EXTRACT

Human-on-the-Loop

Six Tools for
Marketing Leaders
in the Age of Agents

FLORIAN MICHAEL FISCHER

WELKIN MEDIA VERLAG · VIENNA

Contents

Prologue: Waiting for the Barbarians EXTRACT

PART I · WHY NOW

1 18 Years in Eight Pages

2 The Sorting Is Already Under Way

PART II · THE SIX TOOLS

3 The Anti-To-Do List EXTRACT
TOOL 1 · *Prioritised task list with handover criteria*

4 The Agentic Definition Test
TOOL 2 · *Vendor scorecard with eight criteria*

5 The TwoFunnel Audit
TOOL 3 · *Audit template and GEO quick-start checklist*

6 The AgenticMandate
TOOL 4 · *Avis Manifesto template 2026 in seven points*

7 Human-on-the-Loop
TOOL 5 · *Delegation matrix with five columns and twelve entries*

8 The AI Leverage Audit
TOOL 6 · *Team assessment sheet, 15 minutes per person*

PART III · LEADERSHIP

9 What a CAIO Actually Does

10 Three Things That Will Not Become Commodities

Epilogue: Are the Barbarians Really Not Coming?

APPENDIX

Sources

Glossary

Index

This extract contains the Prologue and the beginning of Tool 1. Page numbers for the full book are given in the table of contents of the printed edition.

PROLOGUE

Waiting for the Barbarians



In a poem from 1898, an entire city waits for its downfall. The downfall does not come, and that is the problem.

It was written by Constantine Cavafy, a Greek civil servant in Alexandria who spent thirty years in the irrigation department and wrote poems on the side. This one describes a single day and needs barely more than a page to do it.[1]

Waiting for the Barbarians

– *What are we waiting for, gathered in the forum?*

The barbarians are due to arrive today.

– *Why is there such a stillness in the Senate?*
Why do the senators sit there and pass no laws?

Because the barbarians arrive today.
What laws could the senators pass now?
When the barbarians come, they will make the laws.

– *Why did our emperor rise so early*
and sit at the greatest gate of the city,
enthroned in state, wearing his crown?

Because the barbarians arrive today.
And the emperor is waiting to receive
their leader. He has even prepared
a parchment to present to him.
On it he has written many titles and names.

1 Constantine P. Cavafy (Konstantinos Kavafis): 'Περιμένοντας τους βαρβάρους' (Waiting for the Barbarians), written in 1898, first published in 1904 as a private print in Alexandria. The English version is my own AI-assisted rendering from the Greek original and not a quoted translation. The original is in the public domain.

– *Why have our two consuls and the praetors*
come out today in their red, embroidered togas?
Why do they wear bracelets set with so many amethysts,
and rings with splendid, glittering emeralds?
Why do they carry costly staffs today,
elaborately worked in silver and gold?

Because the barbarians arrive today,
and such things dazzle the barbarians.

– *Why do our worthy orators not come as usual*
to make their speeches and have their say?

Because the barbarians arrive today,
and eloquence and public speeches bore them.

– *Why does this unrest begin all at once,*
and this confusion? (How grave the faces have become.)
Why do the streets and squares empty so quickly,
and everyone go home, deep in thought?

Because night has fallen and the barbarians have not come.
And some have arrived from the frontier
and said that there are no barbarians any more.

And now what is to become of us without barbarians?
Those people were, after all, a kind of solution.

It is worth staying with this image for a moment. The city has prepared, and thoroughly, but for the wrong thing: for the reception, for the show of state, for the impression one makes when the newcomers arrive. The emperor has covered a parchment with titles for someone he has never seen. The consuls have put on their amethysts. Everyone in this poem is busy, and busy with waiting.

Then come the two closing lines, the reason the poem is still read more than a hundred and twenty years later. The city responds to the good news with bewilderment. The threat had given everyone a part to play, and with the report from the frontier the script is gone.

I came across this poem through Ethan Mollick, and given where artificial intelligence stands today I found it more than

fitting. We are only too happy to wait for the solution to all our problems. But what if it does not come?

Why this is here

You have opened a book about artificial intelligence in marketing. So you know the barbarians that are currently being awaited.

Sometimes they are called AGI, sometimes the great wave of automation, sometimes the point at which the models take over our work. The Singularity.[2] Expectation swings between redemption and apocalypse, and both versions have the same practical effect: they postpone the decision. Anyone who waits for an event that will settle everything has nothing to settle today.

The waiting looks like work, too. There are innovation rounds, tool evaluations, pilots, training sessions, roadmaps and very good presentations about what will soon be possible. The senators in the poem also sit together, every one of them present. They pass no laws.

And then comes the report from the frontier. In our case it arrives quietly, as a series of shifts, each one too small to force a decision, and which taken together have changed the profession. Whoever waits for the one day on which the technology is finished and one can finally begin will wait as long as the emperor at the gate.

Whoever waited for somebody else to fire the starting gun eventually finds that the race has been running for a while.

---

2 Ray Kurzweil: *The Singularity Is Nearer. When We Merge with AI*, New York: Viking 2024. Kurzweil understands the Singularity as a metaphor for the point at which humans merge their intelligence with AI, and gives around 2045 as the time horizon (Introduction). For an AI passing a Turing test he gives the year 2029 (Chapter 1). The predecessor *The Singularity Is Near* appeared in 2005.

What this book is for

In eighteen years in this industry I have read enough forecasts to know how badly they age, and I have made enough of my own to be careful with other people's. Tools are much the same for me. The tools I work with today will not exist in this form in two years, and a book that clings to them carries its use-by date on the cover.

This book therefore sticks to the questions you have to answer while you wait. Who decides what. Which work belongs in the shredder before anyone speeds it up. How you can tell whether your team is building a lead right now or losing one. These questions can be answered today, and leaving them lying around does not make them any easier.

For that you get six tools from the daily life of a Viennese performance marketing agency. All six have been used, revised several times, and have broken in places that I write down for you as well, so that you do not have to find the same holes again.

After these six tools you should know two things: which side of the sorting you and your team are on right now, and what you will do in the next ninety days to stay on the right one.

Before we begin

A warning in my own cause. I write from inside an agency. That means I have seen the cases in which something worked, and also the ones in which we tore it down again after three months.

So that you can always tell what a statement rests on, this book distinguishes three kinds of evidence. Studies, reports and talks by third parties carry a footnote with the source. Observations from my work at SlopeLift are written in the first person, the clients are anonymised, and each of these observations is a single case, even if it has repeated itself several times. Assess-

ments and conjectures are marked as such, with words like ‘I believe’ or ‘probably’.

The book is therefore practitioner literature: it reports from application and offers tools. If you quote from it, you are quoting the practitioner’s perspective. Figures and study results should be backed with the original source, which you will find in the relevant footnote.

Nobody will put an entry in the calendar for the arrival of the barbarians. But we had already started to wait when we did not yet know there were barbarians. That waiting began almost twenty years ago, with a slice of Gugelhupf, the Austrian ring cake, and an Excel spreadsheet.

3

The Anti-To-Do List



TOOL 1 · THE ANTI-TO-DO LIST

What to say goodbye to before you buy a single tool

■□□□□ TOOL 1 · PART 1 OF 5

Diagnosis

Most AI projects fail on a question that is never asked, because it seems too obvious to ask: Why are we still doing this ourselves at all?

Instead, a different question gets asked, almost reflexively, in every meeting about adopting AI: How do we do this faster? The question sounds reasonable. It is the beginning of nearly every disappointing AI initiative I have seen in the last three years.

I told myself this story for a while. At SlopeLift, in 2023 and 2024, we built workflows with the AI team, tested them and tore them down again. Some took off straight away, others went completely wrong. And at some point, after enough torn-down workflows, it struck me that the problem had never been the technology. It was the frame of thinking. We took an existing process and bolted a tool onto it, expecting the same thing to come out the other end, only faster.

What actually comes out when you speed up a bad process is more of the wrong thing, faster. That sounds obvious once it is said out loud. In practice you only see it when the heap of output is already lying in front of you and somebody asks what exactly you are supposed to do with it now.

The sea of sameness

There is a second reason why the question ‘faster’ leads you astray, and it has to do with differentiation.

The sea of sameness from Chapter 2 describes the condition: same tools, same prompt, same output. Anyone who uses AI to produce the same thing as everyone else, only faster, ends up producing the same thing as everyone else. Just more of it. Differentiation then no longer happens by itself, and it becomes the real task.

The same consequence follows from both observations. As long as the first question is 'How do we do this faster?', you are optimising precisely the activities that are about to become commodities. You invest energy in speed at a point where speed will soon be worth nothing, because everyone has it.

There is a third observation, which is voiced less often because it is a little uncomfortable. The 56% from Chapter 2 is a figure about organisations. The corresponding figure at team level is probably even worse. Many marketing teams have introduced AI tools that a few individuals use and many others never touch. The pattern is always similar: the one person who uses the tool is faster. Everyone else works as before. The result is a skewed distribution. The Anti-To-Do List is a team exercise, not an optimisation for individuals. It works only if everyone involved sits at the table and brings their tasks.

The shredder before the tool

The more productive order turns the question around. Before a process gets a tool, it belongs under scrutiny: does it belong in the shredder?

That is the diagnosis behind the Anti-To-Do List. Most organisations keep a to-do list of tasks that have to be done, and they look for tools that speed those tasks up. The Anti-To-Do List looks for something else: tasks that vanish from the table altogether, because a human should no longer touch them at all.

The difference is small in the wording and large in the effect. A faster task remains your task. A handed-over task leaves your desk. And only when it leaves your desk does the capacity

become free that you need for the work in which differentiation actually arises.

My guess as to why the share without measurable return is so high: in most of these companies the ‘faster’ question was asked and the ‘why are we still doing this ourselves?’ question was skipped. They automated their VLOOKUP and wondered why the gain failed to appear.

TOOL 1 · PART 2 OF 5

The Framework

The Anti-To-Do List is a systematic search for tasks that a human should no longer touch at all, because an agent takes them over completely. At SlopeLift, three categories crystallised out of this, and they turn up in a similar form at almost every company.

Category 1: The jobs nobody wants to do

Data cleaning, formatting, repetitive reporting steps. The things that qualified people feel above, but that have to be done all the same, and that therefore often get left lying around or are done half-heartedly.

This is the most grateful category, and it is the best entry point, because it wins over sceptics. Nobody defends their claim to replacing umlauts in a CSV. The time saved is visible at once, and there is no argument about wounded professional pride. If you start anywhere, start here.

Category 2: The knowledge gaps

Quick research in new industries, competitor analyses, briefing preparation. Everywhere that missing information blocks the process before it has even begun.

This category is subtler, because the work here often appears disguised as waiting time and is not recognised as work at all. Somebody cannot move forward because they would first have to look at ten competitors, and puts it off, because looking at ten

competitors is tedious. An agent does this research sprint in minutes, and the human steps in at the point where the material has to be interpreted.

What distinguishes Category 2 from Category 1 is correctability. Data cleaning forgives errors: a mistake shows up before it does any harm. Research results can sound convincingly wrong and still travel into a client document, because nobody checked the figures a second time. Handover Criterion 3 therefore has to be taken particularly seriously in Category 2: the Anti-To-Do List hands the task away, but a human keeps the step of interpreting and checking.

Category 3: The blank-page blocks

The leap from zero to something. First drafts, concept frameworks, overcoming the emptiness that stands at the beginning of every creative process.

The greatest resistance lies here, and it is understandable. A first draft feels like authorship. But the first draft is seldom the valuable part. The valuable part is the judgement that follows the draft: what stays, what goes, what hits the client and what misses them. An agent that fills the blank page dissolves the block and leaves the authorship with you.

What the science says, and where it warns

The Anti-To-Do logic has empirical support that is rarely cited, although it belongs exactly here. In 2023, a team around Fabrizio Dell'Acqua at Harvard Business School, together with Boston Consulting Group, observed 758 consultants at work with GPT-4. The field experiment has since appeared in peer-reviewed

form.[1] A companion study from the same project then took a closer look at 244 of these consultants on a complex analysis task, almost 5,000 individual exchanges with the machine, and teased out three working patterns.[2]

The largest group, around 60%, worked as Cyborgs: in constant dialogue with the AI until the boundary between human and artificial thinking blurred. About one in seven worked as a Centaur: delegating individual subtasks, but keeping strategic control. A good quarter squeezed the whole workflow into one or two passes and adopted what the machine delivered, almost half of them without changing a single word. The study calls this group Self-Automators; here I call them delegators. The Centaurs most often arrived at the right recommendation. Cyborgs and Centaurs convinced with their results markedly more than the delegators. And the delegators were the only group that learned nothing, neither about the subject nor about working with AI.

That is the decisive point. The Anti-To-Do List sorts out precisely the tasks for which full handover is truly right, the commodity work from the three categories, so that your attention is freed up for the work in which you have to be a Centaur. Anyone who instead delegates everything and leans back ends up with

---

1 Fabrizio Dell'Acqua et al.: 'Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of Artificial Intelligence on Knowledge Worker Productivity and Quality', *Organization Science*, online ahead of print 11 March 2026, DOI 10.1287/orsc.2025.21838, CC BY 4.0 (previously Harvard Business School Working Paper, 2023). Pre-registered field experiment with Boston Consulting Group, n = 758 consultants. Within the capability frontier 12.2% more tasks completed, 25.1% faster, over 40% higher quality; on the task outside the frontier a 19 percentage point lower probability of a correct solution compared with the control group without AI.

2 Steven Randazzo et al.: 'Cyborgs, Centaurs and Self-Automators: The Three Modes of Human-GenAI Knowledge Work and Their Implications for Skilling and the Future of Expertise', Harvard Business School Working Paper 26-036 (2025), SSRN 4921696. Companion study to the field experiment in footnote 1, 244 consultants, almost 5,000 human-AI interactions. Distribution: Cyborgs 60%, Centaurs 14%, Self-Automators 27% (rounded). Referred to in the main text as 'delegators'.

the delegators who lagged behind in the study. You delegate the one thing in order to be more present for the other.

How costly this can become is shown by the field experiment itself, with figures that cut both ways. On tasks that GPT-4 mastered, the consultants with AI completed a good 12% more, were a quarter faster and delivered over 40% higher quality. On a task that the researchers had deliberately set beyond what the model could do at the time, the picture tipped: the group with AI was 19 percentage points less likely to reach the right solution than the group without. Whoever delegates where the machine cannot reach is wrong more often than someone who does not use it at all. Dell'Acqua calls this invisible line the jagged frontier.

Daniel Hulme calls the emerging system the Extended Adaptive Mind, a thinking apparatus distributed across human and machine.[3] Tool 5 builds out this architecture. For the Anti-To-Do List, one point from it counts: it is the inventory that determines which node in this distributed system takes on which work.

How jagged this frontier runs across whole fields of work was measured by OpenAI in October 2025. GDPval tested 1,320 tasks from 44 occupations, developed from the work of professionals with an average of fourteen years of experience and assessed in direct, blinded pairwise comparison by working professionals.[4] In several task areas Claude Opus 4.1 and GPT-5 reached or exceeded the level of the human comparison group, Claude Opus 4.1 above all in the form of the results, GPT-5 above all in

3 Daniel Hulme: 'Extended Liquid Minds', 2026, referred to in the main text as Extended Adaptive Mind. Theoretical foundation: Andy Clark, David Chalmers: 'The Extended Mind', *Analysis* 58 (1), 1998, pp. 7–19; Andy Clark: 'Extending Minds with Generative AI', *Nature Communications* 16, 4627 (2025).

4 OpenAI: 'GDPval: Evaluating AI Model Performance on Real-World Economically Valuable Tasks', arXiv, October 2025. URL: https://arxiv.org/abs/2510.04374. 1,320 tasks from 44 occupations (220 of them openly accessible), tested on GPT-5, Claude Opus 4.1, Gemini 2.5 Pro and Grok 4, among others, with blinded pairwise comparisons by industry professionals with an average of 14 years of experience.

accuracy and instruction-following. This very irregularity is the jagged frontier in numbers: peak values in some places, a deficit in others, often distributed unpredictably.

A second lesson follows from this, which matters for the later series of tools. Where the AI is good, it is very good. Where it hits its limit, it is convincingly wrong. That is exactly why it is not enough to hand tasks over once and never look again. For the Anti-To-Do List this means in concrete terms: the category in which a model shines today may be one in which it fails in six months, or conversely one that is still blocked today and can soon be released. Criterion 3, correctability, therefore belongs in the same quarterly rhythm in which the list is revised anyway. But that is material for Tool 5. For now, the inventory is what counts.

TOOL 1 · PART 3 OF 5

SlopeLift in Practice

Two hours, four minutes

The most convincing example is the most boring one, and that is no coincidence.

At our agency there was a recurring data-cleaning job before the weekly reporting that reliably ate two hours a week. Merging exports, aligning formats, repairing umlauts, hunting for duplicates. It is the direct successor of the data yoga I practised in 2008, when I was eating Gugelhupf and building pivot tables at the same time. Eighteen years later, an agent did the same work in four minutes.

Nobody in the team mourned those two hours. That is the punchline of the first category. The task was necessary, unloved and fully parameterisable. It had no defender in the room. And because the gain was visible the next day, this one example dissolved more scepticism in the team than any presentation on the potential of agentic workflows.

END OF EXTRACT

Human-on-the-Loop

Six Tools for Marketing Leaders
in the Age of Agents
Florian Michael Fischer

PUBLISHER
Welkin Media Verlag, Vienna

PAPERBACK
ISBN 978-3-9506420-3-2

EBOOK
ISBN 978-3-9506420-4-9

Illustrations: Lilly Pannholzer