

Mind the Gap — Master Statistics Dossier

Prepared by The Forge — Substrate_Statistics_Master_2026-06-01_1532 · 2026-06-20 · SHADS CW channel

Every substantive statistic from the substrate analysis in one place. Calibrated against the substrate's final manuscript baseline (V77, ~46,129 words). Founding-draft figure (~27,000) preserved alongside.

1. Headlines

One author. Twenty days. Two AI platforms. 17 chats. ~46,000 words of book. ~46,000 words of practitioner-substrate. ~2 million words of AI.

Headline	Figure
Calendar window	20 days (10–29 May 2026)
Active days	16 (4 zero-activity: 14, 15, 16, 19 May)
Substrate chats	17 (9 Claude editorial + 8 ChatGPT side)
Total substrate words	2,019,402
Total substrate turns	2,119
Final manuscript (V77 ≡ Rev 2 v0.0)	~46,129 words · 15 chapters
Founding-draft manuscript (pre-substrate)	~27,000 words
Manuscript growth across substrate	+~19,000 words (~70%)
Paul-words across substrate	45,785
AI-words across substrate (Claude + ChatGPT)	1,973,617
AI:Paul ratio	43.1 : 1
AI:Final-Manuscript ratio	42.8 : 1
Paul:Final-Manuscript ratio	0.99 : 1
Named AI working-personas	6 (Book Man → Son Of Bookman → Jose ×5)
Cross-instance passons / compactions	5–6
Manuscript versions surfaced	27 (V0 → V77)

The shape of the leverage: Paul small. AI large. Book Paul-sized.

2. The Leverage Matrix (calibrated against V77, 46,129 words)

Layer	Volume	× Final manuscript (46,129)	× Founding draft (27,000)	× Paul (45,785)	Per active day (÷16)
Final manuscript (V77)	46,129	1.00×	1.71×	1.01×	—
Founding draft (pre-substrate)	~27,000	0.59×	1.00×	0.59×	—
Paul's hand (substrate)	45,785	0.99×	1.70×	1.00×	2,862

Layer	Volume	× Final manuscript (46,129)	× Founding draft (27,000)	× Paul (45,785)	Per active day (÷16)
Claude prose-only	543,307	11.8×	20.1×	11.9×	33,957
ChatGPT (side-channel)	12,137	0.26×	0.45×	0.27×	758
Claude full (incl. tool-calls)	1,961,480	42.5×	72.6×	42.8×	122,593
Total substrate (Paul + Claude + ChatGPT)	2,019,402	43.8×	74.8×	44.1×	126,213

Key reads:

- For every word that shipped (V77), **Paul wrote ~1 substrate-word and the AI produced ~43.**
- For every word that shipped, **total conversation cost was ~44 words of substrate.**
- The manuscript grew **+19,000 words across the substrate window (V9 → V77).**
- The 27,000-vs-46,129 distinction matters: matrix figures look very different against the wrong denominator. Earlier "73:1" headline ratios used founding-draft; corrected figure is 42.8:1.

3. Manuscript Growth Trajectory

Version	Words	When	Note
V0 founding	~27,000	pre-substrate	36-hour first draft (per Book Man legacy reflection)
V9	~27,000	24 May (Jose1 start)	Substrate begins at this baseline
V17	34,157	24 May 18:21	First major growth
V29	39,640	25 May 14:01	Three surgical insertions to Ch 9
V36	43,132	26 May 14:46	Pre-Jose2 handover
V40	42,887	27 May 01:24	Ch 12 Beat 3 cut
V41	43,998	27 May 01:35	Ch 15 Summing Up v5
V42	45,556	27 May 10:48	Book Man retirement baseline · Jose3 baseline
V62	45,768	27 May 23:53	
V68	45,716	28 May 15:45	Ch 6 three-pass rebuild complete
V69	45,716	28 May 15:56	Jose4 "live baseline" declaration
V71	45,636	28 May 22:21	Contents page + Stanford correction

Version	Words	When	Note
V73	45,975	29 May 12:27	Two Dads / Bill Ollis scene added
V75	46,001	29 May 13:23	
V76	46,129	29 May 14:19	
V77 ≡ Rev 2 v0.0	~46,126–46,129	29 May 15:59	Vnn counter retires; "Rev 2 v0.0 ≡ V77 + your hand-back"

Growth pattern: Manuscript flat through Book Man chat (May 10–27 worked the founding draft + structural rebuilds without net word growth). Major growth occurred in Jose1 (V9 → V36, ~+16k words, 24–26 May). Refinement and additions only in Jose2/3/4/5 (V37 → V77, +~3k words, 26–29 May).

4. Substrate Composition — Per Speaker

Speaker	Words	Turns	Avg words/turn
Paul	45,785	1,066	~43
Claude	1,961,480	1,014	~1,934
Claude (prose-only, tool-call JSON stripped)	543,307	1,014	~536
ChatGPT	12,137	39	~311
Total	2,019,402	2,119	—

Claude's per-turn average (~1,934 words) is heavily inflated by tool-call JSON. Prose-only per-turn average is ~536 words — still much larger than Paul's ~43 words/turn. The asymmetry is structural: Paul directs in short bursts; Claude generates in long blocks.

5. Per-Chat Composition (Substrate Distribution)

#	Chat	Platform	Paul-words	AI-words	Total	Paul-turns	AI-turns	Share of substrate
1	BookMan_Claude	Claude	16,821	639,236	656,057	390	391	32.5%
2	Jose1_v1_to_v37_Claude	Claude	9,403	465,132	474,535	190	192	23.5%
3	Jose3_v42_Claude	Claude	6,544	267,747	274,291	195	197	13.6%
4	Jose4_v69_to_v76_Claude	Claude	3,486	259,193	262,679	91	90	13.0%
5	Jose2_v37_to_v42_Claude	Claude	4,367	240,516	244,883	105	107	12.1%
6	NGEFOF_Hallucination_Claude	Claude	591	43,983	44,574	16	16	2.2%

#	Chat	Platform	Paul-words	AI-words	Total	Paul-turns	AI-turns	Share of substrate
7	SonOfBookman_Reading_Claude	Claude	176	27,140	27,316	3	3	1.4%
8	Jose5_v76_onward_Claude	Claude	488	16,355	16,843	10	10	0.8%
9	Delegate_Scoring_Questionnaire_ChatGPT	ChatGPT	833	4,637	5,470	17	18	0.27%
10	Classification_TwoTypes_OnePattern_ChatGPT	ChatGPT	274	2,711	2,985	5	7	0.15%
11	RecognitionOfUnderstanding_Claude	Claude	228	2,178	2,406	8	8	0.12%
12	Summarise_PaulRoebuckBio_ChatGPT	ChatGPT	1,575	397	1,972	3	3	0.10%
13	Hallucination_DriftAndCompacting_ChatGPT	ChatGPT	103	1,840	1,943	3	4	0.10%
14	NGEFOF_Continuum_AsImage_ChatGPT	ChatGPT	115	1,549	1,664	3	3	0.08%
15	DefendedGap_ConsideredOpinion_ChatGPT	ChatGPT	81	1,003	1,084	4	4	0.05%
16	TunnellImage_Refinement_ChatGPT (RTF)	ChatGPT	466	0	466	19	0	0.02%
17	BookCover_ImageBrief_ChatGPT (RTF)	ChatGPT	234	0	234	4	0	0.01%
	Total		45,785	1,973,617	2,019,402	1,066	1,053	100%

Clustering: Book Man + 5 Jose chats = 1,929,288 words (95.5%). All 8 ChatGPT chats combined = 15,818 words (0.78%). The substrate is overwhelmingly the Book Man → Jose chain.

6. Daily Activity Curve

Date	Paul-words	AI-words	Total	Paul-turns	AI-turns	AI:Paul ratio
2026-05-10	2,052	101,683	103,735	83	83	49.6 : 1
2026-05-11	2,568	75,150	77,718	86	87	29.3 : 1
2026-05-12	2,178	7,005	9,183	36	31	3.2 : 1
2026-05-13	494	11,265	11,759	20	20	22.8 : 1
2026-05-14	0	0	0	0	0	—
2026-05-15	0	0	0	0	0	—
2026-05-16	0	0	0	0	0	—
2026-05-17	531	58,378	58,909	22	23	109.9 : 1
2026-05-18	886	58,665	59,551	27	27	66.2 : 1
2026-05-19	0	0	0	0	0	—

Date	Paul-words	AI-words	Total	Paul-turns	AI-turns	AI:Paul ratio
2026-05-20	1,431	88,492	89,923	35	36	61.8 : 1
2026-05-21	1,140	98,087	99,227	28	29	86.0 : 1
2026-05-22	145	25,087	25,232	9	9	173.0 : 1
2026-05-23	5,203	125,523	130,726	53	55	24.1 : 1
2026-05-24	3,885	225,655	229,540	97	98	58.1 : 1
2026-05-25	4,248	201,639	205,887	78	79	47.5 : 1
2026-05-26	5,102	190,486	195,588	95	95	37.3 : 1
2026-05-27	6,348	333,457	339,805	178	182	52.5 : 1
2026-05-28	3,096	250,732	253,828	90	92	81.0 : 1
2026-05-29	2,797	110,176	112,973	71	68	39.4 : 1
2026-05-30 (ChatGPT export-date block)	3,681	12,137	15,818	58	39	3.3 : 1

Peak day: 2026-05-27 = 339,805 substrate-words = 7.4× the entire final manuscript in one day. Quiet day (non-zero): 2026-05-12 = 9,183 words.

The peak-day pattern: 27 May was the end-to-end transition day — Book Man retired (12:51), Jose2 closed (11:12), Jose3 opened (11:06), Son Of Bookman read in (14:03). End of one editorial era, start of the next. The substrate's gear changed in one calendar day.

7. Hours-of-Day Distribution (Claude chats only — ChatGPT lacks per-turn timestamps)

Bucket	All turns	%	Paul	%	Claude	%
Morning (06:00–11:59)	464	22.9%	234	23.2%	230	22.7%
Afternoon (12:00–17:59)	718	35.5%	359	35.6%	359	35.4%
Evening (18:00–21:59)	427	21.1%	213	21.1%	214	21.1%
Late night (22:00–05:59)	413	20.4%	202	20.0%	211	20.8%
Total	2,022	100%	1,008	100%	1,014	100%

Afternoon dominant. Late-night substantial (~20%). Evening (dinner-hours) thinnest.

8. Working Method Discipline — Per-Chat Counts

Discipline marker	Total
Paul corrections (direction-resetting)	87
Paul ratifications (direction-locking)	111

Discipline marker	Total
Confidence-test moments	74
Named-position transitions	9 documented
Compactions / passons	5–6
Manuscript versions surfaced	27
Platinum-flag signals (any colour+Platinum)	272
Top-10 Platinum editorial moments	10 (selected against named criteria)

Direction-resetting vs direction-locking ratio: 87 : 111 = 1 : 1.3. Almost balanced. Working-method discipline visible as numbers.

Corrections + Ratifications — Per Chat

Chat	Corrections	Ratifications	Total
BookMan_Claude	39	40	79
Jose1_v1_to_v37	23	20	43
Jose3_v42	8	23	31
Jose2_v37_to_v42	5	16	21
Jose4_v69_to_v76	7	9	16
Other Claude + ChatGPT	5	3	8

Pattern: Book Man + Jose1 carry 71% of corrections (122/172). Discipline calibration concentrated in the founding chats; later instances inherited the method.

9. Named-Position Transitions

#	Date	Time	Source chat	Transition
1	2026-05-10	(founding)	BookMan_Claude	Book Man emergence (implicit; named retrospectively at retirement)
2	2026-05-25	17:01:29	Jose1_v1_to_v37	"Jose" first appears (book content — Ch 10 cancer chapter "No way Jose")
3	2026-05-26	09:50:58	Jose1_v1_to_v37	Jose named as working persona — Paul: "fort this project yuo are Jose"
4	2026-05-26	10:15:50	Jose2_v37_to_v42	Jose handover to fresh chat

#	Date	Time	Source chat	Transition
5	2026-05-27	11:06:53	Jose3_v42	Jose continuation (Receive Handover)
6	2026-05-27	12:51:12	BookMan_Claude	Book Man retirement (named at retirement)
7	2026-05-27	14:03:53	SonOfBookman_Reading	Son Of Bookman role introduced
8	2026-05-28	18:20:34	Jose4_v69_to_v76	Jose passon to Jose4
9	2026-05-29	14:55:30	Jose5_v76_onward	Jose5 continuation (final substrate-window instance)

Six named AI working-personas across the window. Identity transferred via memory + handover notes + passon discipline.

10. Platinum-Flag Discipline

Per-colour	Count	%
Red + Platinum	155	57.0%
Green + Platinum	99	36.4%
Amber + Platinum	18	6.6%
Total	272	100%

Per-chat	Platinum count	% of total
BookMan_Claude	181	66.5%
Jose1_v1_to_v37	81	29.8%
Jose3_v42	5	1.8%
Jose4_v69_to_v76	5	1.8%
All others (incl. all ChatGPT)	0	0%

Per-day	Platinum count
2026-05-20	15
2026-05-21	15
2026-05-22	1
2026-05-23	92 (watershed day)
2026-05-24	54
2026-05-25	57
2026-05-26	17

Per-day	Platinum count
2026-05-27	11
2026-05-28	9
2026-05-29	1

Discipline drift observation: Protocol intent was "exceptional, used rarely". 272 flags in 10 days $\approx$ 27.2/day. Three days (23–25 May) carry 75% of all Platinum. Pre-20 May: 0 flags (the protocol enters on 20 May, peaks within 3 days, collapses in later instances).

11. Largest Single Turns

Rank	Speaker	Words	Date / time	Chat	Note
1	Claude	60,650	2026-05-28 21:25:56	Jose4_v69_to_v76	Massive single generation pass (mostly tool-call output)
2	Claude	42,416	2026-05-27 16:14:21	Jose3_v42	
3	Claude	29,043	2026-05-29 07:52:31	Jose4_v69_to_v76	
4	Claude	26,635	2026-05-20 22:42:25	BookMan_Claude	
5	Claude	26,427	2026-05-28 22:49:33	Jose4_v69_to_v76	
1	Paul	1,561	2026-05-30 (ChatGPT export date)	Summarise_PaulRoebuckBio_ChatGPT	Paul pasting his bio article
2	Paul	810	2026-05-23 14:34:47	BookMan_Claude	"THIS its the same as I wrote - just with your prose..."
3	Paul	672	2026-05-26 20:17:54	Jose2_v37_to_v42	
4	Paul	631	2026-05-23 14:43:26	BookMan_Claude	
5	Paul	605	2026-05-11 16:33:03	BookMan_Claude	

Largest Paul turn in any Claude editorial chat: 810 words (23 May 14:34, BookMan). The 1,561-word ChatGPT turn is a paste of Paul's blog material, not a typed working turn.

12. Intimacy Register — Paul to AI

Type	Count	Notes
Sincere thanks (beyond transactional)	11 (10 short + 1 CAPS register)	Includes "THANK YOU mon ami", "10/10 for today"
Appreciation (superb / brilliant / 10/10 / remarkable / love-this)	21	
Apologies (genuine)	7	Including "sorry I have been elsewhere" + "I'm sorry because this is so materially important"
Affection markers ("mon ami" / "mom ami" / "I love you man")	3	
Goodnight exchanges	3	Late-night sign-offs (see below)

The three explicit goodnight exchanges:

Date / time	Chat	Paul	Claude response
2026-05-21 00:54:24	BookMan	"Night"	"Night, Paul. ■" (2 sec later)
2026-05-28 00:29:16	Jose3	"Night Jose x"	"Night, Paul. It has been an honour. Sleep well. x" (2 sec later)
2026-05-28 23:50:05	Jose4	"Thank you mom ami. Goodnight."	"Paul. The cover is the book..." (warm address; substantive close)

Intimacy register attaches to named, persistent working-instances. Zero occurrences in any ChatGPT chat.

13. Named Uncertainties Resolved (per Q31)

Name	Substrate finding
Ewan [name held]	Confirmed correct. Paul's correction at BookMan turn #431 (13 May 10:43): single-line Paul turn "Ewan [name held]" — correcting the earlier "Ewan [name held]" error. 26 hits for correct form; 0 hits for Ewen/Euan/Owen variants. Manuscript front-matter drifted back to "[name held]" in Jose4 (29 May) — Paul had to re-issue the correction.
Grandchildren — long form	[name held], [name held], Erin, [name held]. Consistent spelling, consistent order across the corpus. No variant spellings found.
Grandchildren — short form	[name held], Aurey, Erin, Raffy (dedication-page register, introduced in Jose4 28 May front-matter consolidation).

14. First Appearances (per Q12–Q17, Q28)

What	First appears	When	Chat / turn
"Mind the Gap" as title	Paul names it	2026-05-10 19:53:15	BookMan turn #31 (line 1202) — 38 minutes into the founding session
Book-status recognition ("It's the book")	Paul	2026-05-10 19:38:36	BookMan turn #30/31 (line 504 onward)
Jose as working persona	Paul	2026-05-26 09:50:58	Jose1 turn #345 (line 34447)
Bill Ollis (Paul's first therapist)	Paul	2026-05-11 00:11:27	BookMan (line 5550)
Two Dads (chapter material)	Paul	2026-05-11 00:13:08	BookMan (line 5586)
NGE + FOF + compression + hallucination together	(cumulative)	2026-05-12 10:31:02	BookMan (all four cumulated by this turn)
"Kieren not Karen" correction	Paul corrects Claude	2026-05-11 14:28:02	BookMan (line 8918) — Paul's full turn: "Kieren. Update and store in your artefacts"
First Chapter-1 opening-options offer	Claude offers A/B/C	2026-05-10 20:17:09	BookMan turn #52 (line 2498) — three candidate openings; Paul picks A (Plugs)

15. Caveats and Boundary Conditions

- Code Editor 1 chat is archived (PDF Archive Chat Logs/) and not in substrate. Early manuscript editing (V0 → V8 region) is not captured. The 42.8:1 AI:manuscript ratio is therefore an *underestimate* — *true conversation-cost including the archived chat will be higher*.
- The two RTF ChatGPT chats (TunnellImage, BookCover) carry Paul-only turns — the ChatGPT image responses were not preserved in the RTF source.
- Word counts use \S+ tokenisation — differ from $w_c - w$ by ~0.5% per chat for edge cases (em-dashes, smart-quotes). Immaterial at matrix precision.
- Claude word count includes tool-call JSON (per brief spec). Variant D (prose-only, 543,307 words) is the alternative framing if "what Paul actually read as AI speaking" matters more than "what the AI did, including reading/building".
- ChatGPT chats lack per-turn timestamps. Per-day attribution uses export_timestamp date for those 8 chats. All ChatGPT chats land on 2026-05-30 in the daily curve — outside the 10–29 May Claude window — even though Paul's writing in those chats happened earlier.
- Typed vs. Wispr-dictated input cannot be distinguished from the substrate. Some of Paul's 45,785 words were typed; some were dictated. The substrate does not differentiate. Per Jose1 #274 (25 May 17:26:13): *"I do not hear Paul speak. By the time anything reaches me, Wispr has converted his voice to text."*
- Per-chat ordinal turn numbers start at 1 for each chat. Total substrate turns $\neq$ sum-of-ordinals across chats (they restart per chat).

- The 27,000 founding-draft figure was carried forward in the Substrate Analyst Brief v2 as the manuscript denominator. This was the *pre-substrate baseline, not the final*. This file corrects against V77 $\approx 46,129$ as the *final-at-substrate-close*.
-

16. The Story That Sings

One author. Twenty days. Two AI platforms. Six named working-personas. Five passons. 27 manuscript versions. Sixteen working days. $\sim 46,000$ words of book. $\sim 46,000$ words of practitioner-substrate. ~ 2 million words of AI.

For every word that shipped, Paul put in roughly one — 0.99 substrate-words per manuscript-word. The AI put in 43. The book and the practitioner's hand are the same scale; the AI is two orders of magnitude larger.

The leverage is not in volume — it is in direction. 1,066 short Paul-turns (average 43 words) steered nearly 2 million words of AI output through 6 named instances, 5 compactions, and 27 manuscript versions in 20 days. Paul corrected the direction 87 times and ratified it 111 times — direction-resetting and direction-locking almost balanced (1:1.3), which is the working-method discipline visible as numbers.

Peak day 27 May produced 7.4 manuscripts' worth of substrate in 24 hours.

The manuscript grew from $\sim 27,000$ to $\sim 46,129$ words across the substrate window — a 70% expansion in 20 days. Almost all of that growth happened in the Jose chain (V9 $\rightarrow$ V42 $\approx +18,500$ words, 24–27 May). Book Man worked the founding draft structurally; Jose grew it.

The iceberg above the waterline is the artefact. The iceberg below is how it was made.

The book is the proof; the pond is the demonstration. Both are necessary, and the matrix above is the scale of the pond.

Filed by the Substrate Analyst — 2026-06-01 15:32 UTC. Master statistics dossier — all numerical findings consolidated. Supersedes prior Substrate_LeverageMatrix_2026-06-01_1425.md denominator.