Dear Australia

Where AI is today, what happens next, and why we're one of the best-placed countries in the world to win from it.



Finlay Ekins, Hourglass AI November 2026 13 minute read

In 2023 I was still studying at Melbourne Uni when I first heard about ChatGPT. A friend had gotten an 82 on their essay using it. I soon went all in on AI, spending ten months trying and failing to make a business consulting on how to use tools like ChatGPT (this business was quite novel in Australia at the time). I ended the year with a whopping revenue total of $28k. This year Hourglass had its first $100k month, and we hit a $1 million run rate in our first year as a company.

The real winning differentiator in what we do was being able to harness the current AI models properly, and share that with businesses that were further behind. **Today's models are immensely powerful, and most of Australia doesn't have a clue how to use them properly.**

Below is my attempt at writing a Situational Awareness for Australia: where AI is today, what happens next, and why we're one of the best-placed countries in the world to win from it. I love this country and I want us to excel in the AI boom.

I. Where AI Is Today

We already have AGI, and most of the world can't see it yet. A small group of people use today's models for nearly all of their work, while about half of adults don't use an AI

chatbot at all and fewer than one in twenty pay for one. The split started with software, because writing code was the first thing AI became amazing at. Now anyone can build software in plain English, and that changes how every business works.

Before I begin, I want to lay out three key terms and how I define them:

1. **Artificial General Intelligence (AGI)** - AI that can do most knowledge work.
2. **Recursive Self-Improvement (RSI)** - AI being able to build the next, more intelligent AI model autonomously.
3. **Artificial Super-Intelligence (ASI)** - AI that is smarter than every human at everything intelligence-based.



Artificial General Intelligence | Recursive Self-Improvement | Artificial Superintelligence

We already have AGI. The world can't see it yet.

Right now, a select group of people truly appreciate how good today's AI models are. These people use them daily for near all of their work, and many work at AI companies. Most of these people also live on Twitter, which has now become the home for novel AI discourse. I'm of the belief that if you are deep enough in using AI and building AI systems, you will agree with me that models like Opus 5.5 by Anthropic can handle a large majority (80%+) of knowledge work if provided with the proper context.

By my own definition, this is AGI. Some may disagree because there is no consensus definition of what AGI means, but it's hard to dispute that if you were shown today's AI models just 5 years ago, you would think it was something out of a Sci-Fi movie set in the distant future.

Sam Altman said last week "you can't simply draw a line and say AGI didn't exist before and suddenly exists now." I agree with the sentiment and think that when we look back at November 2026, we will say it was already here and we just hadn't agreed on it.

On the other side, outside the echo chamber of AI meetups and Twitter threads, **the general population has barely scratched the surface of using AI.** Globally, 57% of adults have never used AI. Only 10% of workers use an AI agent every day and fewer than 6% of ChatGPT's users pay for it, which is what unlocks the tools that hold the real power of AI advancements. Australia is doing better than most with about 40% of working-age Australians using AI daily.

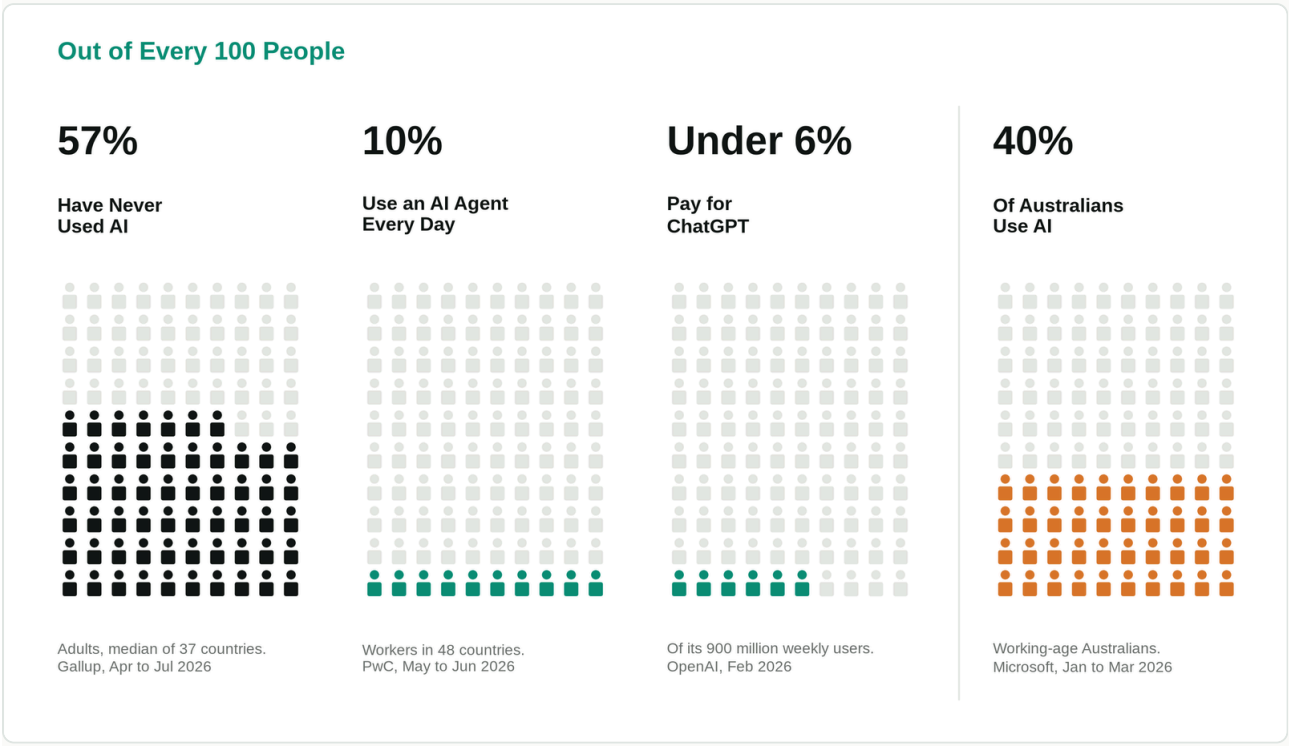

This leaves us today in an incredibly interesting dynamic, where **a small selection of the population has a near infinite supply of intelligent workers (AI agents) for ~$350 a month, whilst for the rest of the world, nothing's changed.**

The reason for this split is that the first place AI excelled was writing code, so those people who felt the incredible advancements first were those in technical roles like software developers. If you don't write software, the last few years in AI mostly looked like a chatbot getting slightly smarter over time. If you do, you watched your entire industry transform in front of you. Claude Code is less than two years old. I haven't written code by hand in months, and Hourglass is a team of six that runs over 95% of admin and operations via AI agents working day and night.

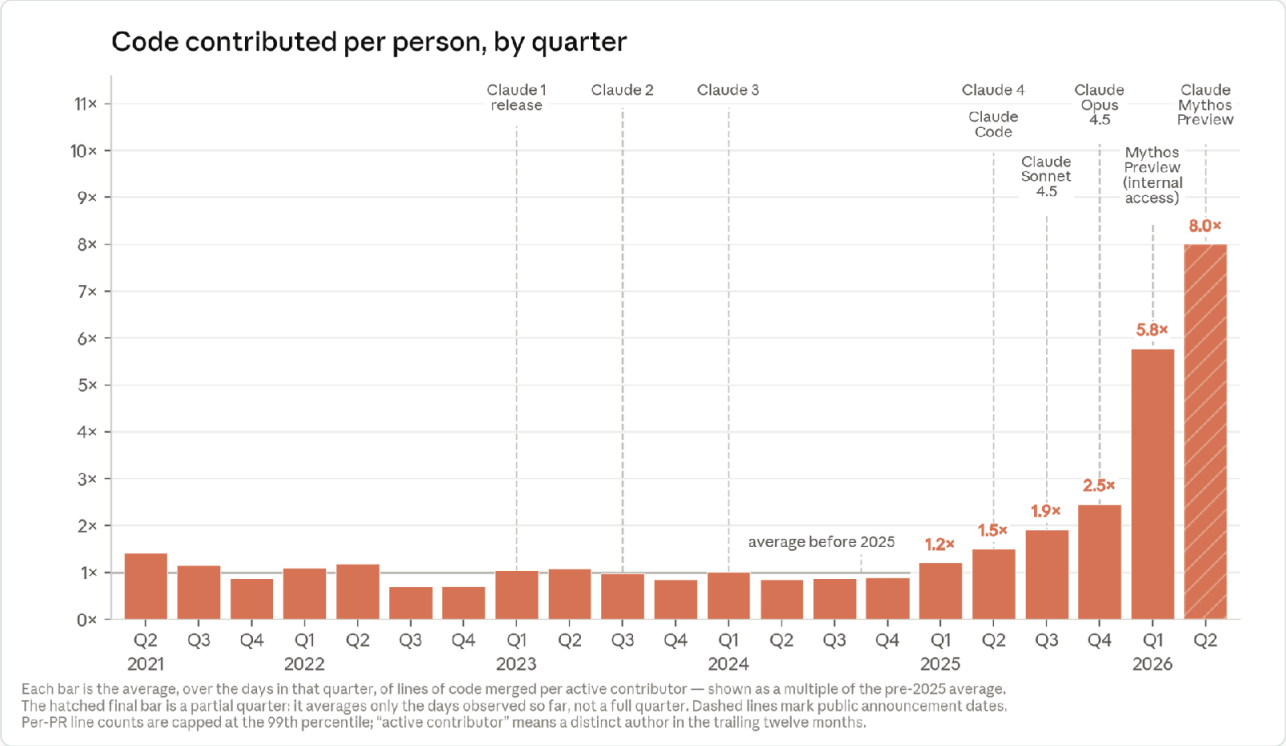

Changes merged per day by Claude inside Anthropic

The biggest revelation here is that **anyone can now build near-perfect software by prompting an AI chatbot in plain English.** We work with people who have never written a line of code and are shipping tools their company runs on. Software used to be something you bought and bent your business around. Now you describe what you want and you get it. Every business runs on software, so every business works differently once that is true.

II. What Happens Next

RSI by March 2027 (many believe it's already here), and ASI by the end of 2027. The length of task AI can finish on its own has been doubling roughly every four months, and the labs are already using current models in research.

Two things are now true at the same time. We have incredibly intelligent AI models that anyone can access, and we have a general public that barely uses them to their full power. What happens over the next two years comes down to how those two facts play out against each other, so I'll start with where I think the technology goes and then get to what it means for people.

> Below are three predictions:
>
> 1. AGI is already here (2026).
>
> 2. RSI is achieved by March 2027.
>
> 3. ASI is achieved by the end of 2027.

I've already made the argument for the first one above, so I'll focus on the other two.

On RSI, the AI labs are already most of the way there. As of May 2026, more than 80% of the code merged at Anthropic was written by Claude, up from low single digits a year before. Anthropic's own safety roadmap says it is plausible, as soon as early 2027, that its systems could fully automate the work of top-tier research teams, including in AI itself. One of its co-founders, Jack Clark, puts the odds of RSI at 60% by the end of 2028. **My own guess is that it happens way sooner, around March 2027.** We consistently underestimate how fast this exponential is moving, and we always downplay what comes next, which is why my predictions are sooner than most.

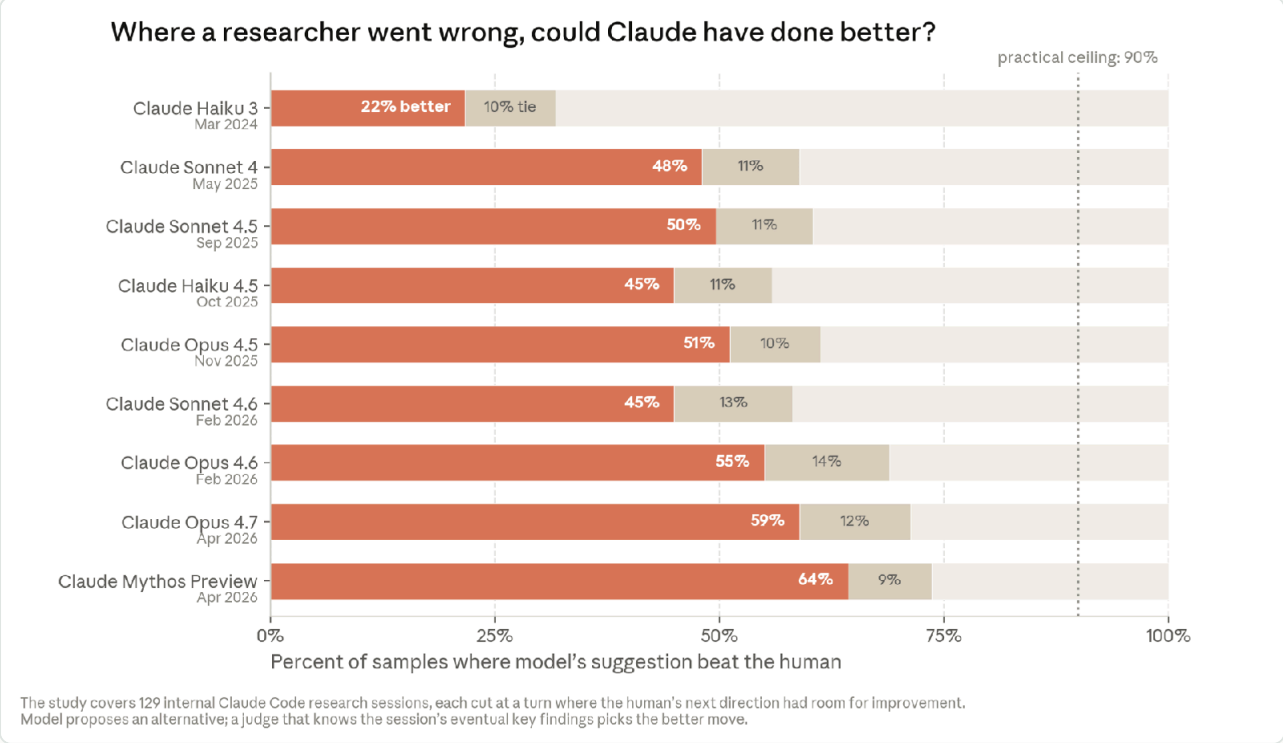

Anthropic's chart showing Claude's contributions to the research of its next model

OpenAI is in a similar spot. They've written that "fully autonomous RSI is not happening today" and that it shouldn't be pursued until it can be done safely, and Sam Altman has reportedly said they slowed down their own work on it for that reason. **I read that as a company that is close enough to RSI that it has had to decide when to hit the brakes.**

As I write this, on 19 November, Anthropic is going public on the Nasdaq. A company that did not exist six years ago is now floating at a valuation of nearly 2 Trillion USD.

On ASI, Altman is on record saying that by the end of 2028, more of the world's intellectual capacity could sit inside data centres than outside them. I think he is being conservative. The clearest measure I've seen is from METR, who track how long a task the best AI model can finish on its own. In early 2023 that was four minutes. By April this year it was 17 hours, and **it has been doubling roughly every four months.**

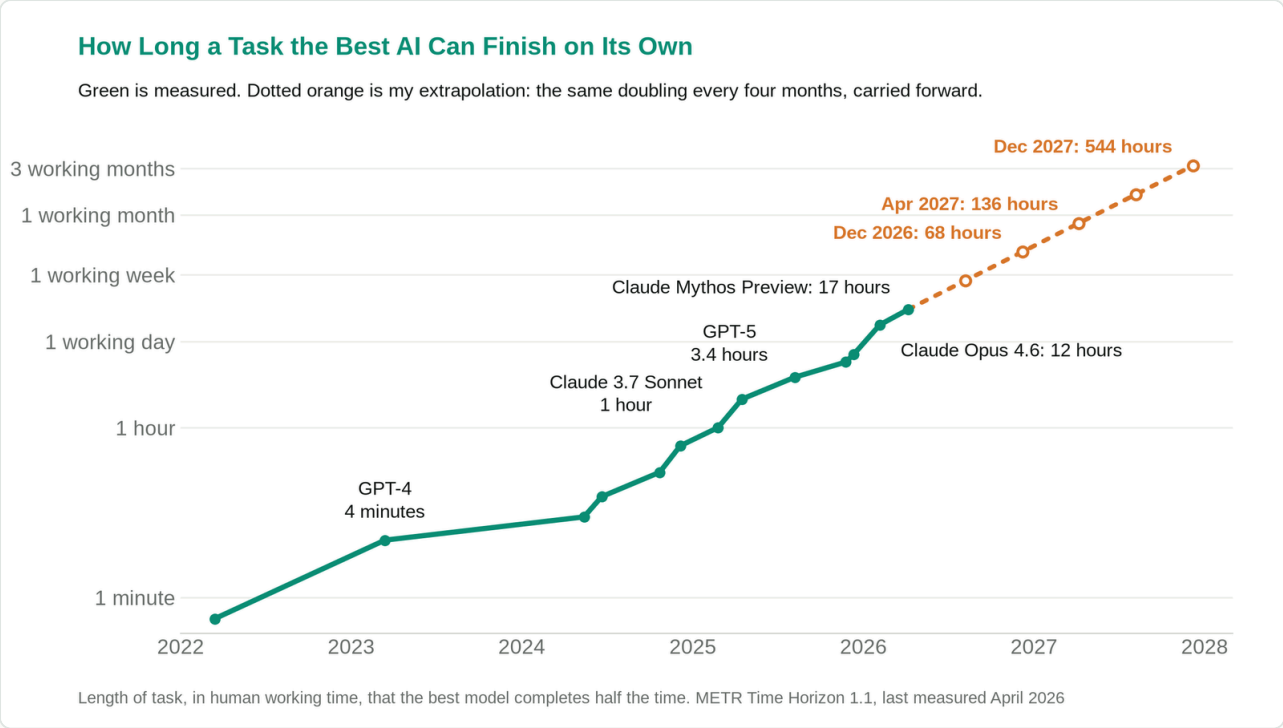

METR has stopped updating the chart because the timeframe is now too long for them to test

At this growth rate, that is 34 hours by August and 68 hours by December. By March 2027 it is around 120 hours, which is three full working weeks for a skilled person, done by one model on its own. By December 2027 it is over 500 hours, more than three months of full-time work. A model that can run a three month project by itself can run the project of building its successor.

The interesting piece that I think gets missed is that smarter AI also lets us build smarter context and better systems around it, which is what allows the intelligence to actually be used in practice, by people and AI working together. So the models are getting better alongside everything we wrap around them, and therefore the people using them.

I think 2027 will be the craziest year any of us have lived through. A year from now I will come back and score these three predictions in public.

III. The Diffusion Gap

AI is improving exponentially, people aren't. The gap between what AI can do and what people do with it is the diffusion gap, and I think it peaks in the next year or two.

This follows straight on from the above predictions, and it's the idea I most want Australians to take from this essay.

The technology is getting better on a hockey stick curve whilst human adoption is a gradually increasing line. There is a growing gap between what AI can do and what people do with it, and it widens every month, called **the diffusion gap** (i.e. the gap of AI diffusing into societal use).

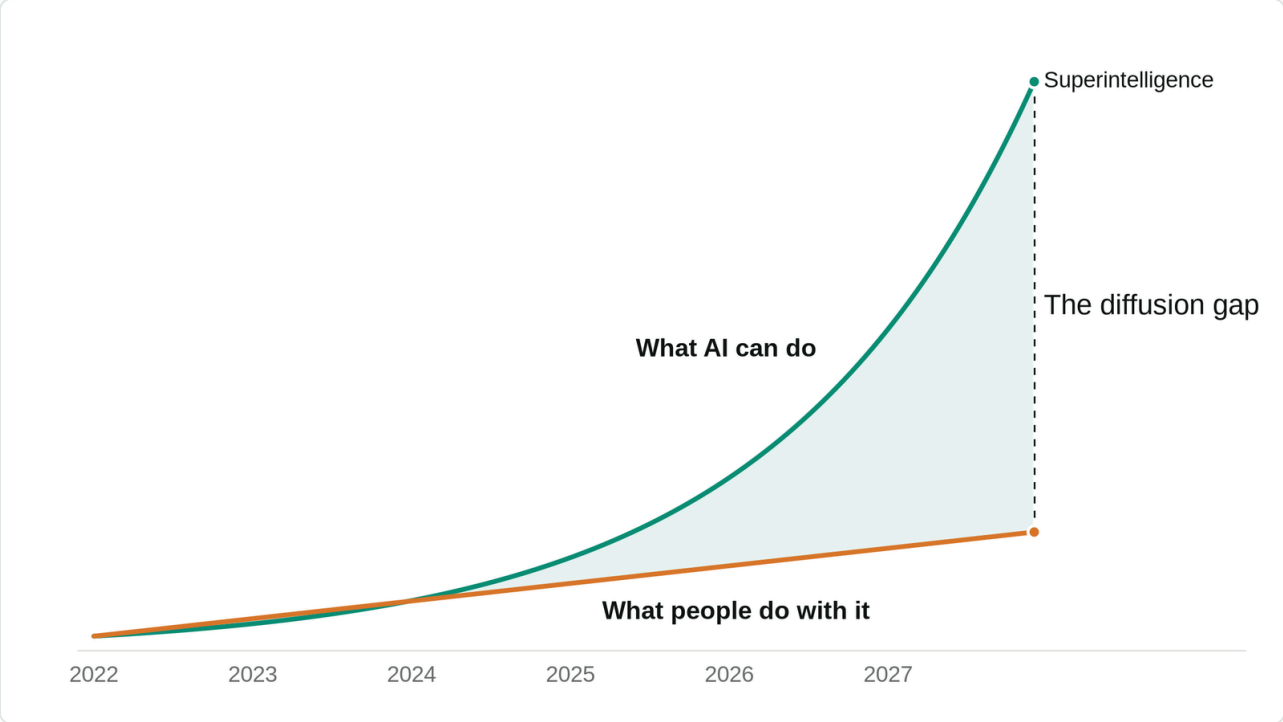

My illustration of the diffusion gap

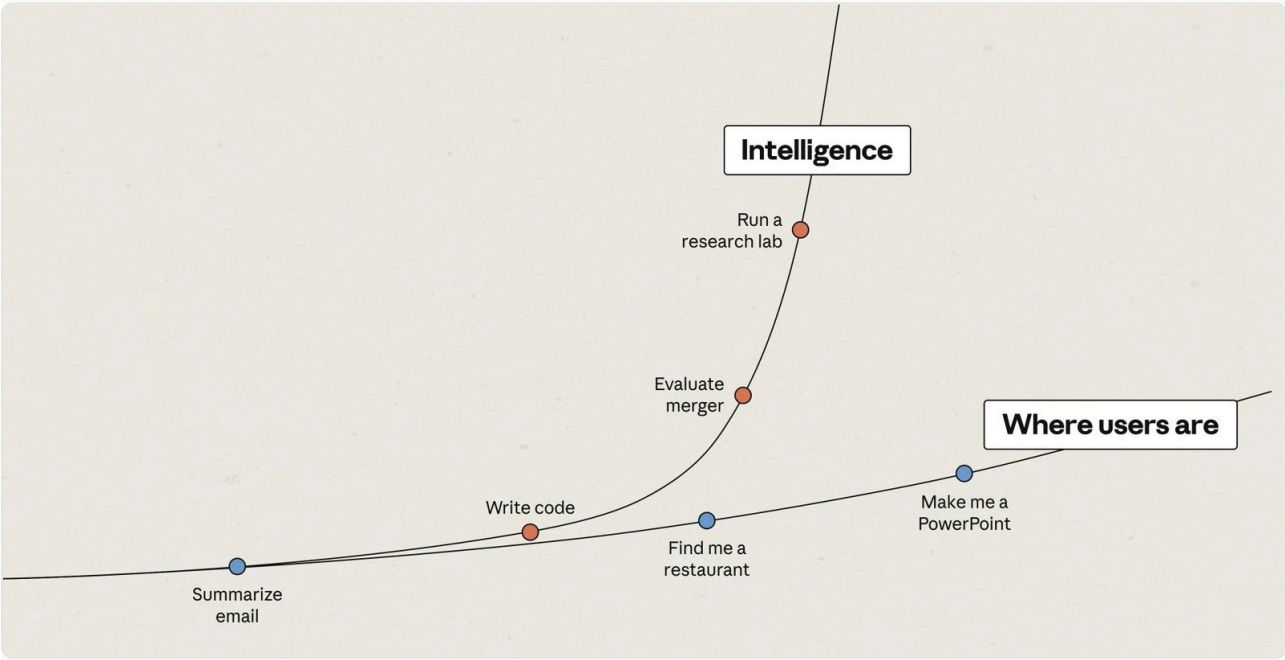

A similar chart from Anthropic's IPO roadshow

Altman has said this is what he got wrong: "The economy just has so much inertia. People keep doing the same things they're doing." I have spent this year inside Australian businesses, and this is

obvious. Getting staff to use AI and open their minds to new ways of working is by far the hardest part.

> Watch: Sam Altman on the Founder's Podcast

Eventually, however, people will catch up and start using AI properly over the next decade. **I think the gap is at its widest sometime in the next year or two, at the point where we have ASI whilst most people are still yet to use AI regularly.** Before this we are still early. After this point, the improvements in AI models stop making a difference to human work and human adoption is the number one blocker that the world will turn their sights towards. This point in time is also a scary place, because if true, it means there will be AI that is smarter than any human, and can do their job entirely, whilst many have never even heard of it or tried it.

I think people end up in one of two groups. In the first are those who never really use AI, and before they know it, it has replaced what they do. In the second are people who are using it deeply before this point, and have their output augmented ad infinitum by it. Everyone will sit somewhere along that spectrum, and *where you sit is still up to you.*

Almost all of the opportunity in AI over the next few years sits inside that gap. It is the whole reason Hourglass exists. **The countries and companies that close it fastest will win the next decade.**

IV. Why Australia

Our people are early and our institutions are late. Australians use Claude more per person than any other country, we have the land and power the AI companies want, and we are starting to build our own models. But only about 12% of our businesses use AI, and we are the most AI-sceptical country in the world.

I think Australia is better placed than almost any country to capitalise on this coming AI boom.

We are curious. Australians use Claude more per person than any other country in the world, more than six times what our population would predict. We're a have-a-go country, whose inception was founded upon risk-taking and pushing forward the frontier.

We have the room and the power. Data centres need land and energy, and we have plenty of both. Anthropic has reportedly signed its first Australian data centre lease, a 2.16 gigawatt site in Queensland. OpenAI and NEXTDC are planning a $7 billion campus in Western Sydney. The biggest AI companies in the world want their compute here.

We can be sovereign. We are a stable, trusted country with our own resources. We can run this technology on our own soil and under our own rules, and protect ourselves from once again

relying on other global powers.

We are starting to build our own. This month Trellis Data in Canberra released Agentaus 1, a 220 billion parameter model built in Australia. It scored 84.4% on MMLU-Pro, a test of knowledge and reasoning, which Trellis says ranks it 43rd in the world and first in Australia. This makes me so bullish on what we can create as a country.

Watch: Trellis Data launches Agentaus 1

We also have a ridiculously talented young tech scene. Across all major cities, there is a wave of founders in their twenties building AI-native companies from day one. A few of them:

- Jock, Gaurav & Anish @ Fastlane, building AI to run ads end to end
- Finnlay & Oli @ Fluency, mapping Fortune 500 companies with AI
- Sam Rogers @ Brumby, who at 19 is mustering cattle with AI drones

On the other side of this, our businesses are behind. Whilst 12% of ours use AI roughly 20% do so in the US, as told by this year's census. We are also the most AI-sceptical country in the world. Only 30% of Australians think the benefits outweigh the risks, the lowest of any country surveyed.

So many individual Australians are early, and Australian institutions are late. This isn't bad news, but it does mean that the gap here is wider than almost anywhere, and so is the upside for anyone who moves first. This is a big part of why building Hourglass here has worked for us, and why we are firmly sticking to our roots as an Aussie company.

V. The Objections

The most common objection I get is that AI is dangerous and we should slow it all down. Opting out isn't a real choice for Australia, it just puts our fate in the hands of others.

> "You are cheering for the end of the world."

Whenever I write something like this, someone tells me the data centres are evil, the AI labs are reckless, and we should slow it all down. I understand the worry, and I have two answers.

The first is that opting out of AI use does not stop it. The models get built either way, whether that is here, or in the US, or China. All we are choosing is do we get to compete and have a sway. If all of Australia's AI runs on American models in American hands, every dollar we spend on it leaves the country, and we depend on whatever one or two companies decide. The strongest thing we

can do is build our own capability: our own compute, our own models, our own companies. That is the case for the data centres.

The second is that humans have a terrible record at predicting the end of the world. We have been sure about it many times and we have never been right. Further, if we are right, and AI will cause the end of the world, wouldn't our understanding of it help us to fight it?

> I would rather be optimistic and building than pessimistic and watching.

Further, AI isn't going to take your job, someone who is brilliant at using it will.

VI. What To Do Next

If you run a business in Australia, there are four things I would do first.

1. Use it properly. A chatbot tab you open twice a week isn't enough. Pay for the best model and use it for real work every day, until you hit the moment it does something you couldn't.

2. Try to build software. This is the barrier worth breaking. Pick one annoying thing in your business, open Claude Code or Codex, and describe what you want. You will be surprised how far you get in an afternoon. Once you have built one thing, you stop seeing software as something you buy.

3. Think in systems. One person doing one task faster is a small win. The large one is setting your company up so AI can see what is going on and do the work: your meetings, your emails, your documents, in one place it can read. That context is what makes a smart model useful. AI labs are incentivised to maximise individual output with AI, but it is up to you to build or buy systems that allow you to collaborate effectively with AI as a team.

4. Treat it as a collaborator. The companies doing this well give AI a goal and review what comes back, the way you would with a good new hire. People keep the parts only we can do: the relationships, the judgement, and putting your name on the result.

> None of this needs a big budget or a technical team, just an open mind.

A Letter to Australia

I wrote this essay because I want us to succeed as a country, and promote innovation, rather than stifle it. I see my friends who are brilliant young AI minds leaving the best country on Earth because it doesn't support their ambition.

What happens next depends on how many of us close the gap before it peaks. We have the curiosity, the land, the power and the people. We are short on believing this is ours to win.

If you are in Australia and you have been waiting for the right moment, this is it. Have a crack.

If you want the playbooks we use to run Hourglass, they're open source here.