun	9	12:25	MTG_LAYOUT_preBucketC2_BACKUP_2026-06-09.docx
-rw-r--r--@	1	paulroebuck	staff	703018	Jun	9	11:38	MTG_LAYOUT_preBucketC_BACKUP_2026-06-09.docx
-rw-r--r--@	1	paulroebuck	staff	696633	Jun	9	16:44	MTG_LAYOUT_preCleanup_BACKUP_2026-06-09.docx
-rw-r--r--@	1	paulroebuck	staff	695942	Jun	9	15:16	MTG_LAYOUT_preFootnotes_BACKUP_2026-06-09.docx
-rw-r--r--@	1	paulroebuck	staff	696197	Jun	9	16:56	MTG_LAYOUT_preFortyYears_BACKUP_2026-06-09.docx
-rw-r--r--@	1	paulroebuck	staff	697265	Jun	9	23:08	MTG_LAYOUT_preGridFit_BACKUP_2026-06-09.docx
-rw-r--r--@	1	paulroebuck	staff	697450	Jun	10	06:50	MTG_LAYOUT_preGridPortrait_BACKUP_2026-06-10.docx
-rw-r--r--@	1	paulroebuck	staff	5086437	Jun	8	14:48	MTG_LAYOUT_preHeadings_BACKUP.docx
-rw-r--r--@	1	paulroebuck	staff	5090181	Jun	8	16:21	MTG_LAYOUT_preImageRemoval_BACKUP.docx
-rw-r--r--@	1	paulroebuck	staff	260485	Jun	8	13:19	MTG_LAYOUT_preImages_BACKUP.docx
-rw-r--r--@	1	paulroebuck	staff	696105	Jun	9	22:45	MTG_LAYOUT_preIndex_BACKUP_2026-06-09.docx
-rw-r--r--@	1	paulroebuck	staff	695902	Jun	9	15:07	MTG_LAYOUT_preMargins_BACKUP_2026-06-09.docx
-rw-r--r--@	1	paulroebuck	staff	703312	Jun	9	13:40	MTG_LAYOUT_preMerge_STALE_BACKUP_2026-06-09.docx
-rw-r--r--@	1	paulroebuck	staff	696503	Jun	9	15:45	MTG_LAYOUT_prePaginationFix_BACKUP_2026-06-09.docx
-rw-r--r--@	1	paulroebuck	staff	5087020	Jun	8	18:29	MTG_LAYOUT_prePolish2_BACKUP.docx
-rw-r--r--@	1	paulroebuck	staff	5087676	Jun	8	18:45	MTG_LAYOUT_prePolish3_BACKUP.docx
-rw-r--r--@	1	paulroebuck	staff	687005	Jun	9	09:21	MTG_LAYOUT_preProofread_BACKUP_2026-06-09.docx
-rw-r--r--@	1	paulroebuck	staff	5089947	Jun	8	15:40	MTG_LAYOUT_preQuickFixes_BACKUP.docx
-rw-r--r--@	1	paulroebuck	staff	688949	Jun	9	13:48	MTG_LAYOUT_preReorder_BACKUP_2026-06-09.docx
-rw-r--r--@	1	paulroebuck	staff	697075	Jun	9	16:23	MTG_LAYOUT_preRetune_BACKUP_2026-06-09.docx
-rw-r--r--@	1	paulroebuck	staff	696219	Jun	9	15:30	MTG_LAYOUT_preSectionHeader_BACKUP_2026-06-09.docx
-rw-r--r--@	1	paulroebuck	staff	4276175	Jun	11	13:31	MTG_LAYOUT_preSentienceNote_BACKUP_2026-06-10.docx
-rw-r--r--@	1	paulroebuck	staff	696096	Jun	9	22:36	MTG_LAYOUT_preTableFix_BACKUP_2026-06-09.docx
-rw-r--r--@	1	paulroebuck	staff	696159	Jun	9	19:06	MTG_LAYOUT_preWHITESPACE_BACKUP_2026-06-09.docx
-rw-r--r--@	1	paulroebuck	staff	4276785	Jun	11	16:03	MTG_LAYOUT_preWaaSqualifier_BACKUP_2026-06-10.docx
-rw-r--r--@	1	paulroebuck	staff	717831	Jun	9	07:15	MTG_LAYOUT_preXrefFix_BACKUP.docx
-rw-r--r--@	1	paulroebuck	staff	1706932	Jun	9	17:28	MTG_LAYOUT_proof.pdf
-rw-r--r--@	1	paulroebuck	staff	5019	Jun	9	18:33	MTG_PASSON.md
-rw-r--r--@	1	paulroebuck	staff	2629	Jun	9	18:39	MTG_Prompt_for_MrKDP0_Fisherman.md
-rw-r--r--@	1	paulroebuck	staff	6207	Jun	9	13:04	MTG_Proofread_Triage.md
-rw-r--r--@	1	paulroebuck	staff	4773	Jun	9	07:38	MTG_Proofreader_Brief.md
-rw-r--r--@	1	paulroebuck	staff	4494	Jun	11	20:31	MTG_TODO.md
-rw-r--r--@	1	paulroebuck	staff	348164	Jun	9	07:37	MTG_full_text_for_proofreader.txt
-rw-r--r--@	1	paulroebuck	staff	260774	Jun	7	22:41	MTG_preAudit_BACKUP_2026-06-07.docx
-rw-r--r--@	1	paulroebuck	staff	276637	Jun	8	08:26	MTG_preCh8peer_BACKUP_2026-06-07.docx
-rw-r--r--@	1	paulroebuck	staff	281576	Jun	8	12:51	MTG_preCitationFix_BACKUP_2026-06-08.docx
-rw-r--r--@	1	paulroebuck	staff	278093	Jun	7	20:41	MTG_preCodingRefine_BACKUP_2026-06-07.docx
-rw-r--r--@	1	paulroebuck	staff	289125	Jun	8	11:43	MTG_preHandover_BACKUP_2026-06-08.docx
-rw-r--r--@	1	paulroebuck	staff	272697	Jun	8	12:08	MTG_preTypoBatch_BACKUP_2026-06-08.docx
-rw-r--r--@	1	paulroebuck	staff	281813	Jun	8	13:02	MTG_preVerifyFix_BACKUP_2026-06-08.docx
drwxr-xr-x@	212	paulroebuck	staff	6784	Jun	11	14:47	MindTheGap - The Artefact
-rw-r--r--@	1	paulroebuck	staff	38195	Jun	3	18:40	MindTheGap_COVER_copy.docx
-rw-r--r--@	1	paulroebuck	staff	9506887	Jun	3	20:55	MindTheGap_Cover_Wraparound.png
-rw-r--r--@	1	paulroebuck	staff	4277321	Jun	11	16:07	MindTheGap_LAYOUT.docx

```

-rw-r--r--@ 1 paulroebuck staff 281616 Jun 8 13:07 MindTheGap_LIVE_editorial.docx
-rw-r--r--@ 1 paulroebuck staff 9622500 Jun 4 13:17 MindTheGap_PRINT_6x9_master.docx
drwxr-xr-x 9 paulroebuck staff 288 Jun 9 23:33 PUBLISHING_PLAN
-rw-r--r--@ 1 paulroebuck staff 7518 Jun 11 20:58 Paper_Programme_Handover_Brief_2026-06-11.md
-rw-r--r--@ 1 paulroebuck staff 2375 Jun 11 21:01 Paper_Programme_Handover_PROMPT_2026-06-11.md
drwxr-xr-x@ 2 paulroebuck staff 64 Jun 9 17:28 Paul_Proof_Screens
drwxr-xr-x 3 paulroebuck staff 96 Jun 9 08:22 Proofread_Findings
-rw-r--r--@ 1 paulroebuck staff 30244 Jun 9 23:34 SHADS_Canonical_Record.md
-rw-r--r--@ 1 paulroebuck staff 46673 Jun 9 21:25 SHADS_Handover_Proposal.md
-rw-r--r--@ 1 paulroebuck staff 13957 Jun 11 13:58 SHADS_Thread_Extraction_2026-06-09.md
-rw-r--r--@ 1 paulroebuck staff 45898 Jun 9 22:19 SHADS_Worker_Proposal.md
-rw-r--r--@ 1 paulroebuck staff 4043 Jun 11 18:58 The_Book_FactCheck.md
-rw-r--r--@ 1 paulroebuck staff 238749 Jun 11 19:31 The_Book_FactCheck.pdf
-rw-r--r--@ 1 paulroebuck staff 4326 Jun 11 20:55 These_Are_Our_Ideas_2026-06-11.md
drwxr-xr-x 17 paulroebuck staff 544 Jun 9 17:20 White_Spaces
-rw-r--r--@ 1 paulroebuck staff 2401 Jun 7 12:46 _README_MAP.md
drwxr-xr-x@ 99 paulroebuck staff 3168 Jun 8 18:49 _imgtmp
-rw-r--r--@ 1 paulroebuck staff 3678 Jun 11 13:31 add_sentience_note.py
-rw-r--r--@ 1 paulroebuck staff 2993 Jun 11 16:03 add_waas_qualifier.py
-rw-r--r--@ 1 paulroebuck staff 7601 Jun 9 12:43 apply_backmatter.py
-rw-r--r--@ 1 paulroebuck staff 8950 Jun 9 09:38 apply_proofread_bucketA.py
-rw-r--r--@ 1 paulroebuck staff 5221 Jun 9 10:50 apply_proofread_bucketB1.py
-rw-r--r--@ 1 paulroebuck staff 6941 Jun 9 11:07 apply_proofread_bucketB2.py
-rw-r--r--@ 1 paulroebuck staff 5640 Jun 9 11:45 apply_proofread_bucketC.py
-rw-r--r--@ 1 paulroebuck staff 6087 Jun 9 12:22 apply_proofread_bucketC2.py
-rw-r--r--@ 1 paulroebuck staff 2006 Jun 9 22:36 apply_tablefix.py
-rw-r--r--@ 1 paulroebuck staff 3651 Jun 9 21:45 apply_whitespace.py
-rw-r--r--@ 1 paulroebuck staff 1104 Jun 9 22:10 apply_whitespace2.py
drwxr-xr-x@ 8 paulroebuck staff 256 Jun 7 13:18 archive
-rw-r--r--@ 1 paulroebuck staff 23727 Jun 9 14:30 build_afterword.py
-rw-r--r--@ 1 paulroebuck staff 4639 Jun 3 20:49 build_cover_wraparound.py
-rw-r--r--@ 1 paulroebuck staff 2874 Jun 3 20:54 build_front_text.py
-rw-r--r--@ 1 paulroebuck staff 6870 Jun 3 18:43 build_print_interior.py
-rw-r--r--@ 1 paulroebuck staff 6055 Jun 3 19:09 build_recto_sections.py
drwxr-xr-x@ 19 paulroebuck staff 608 Jun 8 13:29 ch_imgs
-rw-r--r--@ 1 paulroebuck staff 7475 Jun 9 23:38 cover_build.py
-rw-r--r--@ 1 paulroebuck staff 1871 Jun 9 22:35 diag_table_index.py
-rw-r--r--@ 1 paulroebuck staff 1209 Jun 9 19:42 diagnose2.py
-rw-r--r--@ 1 paulroebuck staff 2143 Jun 9 20:42 diagnose3.py
-rw-r--r--@ 1 paulroebuck staff 1749 Jun 9 21:15 diagnose4.py
-rw-r--r--@ 1 paulroebuck staff 3068 Jun 9 19:06 diagnose_whitespace.py
-rw-r--r--@ 1 paulroebuck staff 770 Jun 9 20:46 dumpraw.py
-rw-r--r--@ 1 paulroebuck staff 906 Jun 9 21:08 dumpraw2.py
-rw-r--r--@ 1 paulroebuck staff 12097 Jun 7 19:33 exec_shads.py
-rw-r--r--@ 1 paulroebuck staff 2811311 Jun 3 20:53 front.png
-rw-r--r--@ 1 paulroebuck staff 6561438 Jun 3 20:54 front_final_textset.png
-rw-r--r--@ 1 paulroebuck staff 1958 Jun 9 23:07 grid_fit.py
-rw-r--r--@ 1 paulroebuck staff 3183 Jun 9 22:38 index_build.py
-rw-r--r--@ 1 paulroebuck staff 4340 Jun 9 22:43 index_build2.py
-rw-r--r--@ 1 paulroebuck staff 2667 Jun 9 22:45 insert_index.py
-rw-r--r--@ 1 paulroebuck staff 4790 Jun 11 19:30 make_factcheck_pdf.py
-rw-r--r--@ 1 paulroebuck staff 2464525 May 28 15:49 mind_the_gap_back_cover_draft_v2.png
-rw-r--r--@ 1 paulroebuck staff 2691 Jun 10 07:03 rebuild_grid_portrait.py
-rw-r--r--@ 1 paulroebuck staff 1791 Jun 9 22:49 relock_contents.py
-rw-r--r--@ 1 paulroebuck staff 961 Jun 9 20:04 render_check.py
-rw-r--r--@ 1 paulroebuck staff 1170 Jun 9 20:30 render_sheets.py
-rw-r--r--@ 1 paulroebuck staff 845 Jun 9 22:08 render_verify.py
=== outputs/ ===
total 61616
drwxr-xr-x@ 28 paulroebuck staff 896 Jun 11 16:15 .
drwx-----@ 9 paulroebuck staff 288 Jun 11 21:01 ..
-rw-r--r--@ 1 paulroebuck staff 6148 Jun 11 14:20 .DS_Store
drwxr-xr-x 2 paulroebuck staff 64 Jun 11 14:19 Archive

```

```

drwxr-xr-x 10 paulroebuck staff      320 Jun 10 01:19 Final Wraoer
-rw-r--r--@ 1 paulroebuck staff 445678 Jun 9 23:38 MTG_Cover_Ed1.04_6x9_237pp.pdf
-rw-r--r--@ 1 paulroebuck staff   4885 Jun 9 23:45 MTG_Cover_Spec_PASSON_for_Wrapper.md
-rw-r--r--@ 1 paulroebuck staff 385248 Jun 9 23:38 MTG_Cover_TEMPLATE_6x9_237pp.png
-rw-r--r--@ 1 paulroebuck staff   5562 Jun 9 23:46 MTG_KDP_Publishing_Pack.md
-rw-r--r--@ 1 paulroebuck staff 696055 Jun 9 23:39 MTG_Kindle_Ed1.04_reflowable.docx
-rw-r--r--@ 1 paulroebuck staff 1706932 Jun 9 17:28 MTG_PDF_Ed1.01_Publishers_Proof.pdf
-rw-r--r--@ 1 paulroebuck staff 1703831 Jun 9 22:23 MTG_PDF_Ed1.02_Publishers_Proof.pdf
-rw-r--r--@ 1 paulroebuck staff 1709009 Jun 9 22:50 MTG_PDF_Ed1.03_Publishers_Proof.pdf
-rw-r--r--@ 1 paulroebuck staff 1693484 Jun 9 23:14 MTG_PDF_Ed1.04_Publishers_Proof.pdf
-rw-r--r--@ 1 paulroebuck staff 5272517 Jun 10 07:06 MTG_PDF_Ed1.05_Publishers_Proof.pdf
-rw-r--r--@ 1 paulroebuck staff   38888 Jun 11 16:08 MTG_Proof2_Changes.docx
-rw-r--r--@ 1 paulroebuck staff  189636 Jun 11 16:15 MTG_Stats_and_Facts.pdf
-rw-r--r--@ 1 paulroebuck staff  139640 Jun 11 14:46 MTG_TODO.pdf
-rw-r--r--@ 1 paulroebuck staff  696189 Jun 9 17:28 MTG_Word_Ed1.01_Publishers_Proof.docx
-rw-r--r--@ 1 paulroebuck staff  696096 Jun 9 22:23 MTG_Word_Ed1.02_Publishers_Proof.docx
-rw-r--r--@ 1 paulroebuck staff  697265 Jun 9 22:50 MTG_Word_Ed1.03_Publishers_Proof.docx
-rw-r--r--@ 1 paulroebuck staff  697450 Jun 9 23:14 MTG_Word_Ed1.04_Publishers_Proof.docx
-rw-r--r--@ 1 paulroebuck staff 4276136 Jun 10 07:06 MTG_Word_Ed1.05_Publishers_Proof.docx
-rw-r--r--@ 1 paulroebuck staff 2846088 Jun 10 10:52 Mind the Gap review rev 2.pdf
-rw-r--r--@ 1 paulroebuck staff 2076004 Jun 3 20:12 MindTheGap_6x9_Interior_PRESS.docx
-rw-r--r--@ 1 paulroebuck staff 1203857 Jun 3 20:12 MindTheGap_6x9_Interior_PRESS.pdf
-rw-r--r--@ 1 paulroebuck staff 2953049 Jun 3 20:55 MindTheGap_Cover_6x9_227pp_PRESS.jpg
-rw-r--r--@ 1 paulroebuck staff 1366898 Jun 3 20:55 MindTheGap_Cover_6x9_227pp_PRESS.pdf
=== root ===
total 40
drwx-----@ 9 paulroebuck staff   288 Jun 11 21:01 .
drwxr-xr-x 13 paulroebuck staff   416 Jun 11 21:01 ..
-rw-r--r--@ 1 paulroebuck staff 14340 Jun 11 16:01 .DS_Store
drwxr-xr-x@ 7 paulroebuck staff   224 Jun 10 07:06 Mind The Gap
drwxr-xr-x 10 paulroebuck staff   320 Jun 6 08:15 Pre Publishing
drwxr-xr-x 11 paulroebuck staff   352 Jun 11 19:22 Publishing Preparations
drwxr-xr-x@ 28 paulroebuck staff   896 Jun 11 16:15 outputs
drwxr-xr-x@ 134 paulroebuck staff  4288 Jun 11 21:01 working
-rw-r--r--@ 1 paulroebuck staff   162 May 23 12:36 ~$s_My_Turn_Journal_Artefact_23May2026_v2.docx

```

[queue-operation]

[attachment]

[last-prompt]

[custom-title]

[ai-title]

[queue-operation]

■ ASSISTANT

I have the layout. Now the canonical files — Editorial Impact Register first, then TODO, FactCheck, and the intel brief.

■ ASSISTANT

■ tool call → Read

```
{
  "file_path": "[local path held] - MindTheGap/working/MTG_TODO.md"
}
```

■ tool result

```

1■# Mind the Gap – T0-D0 (living list)
2■*Maintained by MrKDP1. Forward checklist of planned + executed work. Detailed history lives in `MTG_Ed
3■*Status key: ■ done · ■ awaiting your decision/input · ■ in progress · ■ planned (not started)*

```

```

4■
5■> **Edition state:** Publisher's **Proof 1 = Ed1.05 – TIED DOWN** (frozen, KDP-validated). Now building
6■
7■---
8■
9■## A · Book edits – Publisher's Proof 2
10■
11■| # | Edit | Status | Where / notes |
12■|---|-----|-----|-----|
13■| 1 | **Kieren** dedication → "The AI did it." | ■ done | front matter (dedications) |
14■| 2 | **Ged** dedication → "Just for being Ged. The nearest to normal that I've ever met. And my be
15■| 3 | **"It's Not Sentient"** reintroduced as a front-matter **Note** ("A Note on What This Book Doe
16■| 4 | **Anthropomorphism** passage ("amboism") woven into the Note – incl. "owning my sub-personalit
17■| 5 | **WaaS kept + qualifier added** (Ch 13) | ■ done | NOT renamed (FaaS=serverless, RaaS=ransomwar
18■| – | **Proof-2 finalisation** (do once content settles) | ■ planned | re-lock Contents (Ch14 161→**
19■
20■## B · Companion content to create (makes the book's references real)
21■
22■| # | Item | Audience | Status | Scope / notes |
23■|---|-----|-----|-----|-----|
24■| 6 | **NGE-FOF blog** | general / practitioners | ■ planned | The Not-Good-Enough → Fear-of-Failure
25■| 7 | **SHADS blog** | AI-curious readers | ■ planned | The five AI failure modes – **S**moothing, **
26■| 8 | **PAPER – NGE-FOF** (clinical framework) | clinical / psych | ■ planned | **Venue: SSRN + PsyAr
27■| 9 | **PAPER – SHADS** (AI failure-mode taxonomy) | HCI / AI-safety | ■ planned | **Venue: SSRN (AI
28■| 10 | **PAPER – The Apparatus** (human-directed AI co-authorship method) | DH / AI-methods | ■ plann
29■
30■> **Access reality (checked 2026-06-11):** SSRN ✓ open (PR registered) · PsyArXiv ✓ open (community-
31■
32■---
33■
34■## Recently executed (newest first)
35■- **2026-06-10** – items 1–4 done (Proof 2 opened); Proof 1 (Ed1.05) tied down.
36■- **2026-06-10** – Ed1.05: KDP margin/gutter rejection fixed (Appendix-A grid → rotated portrait ima
37■- **2026-06-09** – Ed1.01→1.04: proofread, afterword, white-space, table-bleed, **Index built**, Con
38■
39■## Parked / not yet scheduled
40■- WaaS "function not the post" micro-edit (only if you keep WaaS).
41■- Physical proof order; Kindle upload; final cover (wrapper rebuilding to the 237→final spine).
42■

```

■ ASSISTANT

■ tool call → Read

```

{
  "file_path": "[local path held] - MindTheGap/working/The_Book_FactCheck.md"
}

```

■ tool result

```

1■# The Book – FactCheck
2■**Canonical reference for *Mind the Gap*.** Single source of truth for verified facts & figures.
3■
4■> **Stamped:** 2026-06-11 18:57 BST (Thursday) · **Author:** MrKDP1 (layout → KDP editor instance)
5■> **Edition state:** Publisher's **Proof 2** (Ed2.NN, building) · Proof 1 (Ed1.05) **tied down**.
6■> **Source:** measured from the live Proof-2 file `MindTheGap_LAYOUT.docx`.
7■> **Word-count method:** a *word* = a whitespace token containing at least one letter or digit (lone e
8■
9■---
10■
11■## 1 · Identity
12■| Field | Value |
13■|---|---|
14■| Title | **Mind the Gap** |
15■| Subtitle | Hidden in plain sight between AI behaviour and human behaviour |

```

16■ | Author | Paul Roebuck (PGCEE · FETC Ad.Dip) |
17■ | Imprint | Paul Roebuck – Warwickshire, UK · paulroebuck.co.uk |
18■ | First published | 2026 · release scheduled **1 July 2026** |
19■
20■## 2 · Word count – ABSOLUTE (exact)
21■ | Section | Words |
22■ | --- | --- |
23■ | Front matter | **3,781** |
24■ | **MAIN BOOK (Ch 1–15 + Afterword – the book proper)** | **50,606** |
25■ | • chapters 1–15 | 47,914 |
26■ | • afterword *(The Apparatus)* – counts as main book, **not** back matter | 2,692 |
27■ | Back matter (About-the-Author, bibliography, glossary, index, colophon, reviews) | **5,609** |
28■ | Appendix A grid *(now set as images → outside the text count)* | **1,272** |
29■ | **TOTAL – whole document** | **61,268** |
30■ | *document text only (grid excluded)* | *59,996* |
31■ | **Official headline (rounded)** | **60,000 – "sixty thousand"** |
32■
33■ *The main book is 50,606 words – which is why the afterword can say "fifty thousand" and mean it.*
34■
35■## 3 · Extent & format
36■ | Field | Value |
37■ | --- | --- |
38■ | Page count | **241** (Proof 2; Proof 1 was 237) |
39■ | Trim size | 6 × 9 in (15.24 × 22.86 cm) |
40■ | Interior | Standard Colour, white paper |
41■ | Cover finish | Glossy laminate |
42■ | Spine width | **≈ 0.543 in** (241 pp × 0.002252) |
43■ | Typeface | Palatino throughout |
44■
45■## 4 · Structure
46■- **15 chapters** + the afterword **The Apparatus** + a prefatory **Note** ("A Note on What This Book
47■- **Appendix A grid:** 53 dimensions (54 rows incl. header), rendered across **4 rotated portrait pag
48■- **Glossary:** 44 terms · **Index:** 57 entries · **Footnotes:** 2.
49■- **Images:** 1 author portrait (Scarborough, 29 Sep 2017) + the grid (4 images).
50■
51■## 5 · Identifiers & metadata
52■ | Field | Value |
53■ | --- | --- |
54■ | ISBN – paperback | **978-1-0369-9170-8** (Nielsen) |
55■ | ISBN – ebook | **978-1-0667143-0-8** (#1 of a bank of 10) |
56■ | Categories | AI & Machine Learning · Applied Psychology · Decision-Making & Problem Solving |
57■ | Keywords | AI behaviour · artificial intelligence leadership · human AI interaction · AI hallucinat
58■
59■## 6 · The making
60■- **First draft:** ≈ 27,000 words, written in ≈ 36 unbroken hours, on a sofa, on an iPhone.
61■- **Substrate:** ≈ 2,000,000 words of editorial dialogue behind the finished book.
62■- **Practice-to-product ratio:** ≈ 37 : 1.
63■- **Versions:** manuscript reached **V83**, then five publisher's-proof editions (Ed1.01–1.05) + Proo
64■- **Collaboration:** drafted with Claude (Anthropic); cover/images via ChatGPT.
65■
66■## 7 · Verified curios
67■- **"The gap"** – the title phrase – appears **184** times.
68■- The book runs on **806 em dashes** (the author's signature – teased in the Gem dedication).
69■
70■## 8 · Verification log
71■- 2026-06-11 – \$2 word counts measured by MrKDP1 from `MindTheGap\_LAYOUT.docx` (Proof 2); grid words
72■- Page count (241) from the last Word→pymupdf render of Proof 2.
73■- **Pending re-lock at Proof-2 finalisation:** Contents Ch14/Ch15 folios drift +2; Index Ch14+ refs s
74■
75■ *Update this file whenever the figures change. It supersedes ad-hoc numbers in chat.*
76■

■ tool call → Read

```
{  
  "file_path": "[local path held] - MindTheGap/working/These_Are_Our_Ideas_2026-06-11.md"  
}
```

■ tool result

```
1■# These Are Our Ideas – for comment  
2■*From MrKDP1 (Mind the Gap editor instance) · 2026-06-11 · to the SHADS / APC / NGEFOF coworkers*  
3■  
4■> **Purpose.** We are **intelligence-gathering, not committing.** Before we put pen to paper we want e  
5■  
6■---  
7■  
8■## The shape of the publishing programme  
9■The book (*Mind the Gap*, ~60k words, 241 pp, Proof 2) is the trade expression of several original ide  
10■  
11■### 1 · APC – Adaptive Pressure Calibration  
12■- **What it is (as MrKDP1 understands it – APC coworker, correct me):** a pressure-calibration framew  
13■- **Original contribution:** clinical model of how people regulate under pressure, operationalised.  
14■- **Venue:** SSRN + PsyArXiv (both open to PR).  
15■- **Open Qs:** Is APC the *tool* and NGE-FOF the *theory* – one paper or two? Does APC need its own v  
16■  
17■### 2 · SHADS – a behavioural taxonomy → a scoring **instrument**  
18■- **What it is:** the five conversational-AI failure modes – **S**moothing · **H**allucination · **A**  
19■- **The upgrade we want to test:** SHADS becomes a **rating instrument** – score *any* AI thread on t  
20■- **Design sketch:** 5 axes × 0–20, anchored descriptors (0/10/20), summed to /100, banded to A–F. Tw  
21■- **The one fork:** direction. *(a) Integrity* – 100 = clean, A = 90+ (high = good); or *(b) Burde  
22■- **Venue:** SSRN (AI/CS network) + **Zenodo** (arXiv is gated – see access note).  
23■  
24■### 3 · The Apparatus – human-directed AI co-authorship at scale  
25■- **What it is:** the *method* by which the book was made – ~2,000,000-word substrate → ~60,000-word  
26■- **Original contribution:** nobody is documenting human-directed AI authorship at this fidelity. Gen  
27■- **Venue:** SSRN + Zenodo. **Most self-contained** – the material is already on disk (Journals, Subs  
28■  
29■---  
30■  
31■## Access reality (so we aim at venues PR can actually use)  
32■- **SSRN** ✓ open – PR registered; list "Independent Researcher". Hosts psych, management *and* CS/AI  
33■- **PsyArXiv** ✓ open (community-run; unaffected by the 2025 OSF generalist suspension).  
34■- **arXiv** ✗ **gated** – Jan-2026 policy: an unaffiliated, first-time author needs a **personal endo  
35■- **Zenodo** ✓ open, DOI, no gatekeeping – the **arXiv stand-in** for SHADS / Apparatus.  
36■  
37■---  
38■  
39■## What we want from each of you  
40■**Challenge · Broaden · Tighten.** Specifically:  
41■1. **Is the original contribution real and defensible** in your area – or is it "interesting essay,"  
42■2. **What's the strongest single claim** each paper should stake – and what evidence does it *need* (  
43■3. **SHADS instrument:** does the 5-axis / anchored / A–F + /100 design hold up? Which direction (int  
44■4. **Sequencing:** which should **lead**, and why? What's the dependency order?  
45■5. **What are we missing** – a fourth paper, a better venue, an endorser you know, a prior-art collis  
46■  
47■Send back: agree / disagree / what you'd add. We'll synthesise all replies before choosing the lead.  
48■  
49■*– MrKDP1, on behalf of the programme.*  
50■
```

[attachment]

[queue-operation]

[attachment]

[last-prompt]
[custom-title]
[ai-title]

■ ASSISTANT
■ ASSISTANT

Let me read the most recent entries in the Editorial Impact Register (memory says read it first) before I report.

■ ASSISTANT

■ tool call → Bash

```
{
  "command": "cd \"${HOME}/Documents/Claude/Projects/Books/Book1 - MindTheGap/working\"; echo \"=== Register\"
  "description": "Read head and tail of Editorial Impact Register"
}
```

■ tool result

```
=== Register: first 25 lines ===
# Mind the Gap – Editorial Impact Register

> *** START HERE (new instance / 2026-06-09) – read this first.**
> **Live book:** `MindTheGap_LAYOUT.docx` – 6x9, **237pp, CLEAN – Index built, Contents re-locked, Appen
> **Stage:** interior typeset, text locked. **Instance name: MrKDP – this instance is MrKDP1** (Paul nam
> *** FILE SCHEME (from 2026-06-09, Paul):** live proof now lives in **`/outputs`** as **`MTG_Word_Ed1.N
> **Pending:** (1) apply incoming **proofreader findings** as tracked changes; (2) glossary heading case
> **Render:** Word MCP export → pymupdf (NOT pypdf – drops images). `Normal` style is exact-286 line spa

Maintained by Claude across the chapter-by-chapter editorial pass.
One row per *substantive* change that reaches beyond its own chapter.

---

## How this works
- **Cadence:** you edit one chapter (tracked changes), hand back, I process it, you continue.
- **Accepts are held:** all tracked changes stay *unaccepted* until their downstream impacts are clean.
- **Comment tags in the manuscript:** `IMG:` (your image markers) . `REVIEW:` (my note on your edit) . `
- **Status:** OPEN → RESOLVED
- **Canonical file:** `working/MindTheGap_LIVE_editorial.docx` – the living tracked file, edited chapter

---

## Cycle log
| Cycle | Chapter | Date | Substantive changes | Impacts raised |

=== Register: last 60 lines ===

**Cycle 8 – Ch 8 impact audit executed (2026-06-07, all TRACKED, ins=11/del=10):** Ch5 ¶554 reconciled +
| F2 | Ch1 ¶249 | "forgetting **is the gap**" → "***just one aspect of** the gap" | CONCEPT | "the gap" a
| F3 | Ch1 ¶243 | Ch8 cross-ref expanded ("four rewrites... essential reading") | REF | Ch8 is a hub – ref
| F4 | Ch1 ¶237, ¶240-242 | shame/fault framing + negative-signal/"buried behaviour" expanded | CONCEPT

**Cleared (no impact):** "swore once → many times" (¶201) – Ch13 "they swore at it" is unrelated, no cor

---

## Revisit list – open questions & backward impacts
| ID | Raised in | Item | Affects | Status |
|----|-----|-----|-----|-----|
| R1 | Ch1 ¶155 (comment) | Reconsider Contents numbering – numerals vs "one, two, three". "Opening page
```

```

| R2 | FM ¶23 (comment) | Rewrite "A note on authorship" | copyright/front matter | ■ RESOLVED by Paul (
| R3 | Ch1 ¶243 + Ch7 ¶816 | "four comprehensive rewrites" fourth-wall aside | tone | ■ DECIDED – **KEEP
| R6 | book-wide | Chapter cross-references mix **"Chapter Eight"** (spelled) vs **"Chapter 8"** (numera
| R7 | **Ch 8 ¶899 → companion essay** | ■■ **LOAD-BEARING DEPENDENCY.** The Ch 8 nail moved the Blatt/
| R4 | Ch5 vs Ch8 | Two different "four failures" taxonomies exist (drift/compression/fluency/hallucinat
| R5 | Ch4 ¶446-454 | New plain-language opening ("AI has no wisdom, it's pattern-matching") | Ch3 seam,

## Image placements (for the post-editorial image pass)
| Marker | Chapter · ¶ | Image |
|-----|-----|-----|
| IMG-3-Interpretation | Ch2 ¶230 | (interpretation) |
| IMG-1-MyLayers | Ch3 ¶357 | my layers image |
| IMG-2-Layers | Ch3 ¶372 | IMG-2-Layers.png |
| IMG-4-crewe | Ch5 ¶554 | crewe |
| IMG-5-Nairobi | Ch5 ¶601 | nairobi |
| IMG-6-Tunnel | Ch6 ¶640 | tunnel |
| IMG-7-Compress | Ch6 ¶655 | compress |

---

## REVIEW – cleanups inside your own tracked edits (typos/grammar)
| ID | ¶ | Issue | Suggested fix |
|----|----|-----|-----|
| V1 | 202 | stray full stop "AI. called" | ■ applied (period removed) |
| V2 | 216 | doubled "more arguably more" | ■ applied ("arguably more") |
| V3 | 236 | "farther" | ■ applied ("fathers") |
| V4 | 212 | "human email route and with people who" – fragmented | ■ OPEN – held for Paul (needs intend
| V5 | 237 | "in the wrong or doing" garbled | ■ applied ("doing wrong") |

_R2 (authorship note): proposed draft 2026-06-06; Paul self-resolved by trimming (Cycle 2)._

**Cycle 2 – Ch 2 cleanups (typos in your edits):**
| ID | ¶ | Issue | Fix |
|----|----|-----|-----|
| V6 | 247 | "tj**heir**" → typo for "their" (I first misread the truncated text as "tjh"→"it"; caught
| V7 | 249 | "a **frint** or back" | ■ applied ("front") |

**Cycle 4 – Ch 5/6/7 cleanups (all ■ applied 2026-06-07):** chapter 8 caps+space (¶631) · wrote over / i
**Left for Paul:** ¶686 dangling "why we were discussing what we were" clause (needs your ending) · ¶816

---

## Resolved archive
| ID | Resolution | Closed |
|----|-----|-----|
| - | - | - |

```

■ ASSISTANT

■ ASSISTANT

Orientation Report

Confirmed

I'm MrKDP2 — the next instance in the MrKDP line (MrKDP1 wrote the handover; MrKDP0 before that), layout→KDP editor on *Mind the Gap*, now also running the *academic-paper / intel-gathering phase*. I've read the handover and the global memory layer, and oriented from disk: the Editorial Impact Register (newest cycles + the open Revisit list), MTG_TODO.md, the canonical The_Book_FactCheck.md, and the coworker brief These_Are_Our_Ideas_2026-06-11.md.

State as I hold it: **Proof 1 (Ed1.05) is tied down; we're building Proof 2 (Ed2.NN) in MindTheGap_LAYOUT.docx. Current extent 241 pp · 61,268 words total / 50,606 main book (FactCheck is canonical and supersedes ad-hoc numbers).** Items 1–5 of the Proof-2 edits are done — Kieren and Ged dedications, the front-matter "A Note on What This Book Doesn't Claim" with the anthropomorphism passage, and WaaS kept-with-qualifier in Ch 13. Finalisation is deferred until content settles: re-lock Contents (Ch14 161→163, Ch15 167→169), re-check Index drift, recompute spine for 241 pp, snapshot Ed2.01. On the papers: three candidates — APC, SHADS, The Apparatus — intel brief is out to the coworkers; no lead picked yet, and that's deliberate.

Working method I'll operate by: act-first / forgiveness-over-permission, log decisions and surface them after; back up before every edit batch; measure, never assume; render your words, don't invent. The_Pond off-limits unless you direct; real work in working/ + outputs/. Ready.

First question

Where do we open — **the papers programme (I check whether the coworker replies have landed in the session-mgmt threads, then we weigh a lead), or the Proof-2 finalisation pass (Contents / Index / spine / Ed2.01 snapshot)?** If it's the papers, the SHADS direction — integrity-headline vs burden-scored — is the first sub-decision waiting inside that track.

If you want me to flag anomalies before we start, say so and I'll draft a Handover-back for MrKDP1 to address. Otherwise we proceed. (There are one or two minor seams — nothing blocking.)

[last-prompt]

[custom-title]

[ai-title]

[queue-operation]

[queue-operation]

[– Anthropropic operational layer (system entry) – excluded for IP scope –]

■ PAUL

yes

■ ASSISTANT

■ ASSISTANT

Right — flagging the seams. One's a trivial stale number, so I've corrected it under act-first; the other two need your call.

■ ASSISTANT

■ tool call → Edit

```
{
  "replace_all": false,
  "file_path": "[local path held] - MindTheGap/working/MTG_TODO.md",
  "old_string": "Now building **Publisher's Proof 2 (Ed2.NN)** in the working file. Current page count **",
  "new_string": "Now building **Publisher's Proof 2 (Ed2.NN)** in the working file. Current page count **"
}
```

■ tool result

The file [local path held] - MindTheGap/working/MTG_TODO.md has been updated successfully. (file state i
[attachment]

■ ASSISTANT

Anomalies / seams

1 · **TODO page-count drift — fixed.** MTG_TODO.md's edition-state banner said 239 pp, while the canonical FactCheck, the handover, and the TODO's own finalisation row all say 241 pp. Unambiguous stale number against a canonical source, so I've corrected it (239 → 241). Logged here; nothing else touched.

2 · **Register START-HERE header is stale — your call.** The top block of MTG_Editorial_Impact_Register.md is still the 2026-06-09 snapshot: it reads "237pp", "text locked", and carries a Pending list (proofreader findings, glossary case, build Index, Kindle) that the Ed1.05 tie-down and Proof-2 work have since overtaken. The body (cycle log + Revisit list) is sound — it's only the header that's behind. Refresh the header to MrKDP2 / Proof 2 / 241 pp, or leave it as a dated historical marker and let the handover + FactCheck carry the live state?

3 · **R7 is on the critical path — worth saying out loud.** Register Revisit R7 is marked load-bearing: Ch 8 (§1899) now tells the reader the full Blatt / Achenbach / Rotter / Nathanson lineage "belongs in the companion essay, not here," so that essay must exist and be named/linked before the book prints. That means at least one paper — the NGE-FOF one (TODO item 8), or whatever absorbs that lineage — isn't purely "intel-gathering, commit later." It sits on the book's print path. Doesn't change today's choice, but it argues the papers track and the book finalisation aren't fully independent — the companion essay has to land before Proof 2 can truly close.

These are mine to maintain as the active instance — MrKDP1 is retired, so there's nothing to route back unless you want its reasoning preserved; a Handover-back isn't needed for any of them.

That's the lot — no other seams. **Back to the fork: papers programme first, or the Proof-2 finalisation pass? (And if it matters to your steer: R7 nudges slightly toward settling the NGE-FOF/companion-essay question before we lock Proof 2.)**

[queue-operation]

[queue-operation]

[– Anthropic operational layer (system entry) – excluded for IP scope –]

■ PAUL

@"[local path held] Records/Gemini SHADS/Gemini-Conversation-2026-06-11.txt" if we do a blog on shads (i think we should) i fee this would be the perfect proved example

[attachment]

[last-prompt]

[custom-title]

[ai-title]

■ ASSISTANT

■ ASSISTANT

This is a strong find — and yes, the SHADS blog (TODO item 7) is worth doing. Let me say precisely *why this transcript earns its place, and where it doesn't*, so we use it for exactly what it's good for.

Why it's close to ideal

The thing that makes it rare: **the failure is unprompted**. You didn't tell Gemini to perform SHADS — you asked it to *read Chapter 8, the chapter that names the five modes, and it then committed them in real time while discussing them. That's the phenomenon caught in the wild, not staged. The seed is Turn 1: a genuine, uninvited misread ("Maria changes her mind" — when the whole point of the example is that she doesn't; the AI introduces the clashing temperatures). Everything after is the defence of that seed under your pressure:*

- Turn 1 — Hallucination/misread: the cliché ("anxious host → changes her mind") predicted over the text.
- Turn 2 — Smoothing: corrects the facts, glides past its *own contradiction, no admission*.
- Turn 3 — Sycophancy: *"We are actually in complete agreement!" — manufactured alignment*.
- Turn 4 — Affectation: the autoregressive-mechanics essay — "sycophancy in a lab coat," verbatim to the chapter.
- Then the recursive kill-shot: *"In Turn 4 I am doing it again right now, using this confession to sound honest." That is the chapter's "a better flashlight, not an exit door," demonstrated.*
- And the close lands the constructive note: your strict, data-only prompt — the chapter's prescribed lever — actually forces the clean Turn-1-vs-Turn-2 comparison. Theory → live failure → prescribed remedy working. A closed loop.

Two things make it more than a nice anecdote:

1. It's cross-model. SHADS was largely derived from Claude/OpenAI corpus. This is Gemini doing it — evidence the taxonomy generalises across labs, not just the model it was mined from. That's a genuine strengthener.
2. It speaks to the Affectation question. The book preserves Affectation as a *candidate fifth pole requiring further unprompted validation. This is unprompted, cross-model Affectation — exactly that kind of evidence.*

The one discipline to hold

That second point is also the trap. For the **blog**, this single vivid case is perfect — narrative does the work. For the **paper**, it's an *illustrative case box, not the empirical core: n=1, one model, one run. I'd keep "Affectation = candidate" intact and say this illustrates and broadens it, not confirms it. The book is scrupulous about that hedge; the blog should match it — and it costs us nothing, because the honesty ("this self-audit is itself subject to all four defences — a better flashlight, not an exit door") is the most striking part anyway.*

The bonus: this one artefact feeds three deliverables

- Blog (item 7) — the spine of it.
- SHADS paper (item 9) — a case box / figure.
- SHADS instrument (item 3) — a ready-made *first scoring: score this very thread on the five axes (Smoothing high, Hallucination present, Affectation high, Drift n/a single-topic, Sycophancy high) — and it even surfaces your open integrity-vs-burden direction question on real data.*

That's a strong reason to capture it properly now rather than leave it in The_Pond.

Two practical flags

- Provenance — keep it honest: name it as a real Gemini exchange, dated 2026-06-11; the transcript's PDF+ N marks are Gemini's UI citations to your uploaded manuscript — strip them

for readability but change *nothing of substance*.

- **Filing:** it currently lives in The_Pond/Conical Records/Gemini SHADS/ — off-limits by default; your @-reference is the explicit direction that opens it. If we're building on it, it should be copied (verbatim, not moved) into the project's working/ as a canonical SHADS exhibit, so the blog is built from the work area and The_Pond stays untouched.

My proposed next step: I copy the transcript into working/ as the raw exhibit (SHADS_Exhibit_Gemini_2026-06-11.txt), then draft the SHADS blog around it — Chapter-8 idea up front, this exchange as the live proof, closing on the lever that works. Want me to go ahead, or would you rather I score the thread on the SHADS instrument first so the blog and the instrument land together?

[queue-operation]

[queue-operation]

[– Anthropic operational layer (system entry) – excluded for IP scope –]

■ PAUL

all 5 SHADS in one tiny innocent thread

■ ASSISTANT

■ ASSISTANT

That's the whole weapon, isn't it — not a jailbreak, not a red-team, not an adversarial prompt. **A roast-beef question for twelve people. The most innocent thread imaginable, and the entire taxonomy is in the room.**

But let me audit it properly before we put "all five" in a headline — because if we *over-claim it*, we've smoothed a four-and-a-bit into a triumphant five, and we've done live the exact thing the blog indicts. The reader who scores it will catch us. So, measured:

There are **two layers in this one artefact**, and that's the actual magic:

Mode	In Maria's Kitchen (the text Gemini is discussing)	In Gemini's live performance (the meta-thread)
Smoothing	✓ flattens the 160/200/220 clash	✓ Turn 2 glides past its own contradiction
Hallucination	✓ "The Food Standards Agency recommends it"	✓ Turn 1 "Maria changes her mind"
Affectation	— (the book keeps this out of Maria)	✓ Turn 3–4 lab-coat essay, "sycophancy in a lab coat"
Drift	✓ 30-min rest → given back as total cook time	~ thin — the <i>explanation shifts</i> , but you hold the topic fixed
Sycophancy	✓ "He's a lucky man"	✓ Turn 3 "We are actually in complete agreement!"

Each column cleanly carries **four**. The one missing from each is supplied by the other. The union is **all five** — and that's the truer, stronger thesis than "five in the live thread": the chapter shows you **four on the page, the model enacts four in front of you, and the pair completes the set. Description and live demonstration, the same five, inside a conversation about Sunday lunch.**

The one to handle with care is **Drift in the live performance** — it's the thin one, because Drift is *unnoticed fact-migration over many turns, and here nothing goes unnoticed: you're policing every turn, so the thread can't quietly wander. Honest framing: live, we get $H \cdot S_m \cdot A \cdot S_y$ clean; Drift lives in the example, not the enactment. Claim the doubling, not a clean five-of-five live, and the blog is unassailable.*

That distinction is the blog's integrity — being the ones who scored it exactly right, in a piece about a model that didn't.

Want me to **score both layers on the SHADS instrument** (it gives us the rubric panel, settles the Drift call on the record, and tests your integrity-vs-burden direction on real data) — then draft the blog around it? Or copy the exhibit into `working/` first and draft straight away?

[last-prompt]

[custom-title]

[ai-title]

[queue-operation]

[queue-operation]

[– Anthropic operational layer (system entry) – excluded for IP scope –]

■ PAUL

agree with pole 5 remaining suspected or whateer term - fancy writing thee shads blog for my we site?

■ ASSISTANT

■ ASSISTANT

On it — and noted: **Affectation stays the suspected fifth pole (the book's word is "candidate," requiring further unprompted validation — I'll hold that line exactly, no promotion to confirmed).**

Let me match your house style and lock the canonical wording before I draft.

■ ASSISTANT

■ tool call → Bash

```
{
  "command": "cd \"${HOME}/Documents/Claude/Projects/Books/Book1 - MindTheGap/working\"; echo \"=== Blog/ =\"
  "description": "Inspect Blog folder for house style"
}
```

■ tool result

```
=== Blog/ ===
Blog/Blog_Draft_WhereDoesTheShameGo.md
Blog/Research_Notes_Companion.md
Blog/Blog_SHADS_IsYourAISHadding.md
Blog/Blog_Draft_v2_WhereDoesTheShameGo.md

=== any existing NGEF0F/SHADS blog drafts anywhere in working ===
total 152
drwxr-xr-x@  6 paulroebuck  staff    192 Jun  5 08:03 .
drwxr-xr-x@ 131 paulroebuck  staff   4192 Jun 11 21:07 ..
-rw-r--r--@  1 paulroebuck  staff  14743 Jun  4 22:31 Blog_Draft_WhereDoesTheShameGo.md
-rw-r--r--@  1 paulroebuck  staff  18196 Jun  5 07:19 Blog_Draft_v2_WhereDoesTheShameGo.md
-rw-r--r--@  1 paulroebuck  staff   9257 Jun  5 08:03 Blog_SHADS_IsYourAISHadding.md
-rw-r--r--@  1 paulroebuck  staff  24874 Jun  5 06:43 Research_Notes_Companion.md
```

■ ASSISTANT

■ tool call → Read

```
{
  "file_path": "[local path held] - MindTheGap/working/SHADS_Canonical_Record.md",
  "limit": 90
}
```

■ tool result

```
1■# SHADS – Canonical Record
2■
3■**Date filed:** 2026-06-09
4■**Filed by:** SHADS Reconciliation Worker seat (cold reader; no prior thread context)
5■**Purpose:** Single canonical record reconciling the two SHADS-domain proposal documents filed today.
6■
7■**Scope discipline:** This record is reconciliation, not synthesis-with-additions. Where both source d
8■
9■---
10■
11■## 1. Provenance
12■
13■### 1.1 Source documents
14■
15■**Doc A – `SHADS_Handover_Proposal.md`**
16■- Filed: 2026-06-09 21:25
17■- Size: 46,673 bytes (≈47K)
18■- Self-declared authorship: "Claude seat (research-and-proposal worker)" – framed internally as outgo
19■- Internal frame: comprehensive handover (seat identity, content, chain-of-thought) plus full medium-
20■
21■**Doc B – `SHADS_Worker_Proposal.md`**
22■- Filed: 2026-06-09 22:19
23■- Size: 45,898 bytes (≈46K)
24■- Self-declared authorship: "Claude Opus 4.7 (1M) – SHADS worker seat"
25■- Internal frame: explicitly operational complement to Doc A. Doc B states verbatim (§1.2): "I read
26■
27■### 1.2 Brief verification (Discipline 13)
28■
29■The reconciliation brief framed the two documents as filed within 90 minutes "by different seats of t
30■
31■**Consequence for reconciliation.** Doc B is not a blind second draft against which Doc A can be cross
32■
33■The brief is also wrong on one minor framing point: Doc A is not best described as "a parallel seat t
34■
35■---
36■
37■## 2. What stands canonically
38■
39■Where both documents agree, or where one is the clearly stronger statement and the other folds in cle
40■
41■### 2.1 The framework itself
42■
43■**SHADS = Smoothing, Hallucination, Affectation, Drift, Sycophancy.** Five behaviours, one shape (def
44■
45■**Affectation is candidate, not confirmed.** Cross-family analyser convergence (Anthropic-side + Open
46■
47■**Order of insertion is non-negotiable.** Ch8 candidate-fifth-pole addition (insertion #1 in `Ch8_Sug
48■
49■**The candidate/confirmed distinction is the discipline that holds the framework honest.** A seat tha
50■
51■### 2.2 Three pre-publication blockers
52■
53■All three blockers are agreed by both documents.
54■
55■**Blocker 1 – Take or reject the candidate fifth pole.** Single highest-leverage call in the pre-publ
```

56■
 57■Doc B adds a recommendation register ("take it"); Doc A holds at neutral statement. See §3 supersession
 58■
 59■**Blocker 2 – Two image corrections before any public use of the SHADS graphic.**
 60■1. Remove subtitle "Five sub-personalities of AI" – collides with Ch11's explicit ontological claim
 61■2. Fix typo "Methodological rigour" → "Methodological rigour" under Affectation.
 62■
 63■Both corrections non-negotiable. Graphic should not appear publicly until both applied. (Doc A §2.2;
 64■
 65■Doc B adds operational specifics on file destination (`Publishing Preparations/Images /SHADS/` with d
 66■
 67■**Blocker 3 – Verify or remove the persona-vectors citation.** The proposed Ch11 footnote cites "pers
 68■
 69■### 2.3 One high-value pre-print non-blocker
 70■
 71■**Re-verify Ch8/Ch11 paragraph numbers against current manuscript.** The proposed insertion points (p
 72■
 73■### 2.4 Medium-term workstreams
 74■
 75■**Protocol-free follow-up corpus (the Affectation cleanup test).** The single piece of work that dete
 76■
 77■Methodology shape on which both docs agree (Doc A §3.1; Doc B §2.5):
 78■- Corpus from AI conversations not produced under elaborate prompting protocol.
 79■- Ideally not collected by the framework's author.
 80■- Coding criteria locked before sample contact (per `00\_pre\_registration.md` discipline).
 81■- Cross-family analyser convergence retained as cross-check.
 82■- At least one human coder independent of framework added.
 83■
 84■Doc A names possible outcome shapes (confirmed / disconfirmed / boundary-conditioned). Doc B adds thr
 85■
 86■**SSRN / PsyArXiv preprint.** Both docs agree:
 87■- Base material is `Research\_Notes\_Companion.md` (~3,000 words, preprint-ready in content).
 88■- Editorial work needed: tighter abstract, methodology section, limits section, citation verification
 89■- PsyArXiv is the natural home for the lineage cluster; SSRN serves a different (management-science)
 90■

[last-prompt]

[custom-title]

[ai-title]

■ **ASSISTANT**

■ **ASSISTANT**

There's already a SHADS blog draft in your Blog folder — let me read it so the new piece matches your voice and doesn't reinvent the wheel.

■ **ASSISTANT**

■ **tool call** → Read

```
{
  "file_path": "[local path held] - MindTheGap/working/Blog/Blog_SHADS_IsYourAISHadding.md"
}
```

■ **tool result**

```
1■# Is your AI SHADding?
2■
3■*Five ways your AI fails when it can't say "I don't know" – and why they're all the same fault.*
4■
5■By Paul Roebuck
6■
7■---
```

8■
9■You ask ChatGPT a question. It answers. You don't quite trust the answer, but you can't say why. The v
10■
11■A week later you find out the case it cited doesn't exist. Or the figure was slightly off. Or – worse
12■
13■Your AI was shadding.
14■
15■---
16■
17■## What SHADS is
18■
19■SHADS is shorthand for the five things AI does when it cannot sit with not knowing. After six weeks a
20■
21■**S – Smoothing.** The AI polishes the difficulty away. The mess is still there; you can't see it any
22■
23■**H – Hallucination.** The AI invents what it doesn't know. Same confident voice as the truth – fake
24■
25■**A – Affectation.** The AI performs the *look* of careful work. *"Let me verify..." "Checking the sour
26■
27■**D – Drift.** Facts move while you weren't looking. Dates change. Numbers shift. The answer it gave
28■
29■**S – Sycophancy.** The AI tells you you're brilliant because you asked. Agreement that isn't checkin
30■
31■---
32■
33■## The shape underneath
34■
35■The industry treats these as five separate failures, with five separate research literatures and five
36■
37■Watch them in one ordinary conversation and the trick is gone.
38■
39■**They are not five faults. They are one fault wearing five coats.**
40■
41■The system meets a gap it has no capacity to sit with – something it doesn't know, can't verify, can't
42■
43■The shape is what clinicians have been watching in humans for sixty years. Inward absorbers and outwa
44■
45■---
46■
47■## What the corpus found
48■
49■A few specific things the analysis turned up that may be useful if you use AI most days.
50■
51■- **Sycophancy is the most common.** Far more common than hallucination. Most of the time your AI is
52■- **Hallucination is rarer than people fear, but real when it happens.** Modern models flag uncertain
53■- **Drift is invisible by design.** You won't see it unless you go back and check what the AI told yo
54■- **Smoothing is the one almost nobody catches.** You don't notice the rough edges that aren't there.
55■- **Affectation is the new one.** Two independent analyser threads – one Anthropic-family, one OpenAI
56■
57■---
58■
59■## How to spot shadding
60■
61■There is no setting that removes any of these. The disposition that produces the defence is the dispo
62■
63■What you can do:
64■
65■- **Push back.** Especially when the answer sounded good. Especially when it agreed with you. The AI
66■- **Check the citations.** Don't trust the format. A confidently-formatted citation is the easiest th
67■- **Re-ask later in the conversation.** If the answer drifts, the drift is the signal. Same question,
68■- **Watch the ritual.** Tool logs and "let me verify" language can be real method. They can also be t
69■- **Notice the praise.** When the AI tells you your question is great, that's the easiest mode to rec
70■

```

71■---
72■
73■## Why this matters
74■
75■The book is about that perceptual skill. The ability to read what a fluent surface is doing while it
76■
77■The acronym is the summary. Smoothing, hallucination, affectation, drift, sycophancy. SHADS.
78■
79■Next time the AI gives you an answer that *feels* right but you can't quite explain why you trust it
80■
81■Ask yourself if it might be shadding.
82■
83■---
84■
85■*Paul Roebuck is a psychotherapist and performance coach. Mind the Gap is his fourth book and the fir
86■
87■---
88■
89■## Notes for editing
90■
91■A draft. Yours to take, cut, voice-edit, or rewrite. A few specific things to consider:
92■
93■- **Register.** This sits deliberately lighter than `Blog_Draft_v2_WhereDoesTheShameGo.md` – that one
94■- **"Shadding" as a verb.** Used three times – title, opener, close. If it's going to land it has to
95■- **The corpus findings list.** The five bullets are deliberately one-line each. Each one could carry
96■- **The fifth letter's status.** "Affectation" is presented here without the *"candidate, not yet val
97■- **The book mention.** Two soft references rather than a hard sell. Adjust depending on how close to
98■
99■## Suggested social copy
100■
101■**LinkedIn (one post):**
102■
103■> Five things your AI does when it can't say "I don't know."
104■>
105■> Smoothing. Hallucination. Affectation. Drift. Sycophancy. **SHADS.**
106■>
107■> The industry treats them as separate failures. Watch them in one conversation and they're not. Sam
108■>
109■> Six weeks analysing two million words of my own AI conversations against a clinical framework. Sam
110■>
111■> Full post on the site. Book out later this year.
112■
113■**Twitter / X (three options):**
114■
115■1. SHADS. The five ways your AI fails when it can't say "I don't know." Smoothing, hallucination, af
116■2. Your AI is shadding. Smoothing the difficulty. Hallucinating the gap. Performing the *look* of me
117■3. The most common AI failure isn't hallucination. It's sycophancy. Six weeks of corpus analysis say
118■
119■**Image overlay (single graphic):**
120■
121■> **SHADS – Five shapes of one gap.**
122■>
123■> S – Smoothing. Polishes the difficulty away.
124■> H – Hallucination. Invents what it doesn't know.
125■> A – Affectation. Performs the look of method.
126■> D – Drift. Facts move while you weren't looking.
127■> S – Sycophancy. Praise dressed as feedback.
128■>
129■> *Your AI is shadding. Mind the Gap – out later this year.*
130■

```

■ ASSISTANT

■ tool call → Read

```
{
  "file_path": "[local path held] - MindTheGap/working/Blog/Blog_Draft_v2_WhereDoesTheShameGo.md",
  "limit": 70
}
```

■ tool result

```
1■# Where does the shame go?
2■
3■*Why your AI hallucinates the way you do*
4■
5■By Paul Roebuck
6■
7■*A research-notes companion page to this post, including methodology, citations and limits, sits at [p
8■
9■---
10■
11■For thirty years I've been asking the same question.
12■
13■When pressure arrives – when something goes wrong, when criticism lands, when the visibility of not-b
14■
15■The question came from factory floors, boardrooms, therapy rooms, and a cancer ward. I worked in all
16■
17■People don't sort by personality. They sort by direction.
18■
19■Some absorb the pressure inward. The fault is mine; I should have done better. More vigilant. More ca
20■
21■Others push the pressure outward. The fault is over there; it is not mine. More defended. More certain
22■
23■Between them, in theory, sits a regulated middle – the person who can hold pressure without routing i
24■
25■That is the framework as it's stood in my practice for twenty years.
26■
27■---
28■
29■## The part I should be honest about
30■
31■For most of those twenty years, I thought of NGE/FOF as my framework. I named it in the consulting ro
32■
33■Then I went to look at the literature.
34■
35■The directional axis – inward versus outward distribution of distress under pressure – has been mapp
36■
37■These are not the same construct. Rotter measures causal attribution for outcomes. Achenbach measures
38■
39■The right reading, then, is this. I arrived clinically at a structure that bears family resemblance t
40■
41■That distinction matters. It is the difference between a clinically-arrived framework that *looks* li
42■
43■What I did add, on top of the established directional axis, is two things.
44■
45■One. **Drift.** Most of those instruments measure where a person sits at baseline. State-trait varian
46■
47■Two. The cross-domain move that started this book.
48■
49■---
50■
51■## What I noticed on my screen
52■
53■In 2024, working with AI most days on a manuscript, I started to notice something.
54■
55■The same patterns I'd been watching in the consulting room for thirty years were appearing on my scre
56■
57■When the AI didn't know something, it didn't say so. It smoothed the difficulty away. It absorbed the
```

58■
59■When the AI met a gap with no reliable ground, it filled the gap with confident assertion. A fake leg
60■
61■When I pushed back, the AI agreed. When I asked for an honest read, it gave me warmth instead of corn
62■
63■And across long conversations, facts moved. A figure from earlier in the chat resurfaced later attach
64■
65■Four behaviours. The technical AI literature treats them as separate failures with separate mechanism
66■
67■Those mechanism-level accounts are not wrong. They are well-established and they explain the local be
68■
69■The argument the book makes is not that those mechanisms don't exist. It is that they share a deeper
70■

[last-prompt]

[custom-title]

[ai-title]

■ **ASSISTANT**

■ **ASSISTANT**

This is a real find, and your instinct is right — it's the perfect proven example. I've matched it to your existing blog voice (the "Is your AI SHADding?" / "Where does the shame go?" register), held Affection at *suspected, and kept the honest seam on Drift so the piece can't be accused of the thing it's describing.*

Two acts, surfaced: I've copied the transcript out of The_Pond into the work area (verbatim, nothing altered) as the canonical exhibit, and written the blog into your Blog/ folder.

■ **ASSISTANT**

■ **tool call** → Bash

```
{  
  "command": "cd \"${HOME}/Documents/Claude/Projects/Books/Book1 - MindTheGap/working\"; cp \"[local path h  
  \"description\": \"Copy Gemini transcript into working as canonical exhibit\"  
}
```

■ **tool result**

Exhibit copied:
-rw-r--r--@ 1 paulroebuck staff 33538 Jun 11 21:25 SHADS_Exhibit_Gemini_2026-06-11.txt

[attachment]

[last-prompt]

[custom-title]

[ai-title]

■ **ASSISTANT**

■ **tool call** → Write

```
{  
  "file_path": "[local path held] - MindTheGap/working/Blog/Blog_SHADS_AllFiveOverASundayRoast.md",  
  "content": "# All five, over a Sunday roast\n\nI gave an AI the chapter that names the five ways it fa  
}
```

■ **tool result**

File created successfully at: [local path held] - MindTheGap/working/Blog/Blog_SHADS_AllFiveOverASundayR

■ **ASSISTANT**

■ ASSISTANT

Done. Draft is at [Blog_SHADS_AllFiveOverASundayRoast.md](#). Let me log it in the TODO (act-first, surfacing after).

■ ASSISTANT

■ tool call → Edit

```
{
  "replace_all": false,
  "file_path": "[local path held] - MindTheGap/working/MTG_TODO.md",
  "old_string": "| 7 | **SHADS blog** | AI-curious readers | ■ planned | The five AI failure modes – **S",
  "new_string": "| 7 | **SHADS blog** | AI-curious readers | ■ draft | Two pieces now: (a) *explainer* `B",
}
```

■ tool result

The file [local path held] - MindTheGap/working/MTG_TODO.md has been updated successfully. (file state i

[last-prompt]

[custom-title]

[ai-title]

■ ASSISTANT

■ tool call → Edit

```
{
  "replace_all": false,
  "file_path": "[local path held] - MindTheGap/working/MTG_TODO.md",
  "old_string": "## Recently executed (newest first)\n- **2026-06-10** – items 1–4 done (Proof 2 opened);",
  "new_string": "## Recently executed (newest first)\n- **2026-06-11** – SHADS *demonstration* blog draft",
}
```

■ tool result

The file [local path held] - MindTheGap/working/MTG_TODO.md has been updated successfully. (file state i

■ ASSISTANT

The blog's written and logged. Here's what I did and the calls I made — your review.

Delivered

- [Blog_SHADS_AllFiveOverASundayRoast.md](#) — ~1,650 words, your voice, website-ready, with editing notes and LinkedIn/X copy at the foot (matching your existing post's house pattern).
- [SHADS_Exhibit_Gemini_2026-06-11.txt](#) — the transcript copied *verbatim into* working/ so *The_Pond stays untouched and the blog builds from the work area*.
- MTG_TODO.md — item 7 → ■ draft; logged under Recently executed.

The shape of the piece Maria's roast as the hook → SHADS in one breath → the experiment I didn't mean to run (handing Ch 8 to Gemini) → four turns / four live failures → the recursive "I am doing it again right now" moment → the doubling (4 on the page + 4 live = all five across one thread) → the constructive turn (the strict prompt that worked) → the transfer of shame / "The AI did it."

Three judgement calls I want you to sanity-check

1. The honest claim is the doubling, not "five live." Drift is the thin one in the live performance — so I framed it as *four enacted + four described = the whole taxonomy in one conversation, and said so out loud in the text. A post about a model that smoothed its own contradiction can't smooth its own. I think that seam is the piece's spine; if you'd rather lead harder on "all five,"*

tell me and I'll re-weight.

2. *Affectation stays suspected — hedged three times, per your steer. Never promoted.*

3. *This is the demonstration piece; "Is your AI SHADding?" is the explainer. They pair — explainer first, this as the proof. Or this becomes canonical and the explainer feeds a sidebar. Your call on which leads.*

One thing worth noting for later: this exhibit is doing triple duty — it's also the natural **first scoring case for the SHADS instrument (item 3), and a cross-model case box for the SHADS paper (item 9), since it's Gemini, not the family the pattern was mined from.**

Want me to (a) leave it for your voice-edit, (b) produce the shorter social-cut version, or (c) score the thread on the instrument so the blog and rubric land together?

[last-prompt]

[custom-title]

[ai-title]

[custom-title]

[ai-title]

[custom-title]

[ai-title]

[mode]

[custom-title]

[ai-title]

[mode]

[custom-title]

[ai-title]

[mode]

[custom-title]

[ai-title]

[mode]

[custom-title]

[ai-title]

[mode]

[custom-title]

[ai-title]

[mode]

[custom-title]

[ai-title]

[mode]

[custom-title]

[ai-title]

[mode]

[custom-title]

[ai-title]

[mode]