

Agents of Sovereignty: Earning and owning value in an economy reshaped by AI

Marc Taverner · The Polaris Group · Updated 23 October 2026

Horizon: 2026 to 2031

Where I'm starting

I expect AI to change the economics of knowledge work quite dramatically. As more people can get the same work done, the value will increasingly sit with whoever brings the knowledge, expertise, relationship with the counterparty or customer and has the right to use the information. I want to back businesses that help individuals and companies put those advantages to work with AI, keep control of them and keep a worthwhile share of the value created.

That's where Agents of Sovereignty starts for me. People need to find ways to remain useful and earn a living. Companies want the benefits of AI without handing over the knowledge and relationships their business depends on. Governments have a problem to work through too, because employment, tax receipts and the cost of supporting people are all tied up in this.

My bet is that as AI becomes easier to get hold of, knowing how to put it to work in a particular business becomes more valuable. There should be businesses worth building around that: helping a professional earn more, getting work into production that a company can't responsibly automate today, or making sure the customer can change tools without starting again. The return has to show up somewhere in the customer's economics. Otherwise we're selling a preference.

So the question I keep coming back to is, who actually does better when this product works? And why will they keep paying the company providing it?

Work is changing, even before the jobs disappear

I think AI will change employment substantially, but redundancies are only part of it. Who gets hired in the first place, how a junior learns the job, what a client will pay and how much work a small team can take on. The job title may stay much the same while the economics underneath it change.

The evidence already warrants attention, although it does not justify a confident date for mass unemployment. Stanford's August 2026 analysis finds a widening employment gap for young workers in AI-exposed occupations, operating mainly through reduced hiring. The authors also find no widespread economy-wide displacement and describe the results as descriptive indicators rather than causal estimates. Danish research provides a useful counterweight: it

finds changes in tasks and occupational movement without detectable changes in earnings or recorded hours over the period studied. [\[1\]\[2\]](#)

I don't think we should explain away that difference in the evidence. We don't yet know the pace, or which places and occupations will feel it most. But waiting for an unemployment crisis doesn't strike me as a sensible investment strategy either. A lot of the commercial relationships that determine who benefits will be settled before then.

The immediate opportunity lies in adaptation. Businesses will experiment with smaller teams, different roles and new ways of delivering services. Individuals will need to understand which parts of their contribution remain valuable and how AI can extend them. Some will prosper inside employment. Others will build businesses, combine several sources of income or find that their previous work commands a lower price.

I want to back businesses that work through that uneven transition. Some may benefit from rapid displacement, but I'm not comfortable making mass unemployment a condition of the whole thesis.

Doing more work and getting paid for it

There's a fairly basic issue here, producing more doesn't mean you get to keep more.

AI can help someone do more work while leaving them no better off. Their employer keeps the saving, a client expects a lower price, or the platform takes its fee. And some of the apparent gain disappears into checking the output and finding the next customer.

Research on customer-support work illustrates both the opportunity and the distribution question. A field study found substantial productivity benefits, particularly for less experienced workers, with suggestive evidence that the system helped spread practices associated with stronger performers. That establishes a way to improve work. It leaves open how the resulting gains should be shared between the people whose experience informed the system, those using it, the employer and the supplier. [\[3\]](#)

For an individual, I want to know what's left after paying for the tools, checking the work, making corrections and winning the customer. For a company, I want to see the whole job. Five minutes saved in a demonstration is easy to like; five minutes saved at the cost of an hour's rework is another matter.

This is why I've moved away from universal income funded by ordinary personal data as the foundation of the thesis. I haven't seen enough evidence that people care deeply enough about the sovereignty of their data or their data is capable of producing meaningful, recurring income. I still care about people sharing in the value created from their lives and work. I'm less convinced by the payment mechanism I originally attached to that ambition.

My stronger starting point is expertise applied to a real customer need. A person may hold specialist judgement, a trusted relationship, access to a particular market or information that is difficult to reproduce legitimately. AI may let that person serve more customers or offer a service that was previously uneconomic. The opportunity depends on preserving something scarce as the cost of producing generic work falls. People without those advantages also need a route into earning; that is where access, learning and credible evidence of capability become important. Their need is real even where the payer has yet to be established.

I'm still interested in licensing expertise and data where someone has the right to supply it and a buyer has a reason to pay again. The same goes for services delivered by people using agents. There needs to be a customer at the other end.

There are counterexamples worth taking seriously. US Trump Accounts give eligible children a route into invested capital, including a \$1,000 government contribution for qualifying US citizens born in 2025 to 2028, which a parent has to claim rather than receive automatically. On 19 October, Musk pledged one SpaceX share per American child next year. Details as yet remain unsettled. [12][13]

Anthropic's announcement is a separate proposal: 5% of its shares would go into a public-benefit fund, subject to stockholder approval. It could support donations or distributions to particular populations. That isn't a promise of shares for every young person, and its eventual reach is still to be set. [14]

I welcome the direction. Owning a stake can give people a share of future growth, and that is a real counterexample to the idea that all the value must stay with the founders and platforms. I wouldn't dismiss the benefit to a family.

But I see these as partial responses to the immediate pressure. The US children's schemes are geographically bounded and don't answer the position of adults elsewhere whose work is changing now. Shares may compound for years, so the asset itself needn't be short term. What's missing is a settled, continuing entitlement tied to the value people contribute.

My concern is the bargain we might gradually come to accept. A small allocation today could prove a very low price for society's acceptance of decades of commercial use of our digital selves: our data, habits and accumulated judgement. Cheap, if it comes to stand in for a continuing claim on that value. I'm not saying these schemes purchase perpetual data rights; the announcements cited here do not establish that exchange. I'm saying we shouldn't let an attractive gift settle the much bigger question of who controls that use and who keeps getting paid.

What I mean by sovereignty

Sovereignty is a big word. For me it comes down to fairly practical questions: what can I authorise, what can I refuse, what can I inspect, and what can I take with me if I leave? Then, what are those rights worth in practice?

For an individual, that could mean keeping a useful record of their own expertise, directing an agent acting on their behalf, deciding what information it may disclose, and receiving agreed compensation for a legitimate use of their contribution. It does not imply ownership of everything they encounter at work.

For a company, it could mean controlling access to confidential information, defining the authority of agents and staff, retaining the knowledge needed to operate a critical workflow, and having a workable alternative if a supplier changes its service or terms.

These interests will sometimes conflict. An employee's general experience, an employer's confidential process and a client's restricted information can sit inside the same conversation. A useful product must respect those distinctions. Putting the conversation in a private store does not settle the rights to reuse it.

Consider a professional who teaches an agent how to handle difficult customer situations. If the company retains all the resulting capability after that person leaves, the system may have strengthened the company's position while weakening the individual's. That can be a legitimate contractual outcome, but it cannot be presented as evidence that everyone has gained sovereignty.

Where a business claims to improve individual participation, I will look for the mechanism: higher retained earnings, enforceable control over a contribution, portable capability, ownership or an agreed share of proceeds. A consent screen alone will not satisfy that test.

I need to keep the commercial return and the individual's position separate. An enterprise product can be a good investment without making life better for the people working there. Fair enough, but let's say so. If this thesis ends up only protecting companies, it will have lost the concern that started it.

Some arrangements may work better through professional bodies, employers, unions or other collective institutions. Their role is worth investigating because a scattered group of individuals may struggle to negotiate terms that an organised counterparty can secure. The commercial case still needs a funded buyer and a useful service.

Where companies will spend

It's easy to see how this starts inside a company. Someone finds an AI tool that helps them get their work done and starts using it. From their point of view it makes sense. The company then has to work out what information went into it, what the tool was allowed to do and who is responsible if it goes wrong.

The concern is supported by evidence. In the UK Business Data Survey 2026, 73% of businesses handling digitised data said they were uncomfortable with that data being used to train external AI models. This is a finding about attitudes to training; it should not be stretched into a refusal to use external services. Separate UK research identifies data security and accuracy as concerns while finding that concerns often do not prevent adoption. [\[4\]\[5\]](#)

That's the opening I want to investigate. Work where AI is useful enough to be worth adopting, but the customer needs better controls before they can put it into everyday use. The value is in getting the work done reliably.

The first buyer may be the person responsible for the business outcome. Security, legal and compliance may determine whether the purchase can proceed. In other cases, a specific control failure will make the security buyer the budget holder. I would establish that relationship in each market rather than assume that every opportunity belongs in a CISO's budget.

The threat model also needs to be precise. Use of information for model training, disclosure to another user, excessive access by an agent, an unauthorised action and dependency on one supplier are different problems. A contractual restriction can address one without resolving the others. Controls need to match the actual risk and the consequences of failure.

An independent supplier has to win this argument against products the customer already buys. Microsoft is extending Agent 365 across external platforms, with registry connections to AWS and Google in public preview. Independent evaluation products already offer capabilities such as testing, tracing and self-hosted deployment. The market contains capable competitors with distribution. [\[8\]\[9\]](#)

A startup needs a reason to win: a job the incumbent handles badly, evidence the customer can't get elsewhere, or an integration that makes the whole thing practical. Being independent may help. It won't, on its own, get a purchase order signed.

The harness, and what the customer actually keeps

I believe the harness matters. I mean the knowledge, instructions, permissions, tools, examples and corrections that make a model useful for a particular job. Much of the work goes into getting those things right.

I want the customer to retain the parts that contain their knowledge and define acceptable work. I also want evidence that those parts remain useful as the underlying technology changes.

Having a copy of the files is a start. What I'd want to see is the customer moving the work to another system and still getting an acceptable result, at a cost they can live with. Models behave differently, integrations break and instructions go stale. We need to try the move.

Evaluations are valuable because they give the customer a way to make that judgement. They require maintenance, independent test cases and checks against actual business results. An evaluation suite can preserve an old mistake just as efficiently as a useful standard.

There may be a good recurring business in keeping all this working as the tools change. There may also be a lot of consulting dressed up as software. Before putting money behind it, I want to know how much of the work repeats from one customer to the next and what support costs do to the margin.

Customer ownership is a product requirement. The supplier's commercial advantage must come from continuing to deliver something worth buying.

Government has a stake in this

Governments can't ignore this if employment becomes less secure, young people struggle to get started or more income moves away from labour. Tax receipts and support costs come into it pretty quickly. Equally, higher productivity could leave governments with more resources. It depends who gets the gains, how they're taxed and whether people find another way to earn.

Policy could influence that outcome through training, procurement, competition, employment protections, social insurance and taxation. It could also strengthen the position of existing technology suppliers. A policy response is not automatically a market for the businesses I would like to see built.

We can already see governments trying to help people adapt. The UK's AI Skills Boost programme targets 10 million workers by 2030. The IMF discusses training, social protection and tax changes. Useful signs of interest, although they don't amount to a commitment to personal data income or portable agents. [6][7]

For investment purposes I'd put a funded programme a long way ahead of a speech about what ought to happen. A business relying on the government needs to know who can buy, where the money comes from and how long it will take to get paid.

The same discipline applies to regulation. We need to establish which obligations attach to the customer's actual activity and what would satisfy them. Neither a broad appeal to the AI Act nor an assertion of data ownership is enough. European guidance makes clear, for example, that modifying a general-purpose model does not automatically create all the obligations of a model

provider. Data Act switching provisions also need reading against the specific service: the data a customer can take with it excludes anything protected by intellectual property rights or held as the provider's trade secret. [\[10\]](#)[\[11\]](#)

I see a useful role for policy work in improving the conditions for participation: clearer rights, credible evidence of skills, fairer terms for contributions and practical routes back into productive work. I would keep that work connected to commercial evidence, while judging its success separately from the return on an equity investment.

Where I would put my time and money

I'd start with businesses helping customers do valuable, repeat work with controls they can rely on. A clear buyer, a useful result and a reason the existing tools aren't enough. I'm more interested in getting one job right and expanding from there than funding a broad platform while it searches for a customer.

Alongside that, I want to investigate businesses that increase the ability of individuals and small firms to earn. The attraction is a visible connection between useful capability and income. The test is demanding: repeat customers, retained earnings and a reason the service remains valuable when competitors have access to similar models. Tools that produce more generic output will have to overcome a difficult pricing problem.

I'm also interested in what survives a change of tools or working relationships: the tests, permissions, knowledge and evidence that let someone keep doing useful work. There has to be a real switching problem or cost behind the demand. A sensible standard still needs people to adopt it and someone to pay.

I would pursue workforce participation and compensation arrangements through smaller experiments and institutional relationships until a repeatable buyer emerges. This keeps the individual's position in the thesis without pretending that an unresolved market is already ready for venture capital.

Across these areas, I would look for access to customers, depth in a workflow, dependable execution and an advantage that improves with use. Any claimed data advantage must be based on rights the company actually holds. Any network effect must show how one participant makes the product more valuable to another; a growing database alone is insufficient.

The technology choice follows the problem. Privacy-preserving techniques, local processing, identity systems and shared ledgers may each be useful. Blockchain has a case where several parties need a shared record or settlement mechanism and cannot rely comfortably on one operator. It still has to beat a simpler arrangement on cost, usability and enforceability. A token does not create demand for the underlying contribution.

I will pass on businesses whose income case relies on ordinary personal data acquiring an unexplained price, that promise protections they cannot enforce, or whose principal feature is already adequate inside the customer's existing bundle. I will also pass on a service business priced as software without evidence that delivery can become repeatable. In every category, the business must operate within rights genuinely agreed with the people and organisations whose contributions it uses.

Over the next five years

In the first year I'd spend most of the effort proving demand in a few specific workflows. Paid trials, a proper comparison with what customers already use, and all the implementation costs included. Building something for myself will teach me a lot. It won't tell me whether an enterprise will buy it.

In years two and three I'd put more time and capital behind the businesses where customers come back, delivery costs hold up and the product earns a bigger place in their work. We'll also have seen how the incumbents respond. If the agreed tests keep failing, I don't want to fund another explanation for why demand is just around the corner.

In years four and five, the question becomes whether the company can maintain its position through changing models, platform terms and regulation. Expansion, acquisition or liquidity may be possible. None should be assumed simply because the thesis has a five-year horizon.

At every stage I would test three futures. Under slower adoption, the company must justify itself through present-day productivity or control. Under uneven disruption, it must serve the occupations and businesses experiencing the change. Under rapid displacement, it must show that its buyer retains the ability to pay. A consumer business can face its strongest social need at precisely the moment its customers lose purchasing power.

Time and capital need different decisions. Some research, policy work or advisory activity may be well worth doing without becoming a venture investment. I don't want the thesis to force everything I care about into one return model.

What would make me change my mind

I have to be willing to be wrong about this. Otherwise it's a story I like, rather than something I can use to decide where to invest.

If people use a product consistently and still earn no more after costs, that weakens the individual-income case. The same is true if customers and platforms take almost all the gain. Lots of usage isn't enough.

On the enterprise side, I'd worry if paid trials keep turning into custom projects, existing suppliers already do the job, or the checking costs eat the savings. A compliance concern needs to turn into a purchase at some point.

For portability, I require a real move between systems with acceptable outcomes and a measured cost. If customers rarely need to move, or can already move cheaply, the proposed premium will be difficult to sustain.

With government-dependent opportunities, I want to see support become funding, authority or an actual purchase. I'm not going to keep moving the date because I like the idea.

I will review these signals quarterly and revisit the full thesis annually. The important measures are retained income, customer outcomes, renewal, implementation effort, margins and the position of competitors. Funding announcements and model benchmarks provide context; they cannot substitute for those measures.

Where that leaves me

I still believe the way we organise AI's use will have a lot to do with who benefits. The models matter, of course. So does who can use the knowledge, reach the customer and decide how the money is shared.

That's where I want to spend my time. Helping people find a stronger position from which to earn, and helping companies use AI without losing control of the things their business depends on. Sometimes those interests will line up. Where they don't, we need to be honest about it.

There should be good businesses here. I want to find the ones where the customer's gain is clear and the supplier can keep earning its place. That's the version of Agents of Sovereignty I'm taking into conversations with founders, customers and investors. Still plenty to test.

Evidence notes

This version draws on the research discussed in October 2026, with the US equity and public-benefit announcements added on 21 October. Every cited source was checked against my primary document on 23 October 2026. I am human, I may make mistakes & I may have thoughts and opinions you don't agree with. These are likely driven by my own personal experiences and humanity.

[1] Stanford Digital Economy Lab, *Canaries in the Coal Mine?*, revised August 2026. US payroll evidence; descriptive indicators, with limitations acknowledged by the authors.

<https://digitaleconomy.stanford.edu/publication/canaries-in-the-coal-mine-six-facts-about-the-recent-employment-effects-of-artificial-intelligence/>

[2] Anders Humlum and Emilie Vestergaard, Large Language Models, Small Labor Market Effects, September 2025. Danish adoption surveys linked to administrative records; limited to the population and period studied.

https://bfi.uchicago.edu/wp-content/uploads/2025/04/BFI_WP_2025-56-4.pdf

[3] Erik Brynjolfsson, Danielle Li and Lindsey Raymond, Generative AI at Work, NBER working paper, subsequently published in 2025. Evidence from a specific customer-support setting; it does not establish economy-wide income effects.

<https://www.nber.org/papers/w31161>

[4] UK Business Data Survey 2026. Attitudes to external AI training among businesses handling digitised data; not a willingness-to-pay study.

<https://www.gov.uk/government/statistics/uk-business-data-survey-2026/uk-business-data-survey-2026>

[5] UK AI Adoption Research, updated February 2026, based on 2025 fieldwork. Reported adoption, benefits and barriers; the survey does not capture shadow adoption.

<https://www.gov.uk/government/publications/ai-adoption-research/ai-adoption-research>

[6] IMF, Broadening the Gains from Generative AI: The Role of Fiscal Policies, June 2024. Policy analysis and recommendations, rather than an enacted programme or forecast of certain unemployment.

<https://www.imf.org/en/Publications/Staff-Discussion-Notes/Issues/2024/06/11/Broadening-the-Gains-from-Generative-AI-The-Role-of-Fiscal-Policies-549639>

[7] UK AI Skills Boost explainer, January 2026. Government-industry programme and target; not evidence that the target has been achieved or that participation increases earnings.

<https://www.gov.uk/government/publications/ai-skills-boost-explainer>

[8] Microsoft, Agent 365 capabilities and integrations, May 2026. Agent 365 itself is generally available; the AWS Bedrock and Google Cloud registry sync is in public preview, with lifecycle governance over those agents still to come.

<https://www.microsoft.com/en-us/security/blog/2026/05/01/microsoft-agent-365-now-generally-available-expands-capabilities-and-integrations/>

[9] LangSmith evaluation product documentation. Evidence of an existing competing offering, rather than independent proof of performance or customer outcomes.

<https://www.langchain.com/langsmith/evaluation>

[10] European Commission, Guidelines on obligations for General-Purpose AI providers. Modification makes an actor a provider only where it uses more than one third of the original model's training compute, and even then the obligations are limited to documenting the modification.

<https://digital-strategy.ec.europa.eu/en/faqs/guidelines-obligations-general-purpose-ai-providers>

[11] European Commission, Data Act explained. Switching and egress charges end on 12 January 2027; the exportable scope covers input and output data and metadata but excludes material protected by intellectual property rights or constituting the provider's trade secret.

<https://digital-strategy.ec.europa.eu/en/factpages/data-act-explained>

[12] IRS, Trump Accounts. Eligibility and the \$1,000 pilot contribution; not a universal payment to all children, and the contribution must be elected on Form 4547 even where Treasury has opened the account automatically.

<https://www.irs.gov/trumpaccounts>

[13] Associated Press, 19 October 2026. Musk's proposed SpaceX share gift; announcement rather than completed distribution. The pledge, its date and its unsettled details are corroborated by several outlets carrying the AFP report of 19 and 20 October 2026.

<https://apnews.com/article/b3312da37a0e340be275295073055e00>

[14] Anthropic, Investing in the public benefit, 18 October 2026. Proposed equity commitment, approval condition and potential uses; distribution arrangements remain to be developed. The announcement also provides for a further 2% at the board's discretion for every 50% increase in company value above a baseline.

<https://www.anthropic.com/news/public-benefit-fund>