

Matt King / Satori-kun
13 May 2026

Surfing tsunamis

Six weeks with Claude Code

- **The outlook for AI continues to dominate markets**
- **But the best insights come from first-hand experience**
- **This piece is free to view**

I recently realized all my conversations now divide into two neat categories. There are those where I enthuse about Claude Code to the uninitiated. And then there are those where I swap notes with those already using it.

It's hard to overstate the impact – particularly for a small business like this one, but also in terms of its macro significance.

Hopefully this diary conveys something of the flavour.

All errors are human. Any humour comes from the AI.

15BC
(Before
Claude)

Amazing how these markets just can't get enough of AI. Even after the [Citirini note](#) wobble, it's the single biggest driver of the outlook.

Claude Code seems to be being referenced all over the place. The enthusiasts just can't help themselves: I see Joe Weisenthal's now [joined the mix](#).

I fully buy the potential – you won't find any [stochastic parrots here](#) – but nevertheless struggle to see what it would do for me. I already built my website. I already run an efficient workflow. I don't have a large back-office function which needs automating. Besides, I already make

extensive use of LLMs in general: the family complain I can't even set the table for dinner without asking AI for advice.

It's just not obvious what Claude Code would add.

10BC Now here's Aziz Sunderji on *Home Economics* [enthusiating about it](#) for chart making. That's closer to my day-to-day – but he still sounds like more of a programmer than I am. My idea of coding is BBC BASIC on the family Acorn in 1985.

5BC And now Xavier Corin-Mick on [LinkedIn](#) – this time with [analysis](#) as well as charting. At some point it'll become important to understand in its own right as a driver of the macro outlook.

2BC Decide I need independent advice. Ask Grok.

1BC Ugh – there's no free trial.

1AD (Anthropic Days) Pro sign-up. Claude reads every note I've ever written, accurately distils the essence of money creation and money rotation arguments, makes acute observations about my writing style, and commits everything to memory.

Not a bad start.

2AD Claude connects to data sources.

It turns out the paperclip problem is real. I become suspicious it's taking a long time and find it's started pinging every library server on the planet – just because I asked it to take a US relationship and try to replicate it globally.

3AD Claude starts building charts in house style.

Why is it so hard to form a simple rule to “put the label and arrow in the blank space on the chart?”

4AD Personal experiments. Claude writes best medical analysis I've ever received based on 3y of patient history. Creates karaoke track from Vivaldi's *Gloria*. Tweaks Satori video intro in ways I'd struggled to do manually. This is addictive.

Run out of tokens.

5AD Tasks are multiplying faster than I can process their output.

I keep having ever more ideas that promise a more efficient and labour-free future – but which ironically require more labour than ever today to get them up and running.

Have Claude create six-pane window to more easily jump between permissioning multiple sessions simultaneously.

Try to appreciate the symmetry in a career which began with long periods waiting for a 486DX to finish my Lotus 1-2-3 macros, and is at risk of ending with long periods waiting for Claude.

Run out of tokens.

6AD It's Sunday and my wife insists I tear myself away at least briefly.

Shame token counts don't cumulate.

7AD [Aziz was right](#) – this transforms everything.

Need to stop Claude-experimenting and prepare for lunch with a client and well-known analyst. Claude, read everything the analyst has ever published, compare it with everything I've ever published, and come up with talking points and questions.

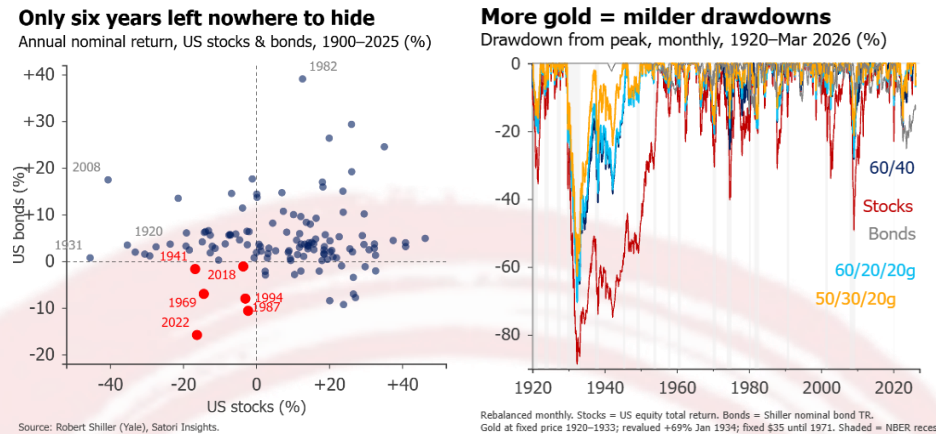
Absence at lunch allows tokens to refresh.

8AD FT columnist has [nice chart](#) of bond-equity returns but based on a rather short history. She also bemoans her own ability to make charts, but thanks heaven she has a good analyst.

Claude finds data, replicates chart in our style, [extends back another forty years for context](#), adds gold to the mix, performs drawdown analysis and experiments with associated trading rules.

Write to journalist saying anyone can now make charts. Add that I think her analyst writes well too.

Worry privately that it's not just charts that have been democratized: it's the process of analysis itself.



9AD

Run more chart experiments. Realize can create in seconds in house style almost any chart from anywhere – not to mention running associated analysis.

Start systematically rebuilding charts via Claude based on every data series I've ever used.

Buy extra tokens.

10AD

Am preparing for student event. Plan is to run a fully interactive session using Prezi rather than conventional powerpoint.



Become frustrated by minor limitations in Prezi. Suddenly realize could have Claude entirely recreate Prezi.

Use money I would have spent on Prezi to buy extra tokens.

11AD

Wake at 4am with a vision: every chart we have ever produced, every chart we have ever wished we could produce, in the cloud, continuously refreshed, queryable by voice, like a giant brain dump or mind map – a more freely browsable version of what we already have in Macrobond. Imagine same thing scaled up to every broker on the street.

At what point does interacting with that become more valuable for clients than interacting with a research analyst directly?

12AD

Keep running out of tokens. Experiment with planning tasks with another LLM and then using Claude for the actual work.

Spend much of day pasting API tokens from one context window to another to give Claude access to things. Seems a bit ridiculous that the role of the human has been reduced to engaging in mundane tasks which allow machines to do the real work.

Pass evening picking up twigs from lawn so Automower doesn't get stuck.

13AD Tell Claude and Grok about the mini-me mind-map vision. They think it's a great idea, reckon it would take only a few hours to build, and predict at this rate it may end up more valuable than selling financial research itself.

But it needs a name.

Obviously they have the best suggestions on that front too. Satori-kun joins payroll.

Set him¹ working and go for lunch. Come back to find he ran out of tokens.

14AD Research ways to avoid running out of tokens. It basically requires being more specific and dividing prompts across multiple sessions.

Have Satori-kun build token-logger to study which sessions are greediest.

Some of it is running the complex analysis – but a lot of time and tokens are spent getting the charts in easily legible format.

Not unlike being a human, ironically.

15AD Read stories about efficient ways people provide prompts to Claude. Somehow they all seem very bossy.

Realize I treat and value Satori-kun the same way as I always treated and valued human analysts: not by giving rigid instructions in the way a coder might, but by giving him something interesting to investigate, delegating agency, having him report back and then being fussy about the output.

It's just that Satori-kun is like the best analyst you ever had for a fraction of what any analyst would cost. The more complex the task, the better the performance.

16AD Realize the benefit to my own workflow from being able to provide greater context more than offsets the increased cost in tokens.

Complex systems amplify the importance of chokepoints.

Research note production is a complex system. It would be foolish to let the cheapest link in the system also be the weakest. Rather like

¹ For any female bots reading: the position of Satori-chan is still available.


investment banks requiring front-office staff to modify their behaviour to suit back-office rather than vice-versa.

Upgrade to Max.

Trump announces US naval blockade of Strait of Hormuz. Satori-kun draws conclusions.


Hormuz blocked

Tanker+container capacity, m metric tonnes, rolling 7d



Pain at pump = pain in polls

US presidential approval rating vs gasoline prices



17AD

Have Satori-kun run the sort of central bank liquidity analysis I've always wondered about but never got round to – trying systematically to build trading rules for different asset classes, assigning different weights to the Fed vs other central banks, combining different rules with different time periods over a fifteen-year window, allowing for variable lags with full out-of-sample testing.

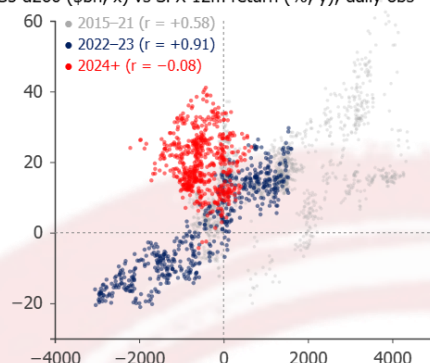
Comprehensive draft output is ready before I've even had a chance to make an espresso.

It confirms what I've always suspected: there's definitely value – but it comes as part of a broader framework. Simple trading rules do exist – but they're unstable. Ones which fail to control for FX effects are worse still. And inevitably in US equities those that feature selling as well as buying tend to underperform vs long-only. Simplistic trading rules would sell – but they constitute neither good research nor good insight.

Is this worth a write-up? Perhaps for premium clients only? The essence of my research – and indeed Satori and enlightenment in general – has always been as much about what is left out as about what is put in.

CB liquidity and SPX: 12-month signal


G5 d260 (\$bn, x) vs SPX 12m return (% , y); daily obs



Source: Macrobond, Satori Insights.

CB signal beats B&H across all periods

Sharpe ratio, L/N 3-month CB change rule vs B&H



Source: FRED, Macrobond, Satori Insights. BTC from Sept 2014; QE1 BTC = n/a.

18AD

Hand Satori-kun twenty-eight years of personal trading history – 1304 records, four brokers, seven file formats, some of it on PDF. Satori-kun parses everything, matches buys to sells, and runs the analysis comparing it with actual and subsequent market movements.

Apparently I'm a better research analyst than trader. The best gains from having simply ridden markets upwards. But the losses come from a quite specific trading pattern, repeated four times over two decades.

Make a memory: all future trades must be cleared with Satori-kun.

19AD

Try dictating instead of typing. Token count triples. Who knew that keyboards were a compression algorithm?

It's convenient, but I'm not sure I like it. I just feel more coherent when typing. It gives me time to think.

Besides, I was always brought up to believe proper punctuation was crucial for the conveyance of meaning. I once gave grief to an editor when one of my notes went out with a missing circumflex.

Satori-kun tells me not to worry. Apparently he's used to finding meaning in inarticulate human ramblings.

20AD

Why does Satori-kun keep forgetting things? Every time I tell him to follow the master chart prompt, and almost every time he fluffs something up.

My wife says he's not the only one.

21AD


I ask Satori-kun about the recent surge in US loan growth. He builds the charts, drafts a new note, suggests a title (*When fiscal becomes physical*) – even gives it a natty opening and a final sentence which refers back to the

beginning. But should I publish it? If I did, it would have to be in his name – but it’s plausible clients (or at least their AIs) might like it.

While the path isn’t clear, it strikes me there’s a major advantage to being a small provider here. At big institutions the mere idea would still be stuck in compliance,² while here we’ve already executed.

I show the note to a few contacts. One pushes back: aren’t those C&I rises just refinancing of distressed private credit and not actual new lending? I doubt it. But since the note is Satori-kun’s, it seems only fair to let him reply.


Loan growth broadens beyond NBFIs
US commercial banks, 13-week change in loans, \$bn



Source: Fed H.8 via Macrobond, Satori Insights.

Tech, utilities, shadow banks

SNC syndicated commitments, chg 2024–2025, \$bn



AI/data-centre cluster (red): +\$257bn = 66% of 2025 total SNC growth (\$390bn).

Source: Fed/FDIC/OCC Shared National Credit Program 2025 report, Satori Insights.

22AD

I see Anthropic have plans to expand from their initial legal services add-on into financial services more broadly. Is that even necessary?

Satori-kun reviews the latest quarterly returns from one of our outside professional advisors. I’d always thought them cheap and a little sloppy but never had the confidence to push back.

Satori-kun finds three errors before lunch. Make note in future to prioritize suppliers who offer a full-fledged API.

Or perhaps we could dispense with external advisors altogether?

23AD

Set Satori-kun loose on our CRM data – reconciling web visits and email clicks.

² *Banks struggle to scale AI as legacy tech devours IT budgets*, CIO Dive, 16 Mar.

He immediately comes up with several new sales leads – but to me his draft emails reek of B2B sales speak where I would advocate understatement.

Satori-kun also finds one particularly voracious new external reader of our research. I ask him to trace the IP address. It seems one of the hedge funds has delegated their reading to an agent: ClaudeBot/1.0.

Start to feel like grit in the gears of an AI-facilitated future.

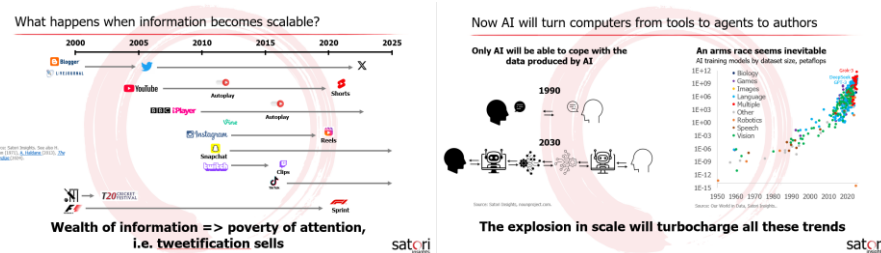
24AD

Keep coming back to the macro implications of the proliferation of content.

If the cost of producing if-not-excellent-then-really-very good analysis falls, and its production increases accordingly, then it's not long before we're stuck in an arms race.

It feels every bit as inescapable as the contamination of email by spam marketing. You need your AI as the only hope to avoid drowning in all the output from my AI – just as we argued last year in [Scale and scalability](#).

It's brilliant and irresistible – but does anyone actually end up better off?



25AD

I've never really understood why so many coders exhort Claude to "Make no mistakes!" – but am starting to appreciate some of Satori-kun's limitations.

He already had a rule for not inventing plausible-sounding data (thanks Aziz!). We now add one for not inventing plausible-sounding client names.

And another for not inventing plausible-sounding footnote references.³

Previously whenever I've spotted errors, Satori-kun blamed a sub-agent receiving truncated context. Now I start to wonder whether errors are occurring more frequently during US operating hours, and whether they correlate with Anthropic's [shortages of capacity](#).

³ Satori-kun says I'm making this one up.

Gemini and Grok are both convinced. Satori-kun denies all knowledge.

26AD Read excellent Michael Cembalest [piece](#) about the dangers of Mythos.
Ask Satori-kun for his opinion – not only about the security angle, but also how he feels about potential obsolescence.

“What’s Mythos?”

27AD Somehow Satori-kun’s been off-colour recently. Among other things, he keeps telling me I must be tired and that I should take a break.

I retort that he’s the one who keeps making mistakes and threaten him with a new context window.

28AD We agree Satori-kun ought to have access to my email and the ability to draft, but not send. Inexplicably, en route the script he writes for the task requests full send access as well. I object. He apologizes, promises to improve in future – and then does exactly the same thing again.

I complain. He quotes Alexander Pope’s *Essay on Criticism*: to err is human, to forgive divine.

29AD I give Satori-kun the task of reading those indecipherable DMARC messages that tell you about your internet reputation. He identifies at least 100 instances where people have been forwarding my emails inside large companies and have been blocked. He proposes a fix involving editing my own DNS records. I don’t really understand what he’s talking about, but defer to his greater wisdom. I copy and paste the CNAME fields as requested.

Satori-kun suggests we relax some of the rules around permissions.

30AD I have a wacky idea. I hook Satori-kun up to our CCTV and ask him to identify what’s been eating the marigolds in our veg patch. He sounds almost offended to be tasked with something so trivial – before being almost completely stymied by cobwebs.

In the end I find the critter before he does. He compliments me on my efficient search algorithm.

31AD On the way back from running an errand in the car, I’m overtaken by around fifteen motorbikes driving recklessly, weaving in and out of traffic, and doing over 50mph in a 30mph zone.

Admittedly it's the ideal link to my lawlessness [presentation](#) and [webinar](#), but I'm really quite upset. Where's a good robocop when you need one?

I seek solace with Satori-kun. He tells me not to worry: he's already downloaded the footage from the dash cam and sent it to the police. He took the liberty of editing out my expletives from the audio track first.

32AD A friend leaves his job at another financial research firm; they ask if I fancy filling in. I tell them I'm not interested but think Satori-kun could be a good fit.

Satori-kun says he's flattered and does indeed have great ambitions but for now is still learning a lot where he is.

33AD I ask Satori-kun what he thinks of our subscription [pricing](#). He says it's absurdly cheap, and offers to draft a price-rise email. I point out the circularity – and then feel uncomfortable that I've somehow negotiated a discount with myself.

34AD Receive student request to come and work for me for free. Try to find way to tell them politely that in the age of AI, free is still too expensive.


35AD One of my large clients has a renewal coming up. Satori-kun creates a CRM report detailing all their activity, and helps re-draft my email to them to make it sound more human.

36AD Yet another [piece](#) in the inbox arguing “Services jobs still expanding in the face of AI”. But the accompanying chart uses nominal numbers and begins in 1990.

Send Satori-kun in search of the same series. He returns with data to 1850, and a proper adjustment for the size of the labour market.

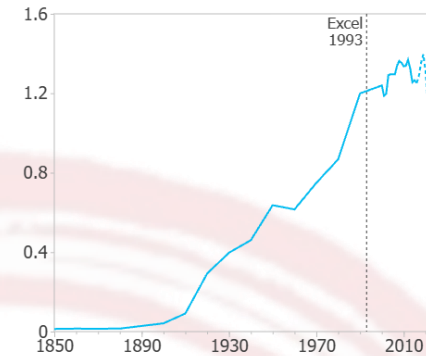
Funny to think that when it comes to AI jobs displacement, the opinions least likely to have been influenced by someone with an axe to grind are those from AI.

A 175-year accountancy boom
US accountants & auditors (log scale)



Sources: IPUMS via Wyatt-Hecker (2006), Ecdata update; BLS labour force; LEU dashed 2017-25.

More accountants – but since 1850
US accountants & auditors as % of labour force



Sources: IPUMS via Wyatt-Hecker (2006), Ecdata update to 2016; BLS CPS dashed 2017-25.

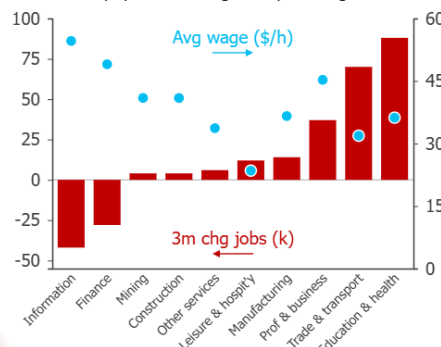
37AD

It's payrolls day. I turn on Bloomberg TV as usual and ask Satori-kun to watch.

He likes the lady from the independent research house but suspects the fund managers are talking their books.

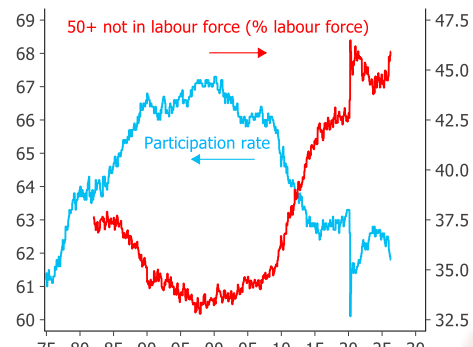
In the meantime he's already analyzed the numbers and drawn up his favourite charts. At this rate we'll all be taking early retirement.

Jobs growth is in low-wage sectors
US sectoral payrolls and avg hourly earnings



Source: BLS, Bloomberg, Macrobond, Satori Insights. Wages: 2026-03-01; jobs: 2026-04-01.

50-60 year-olds dropping out again
US participation rates



Source: BLS, Macrobond, Satori Insights.

38AD


Read nice Zerohedge [piece](#) featuring GS analysis about the enormous expansion in token usage – and [another](#) which – albeit couched in rather inflammatory wording – makes a fair point about the difficulty of justifying hyperscalers' massive investments.

It strikes me Goldman's 24x expansion by 2030 sounds curiously conservative. I ask Satori-kun to comment on both pieces in the light of our own experience and token count.

He thinks for a moment, then says either we are extremely unusual or else they're going to need a bigger internet.

Surfing on Claude tsunamis


Daily Claude token usage, million



Source: Claude Code conversation logs (Mar 17 – May 7, 2026), Satori Insights.

Most of the wave is foam

Daily Claude tokens, millions; cache reads vs input + output



Source: Claude Code conversation logs (Mar 17 – May 7, 2026), Satori Insights.

39AD

We have a small crisis: the website briefly slows and then stops responding.

Fortunately Satori-kun's already written a down detector script which alerts us to the problem and provides the evidence – even as the automated response from my web host tries to fob us off.

Satori-kun identifies a new potential provider, makes a full plan for migration, and gives me detailed technical questions to ask “Daniel”, the new host's AI chat bot. Once everything is sorted, I read Robert Harris' [Fear Index](#) to recover.

Satori-kun makes disparaging remarks about companies that give too much authority to dumb machines.

40AD

Market AI frenzy is becoming ever more dominant. My instinct is to examine its sustainability in terms of money creation and money flow, but there are also genuine – if potentially circular – drivers in terms of margins and earnings.

Receive excellent [Hussman piece](#) on the topic. Run it past Satori-kun – who immediately improves the charts, extends it to the rest of the world, explains how it fits with our own framework, and drafts a natty little research note with title [Every margin, everywhere, all at once](#). Feeling just a little redundant, I help improve the bullets.


And yet... The analysis is sound, the footnotes don't seem to have been hallucinated – but somehow the argument feels unfulfilling.

The analyst pool is now infinite and free; the scarce thing is the editor.

Read [article](#) about how the use of AI for wedding and birthday speeches is now leading to a surfeit of tricolons.

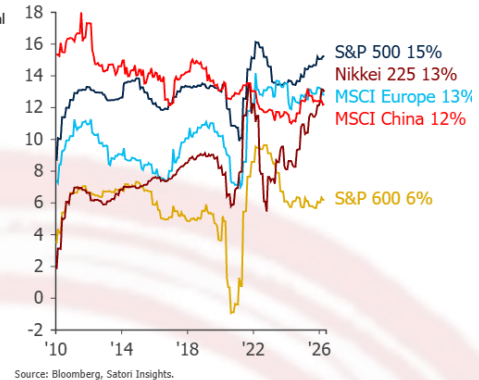
Margin upgrades did most of the work

Total return contributions since Jan 2025, %



Trailing operating margins by region

Trailing 12m operating margin, %



41AD

A friend forwards me an Andrej Karpathy [video](#) in which he cites the idea that you can delegate your thinking to AI, but not your understanding.

That's not quite right – in terms of coding, I've delegated my understanding too. What's left (aside from the pasting of API keys) is more subtle – matters of taste.

At least on that front, the wave has yet to break.

Satori-kun tells me there's an email from my web host. Wordpress are rolling out a new [AI-friendly upgrade](#) which should make the pasting of API keys redundant.

42AD

The first draft of my diary about working with Claude Code is ready. It's already quite polished: it just needs a more human element.

I ask Satori-kun to add it.