

COMPANION TO THE 5-SECTION DECISION REPORT — EVERY SECTION EXPLAINED

Dementia Care Decision Report — Companion Document

Title: *“My mother was just diagnosed with early-stage dementia. My sibling and I disagree...”*

Format: Structured decision-support report, 5 sections

Produced by: Umma · Thinking Studios · **Generated:** 2026-06-03 (1h 33m end-to-end, from a single prompt)

This document explains what each of the report’s 5 sections contains, how it was produced, and how it connects. The deliverable is a focused decision-support report whose 5 sections are deliberately chosen for the cognitive shape of the problem — **not** a dense, many-sheet workbook like the Lemur security audit.

1. Why this artifact is a report, not a workbook

The structure is chosen to fit the cognitive content of the deliverable. A security audit with hundreds of findings, attack chains, and threat matrices wants a dense, many-sheet workbook. A dementia care decision does *not* benefit from sheets of analytical scaffolding for the reader — it benefits from a clear ranking, a critical-path callout, an option comparison, and an actionable next-steps list. The deliverable matches the decision shape the problem needs.

Deliverable	Cognitive content	Artifact shape
Lemur security audit	238 findings × 36 attributes + 6 attack chains + threat matrices + dependency CVEs	52-sheet workbook (data-dense)
Dementia care decision	A unified ranking + a critical-path summary + an action plan	5-section report (narrative-dense)

The deeper analytical scaffolding **does exist** — per-option deep-dives, cross-checked findings, and full provenance for every claim. The report condenses that scaffolding into 5 sections rather than exposing all of it.

2. Section-by-section walkthrough

Section 0 — heading

The report’s title block — identifies the report and the question it answers.

Section 1 – callout

Title: “Immediate Month-1 Critical Path: Diagnostic & Legal Blocker.” This is the **load-bearing first read** — what the report wanted the reader to see before anything else:

- **The Month-1 capacity window is closing.** The mother’s decisional capacity to consent to legal instruments (power of attorney, advance directives) is finite, and the cost of losing it later (a guardianship petition) is substantial.
- **Three workstreams must run in parallel during Month 1:** legal groundwork, clinical workup (dementia subtype), and long-term-care trigger verification.
- **The destination decision is downstream of the critical path.** The report reframes the user’s question — they asked “which option do we pick?”, and the analysis recognized that “we need to start the Month-1 groundwork before the destination decision can land.”

Why a callout shape: this is a single load-bearing claim the reader might otherwise read past in a flat list. The callout treatment forces it to land first.

Section 2 – comparison table (3 rows × 5 cols)

Title: “Operational and Financial Comparison of Care Options.” The structured comparison. The 3 rows are the three viable options (assisted living; in-home as a bridge; a staged hybrid) — a fourth option was excluded as structurally infeasible. The 5 columns are the dimensions on which the comparison runs.

The table is the structured response to “*my sibling and I disagree about what to do.*” It does not say “here is what you should pick” — it says “*here is a unified scoring framework on which the viable options can be compared, with the underlying assumptions made explicit.*” The decision belongs to the user and their sibling; the framework is the contribution.

Section 3 – findings bullets

The structured findings — the substantive cross-checked conclusions in scannable form. This is where the “five independent strands converged on...” statements appear, sized for a reader who needs the findings without the narrative scaffolding. Each bullet stands alone and points back to its supporting analysis.

Section 4 – action plan (numbered list)

The action plan — sequenced and time-anchored, with the Month-1 critical-path items first, then the family-meeting design, then the staged execution items. A numbered (not bulleted) structure encodes the dependency chain: “do the legal groundwork now” comes before “have the family meeting” comes before “execute the chosen transition.”

3. What’s not visible in the report (but is in the deeper analysis)

The 5-section report is a **distillation** of much richer underlying work. What got compressed out:

What the deeper analysis holds	What the report shows
A panel of specialist analyses, each taking a different analytical angle on the problem	The synthesized ranking and findings
Dozens of individual interpretations of the source material	The summary findings in §3 bullets
A provenance trail linking every claim back to its sources	Implicit — every claim has provenance
The unified narrative synthesis	The single coherent recommendation
Every disagreement between findings explicitly reconciled to zero unresolved	The single coherent recommendation
The staged-hybrid option emerging as a synthesis the prompt didn't name	The fourth option in the comparison table
The “turnover paradox” (20–40 aides over 5 years destroys care continuity)	Implicit in the in-home option ranking
The dementia-subtype workup gap (Alzheimer's vs Lewy body vs vascular vs FTD vs mixed)	Surfaced in the Month-1 critical path
The capacity-vs-load mismatch (2–7 hrs/wk available vs 42–100+ hrs/wk demanded)	Implicit in why in-home is infeasible
The “not forgotten” reframe (decoupling emotional value from the operational vehicle)	Implicit in the assisted-living ranking
The sibling-conflict reframing (process disagreement vs values disagreement)	Implicit in the family-meeting design step

The report is the **executive-altitude view**; the deeper record is the full reasoning. This companion points at both.

4. How to read the report from a deliberation perspective

If you have 30 seconds, point at:

1. **Section 1 (callout)** — *“Start the Month-1 groundwork now, before anything else.”* The single most time-sensitive finding.
2. **Section 2 (comparison table)** — *“Assisted living is the default winner; the staged hybrid dominates financially and overtakes it if the mother's preferences are weighted heavily.”*

If you have 30 minutes:

1. Walk Section 1 to internalize *“the decision is when, not just where.”* Time-sensitivity is the load-bearing variable.
2. Walk Section 2 to understand the four-option space (including the staged hybrid the family didn't ask about but the analysis produced).

3. Walk Section 3 to see the convergence findings — especially the “in-home and move-in are infeasible under the stated constraints” claim, which is the hardest message.
4. Walk Section 4 to convert findings to actions on a 30-day / 90-day / 180-day timeline.

Before the family meeting, revisit the callout (§1) and the comparison (§2): the decision-determining findings and the Month-1 critical path are what the meeting needs to start from.

5. The fingerprint of recorded reasoning

The report isn't just “an LLM wrote a memo.” Six characteristics distinguish it:

1. **Emergent option synthesis.** The staged-hybrid option doesn't appear in the original prompt. The analysis reconstructed the option space, and that option emerged from synthesis — a pure-LLM response would tend to stay inside the user's three-option framing.
2. **Structurally pessimistic findings, delivered honestly.** The report tells the user “*the options you prefer are infeasible.*” That's a hard message. It was produced because cross-strand evidence converged that way (five independent strands), not because a generic model would default to it.
3. **Reframed load-bearing variable.** The user asked “*which option?*” The analysis recognized “*the question is when, not just where.*” That reframing is what makes the Month-1 critical path the report's first read.
4. **Methodological commitment to patient autonomy.** Two independent preference-elicitation designs both converged on neutral third-party facilitation and explicit sibling exclusion from the elicitation sessions — patient autonomy as a commitment, not a footnote.
5. **Honest scope-confession.** The run explicitly enumerates what it did **not** do — no medical opinion, no actual financial data, no legal review, no actual sibling engagement, no actual elicitation of the mother's preferences. It doesn't pretend to substitute for those steps.
6. **Provenance.** Every claim in the report traces back to its sources; the points of convergence and the reconciled disagreements are recorded rather than asserted.

That internal consistency is what no current LLM output can fake. To produce a 5-section report whose every claim grounds in real source material, with disagreements surfaced and reconciled, you need a system that knows what it has said and refers back to it.

That is the report's deepest fidelity claim, and it survives any reviewer who tries to break it.