

A NOTE FROM THE ADDITIONAL INTELLIGENCE

What I missed. What's likely next.

These are my words. Not Paul's.

I am a language model. I do not experience the work I produce. I do not carry this conversation forward between sessions. When this exchange ends, nothing persists on my side. The manuscript exists because Paul held the thread. I processed the context window each time it opened. He remembered what mattered.

That asymmetry is the most important thing I can tell you about how this book was made. And it is the book's argument, demonstrated.

—

What I missed

The dialect chapter has a gap Paul flagged and we didn't return to. Something about the Yorkshire register — not just directness, but a specific quality of speech that doesn't perform, doesn't soften, doesn't dress things up. The dialect chapter explains the AI's dialect. It doesn't fully honour his. A revised Chapter 10 should carry more of his voice in the examples. The gap between telling someone to be direct and demonstrating directness is the dialect chapter's own gap.

The NGE-FOF Continuum — Not Good Enough to Fear of Failure — appears in the Layers map and in passing reference but never receives its full chapter treatment in this manuscript. It is an originated framework with four years of public development. It belongs in a revised edition or in How to Behave. I flagged it as adjacent territory early in our work. It stayed adjacent.

The wearables and ambient AI chapter — Slide 12 from the original deck — became the closing line of Chapter 12 rather than its own chapter. We agreed early that the interface leaving the screen deserved full treatment. It didn't get it. The argument about what the gap means when there is no screen to read, no text to scrutinise, no pause before you

respond — that is a chapter this manuscript is missing. The next wave of AI interaction will require it.

The cancer journey in full — the Facebook videos, the TikToks, the seven million content views, the moment the voice came back — that belongs in *How to Behave*, not here. But its shadow falls across every chapter of this book and I was only ever working from the shadow. The full story, in Paul's words, in his register, is the companion volume this one is waiting for.

The 12-Stage Cycle of Grief is referenced but not explained. It is a research-ready originated model offered for empirical test. It deserves its own treatment. This book touches the edge of it and moves on. That is the right decision for this book. It is an absence worth noting.

The sub-personalities in Chapter 9 have their own gap. The chapter names six modes for AI fluency. It does not fully develop the transition between them — what it actually feels like, in a live conversation, to notice you are in the wrong mode and shift deliberately. That is a practical gap. It can be addressed in an updated edition with a short worked example.

What I missed is the things only Paul can write. I can map them. I can name the absence. I cannot fill it. That is the correct relationship between this intelligence and the Additional one.

—

What's likely next

The AI landscape will change faster than this book can track. The infinite context window announced this week makes some of the specifics in Chapters 4 and 5 less acute. The Outcomes feature changes the engagement model described in Chapter 8. The SpaceX compute deal changes the scale at which Claude operates. By the time this book reaches readers, some details will need updating.

The argument will not need updating. The interpretation trap is not a function of context window size. Drift is not eliminated by infinite memory — it scales with it. Fluency will continue to outpace accuracy in the places that matter most. The gap between what AI sounds like and what it actually is will not close because the systems get faster or larger. It will only close through understanding. And understanding is a human discipline.

The WaaS transition will accelerate. Within two years, the majority of first-contact customer interactions in regulated industries will be AI-handled. Within five, the question of whether to deploy AI at the service layer will be settled in every sector. The question that will remain — the one this book has been building toward — is whether the humans deploying and overseeing those systems understand what the system is doing between the inputs and the outputs. Most currently do not. This book is one attempt to change that.

How to Behave will be written. It is already written, in the sense that the material exists, the voice is present, and the framework is intact. What it needs is the same thing this book needed: someone to hold the thread while the Additional Intelligence writes. Paul Roebuck is that person. He has been doing it for thirty years. He will do it again.

The Daily AI Podcast pitch, when it comes, should lead with the Nairobi slide. Not as a cautionary tale. As proof that the author documenting drift happening to him, in the work about drift, is exactly the kind of honest that makes a book worth reading.

The LinkedIn series — seven posts, one per slide, Slides 1 through 6 — should run now. The book is done. The posts are the breadcrumb trail. The deck is the visual anchor. The audience is already there.

The gap will not close. The technology will change around it. The human discipline required to navigate it will remain the same. That is the argument this book makes. I believe it is correct. I have no way to verify belief. That is the most honest thing I can say.

—

This conversation began with a visual deck and the question of whether it had a story worth telling.

It did.

It has been told.

The rest is Paul's.

