


Redburn Review

SEPTEMBER 2026



Rothschild & Co | Redburn

Toll Operators

Rocket Men

Octopus Farming

SEPTEMBER 2026

Contents

Markets

Toll Operators 4
Why the Ratings Duopoly Endures
Charles Bendit

Pigouvian Taxes 6
The Road to Regulating AI
Melissa Davies

Three Conditions 8
A Not-so-‘Super App’
Joseph Barker

Rocket Men 10
The Problems with Space Law
Katherine Stamp

Opportunity 13
AI and Quantum in Drug Discovery
Simon Baker

All That Glitters 16
Gold’s Fluctuating Monetary History
Michael Taylor

Fanfiction 18
The Rise of the Music Superfan
Aleia de Buretel

Tidal Wave 20
The Risks of AI
Omar Sheikh

The Global Retail Bank 22
From Empire to Platform
Harry Bartlett

Diversions

Screen Time 24
Detecting Cancer
Jamie Clark

One Signature or Five? 26
The Reality of Enterprise AI Adoption
Rowan Adley

Regulatory Friction 28
The Road Ahead for Light Vehicles
Tobias Beith

Uphill Climb 30
The Democrats’ Senate Problem
Fred Hill

The People’s Currency 32
Liberalising the Renminbi
James Invine

Tech Talk 34
AI in Legal Workflows
Henry Hayden

Letter from America 38
Hidden Heroes of Gameday
Connor Cheynet

United 40
The FIFA World Cup
Daniel Sepahi

Octopus Farming 42
Our Diet, Their Intelligence
Neil Tyler

The Serpent in the Grove 44
AI in Fiction
Archie Cotterell

Obituary 46
Sir Garfield Sobers

Editor Archie Cotterell **Production Editor** Jon Sever
Digital Director Joe Shadafza **Sub-Editor** Georgia Martin
Design Consultants Phable **Photos** Getty Images

Editor's Letter

In the mid-1980s, at the height of the Cold War, the novelist Martin Amis told his father Kingsley, also a novelist, that he was planning to write a book about nuclear weapons. Kingsley replied, “I suppose... you’re against them?”.

Kingsley Amis was a scourge of *bien pensants*, and as investment advisors we should be too. Fortunately, Rothschild & Co Redburn has a healthy tradition of challenging orthodoxies and received wisdom.

In this edition, Melissa Davies proposes Pigouvian taxes to ensure the AI roll-out is not winner takes all, Joseph Barker rejects the idea that ‘Super Apps’ will become established in the West as they have in China, and Katherine Stamp draws attention to the shortcomings and inconsistencies of space law – good news, as ever, for the lawyers. Elsewhere, Omar Sheikh fears confirmation bias and the outsourcing of critical thinking by those using AI.

Among the Diversions, the same principle applies. Neil Tyler worries about octopus farming, and Connor Cheynet looks behind the glamour to shine a light on the unsung heroes of every sports event – the ground staff. As if to prove we write about subjects most investment banks do not, Aleia de Buretel outs herself as a Lana Del Rey superfan, the singer’s melancholic Americana offering a welcome antidote to the current geopolitical noise.

Some might argue there is a thin line between being a scourge of *bien pensants* and merely trying to *épater les bourgeois*. We, as you might hope, would beg to differ.



impossible to dislodge, for four reinforcing reasons. First, global protocol network effects: ratings function as a shared language, and because crowded information marketplaces are inefficient, issuers and investors coalesce around one or two conventions, with familiarity breeding trust. Second, historical context: information becomes more valuable the more context it carries, and both firms hold over a century of it. Third, ecosystem entrenchment: issuer processes, investor mandates and regulatory workflows are all built around the incumbents.

The fourth, and perhaps most underappreciated, is a principal-agent dynamic. The party buying the rating (the issuer) is not the party consuming it (the investor). The investor in turn is usually an agent, a portfolio manager investing other people's money. That manager has every incentive to rely on a rating the market trusts rather than underwrite the judgement of an unfamiliar upstart, because the established benchmark is the career-safe choice. Issuers, knowing this, buy the ratings their investors expect to see, and regulators then codify those same ratings into capital rules and investment mandates. The result is a self-reinforcing loop in which trust, demand and regulation point to the same two names. It explains why greater methodology disclosure and a longer roster of approved agencies (Nationally Recognized Statistical Rating Organizations, or NRSROs) changed so little.

Hence the market is better described as a duopoly-plus-one than a level playing field. S&P and Moody's each hold roughly 40% of the global market, with Fitch a distant third at c12-15%. Fitch has made incremental gains in selected categories (EMEA, structured finance) but not in ways that affect industry pricing. Morningstar DBRS is strongest in specific niches (Canada, European structured finance) but remains small. Regional agencies operate mainly in local markets like China and Japan.

Why the Ratings Duopoly Endures



Charles Bendit
Capital Markets
Research

The global Credit Ratings industry remains one of the most durable and economically attractive segments within the Financial Services sector. At its core, a rating agency is a toll collector on global debt issuance, a c\$350trn market that has expanded dramatically relative to global GDP over the past four decades. Its fundamental attributes (regulatory entrenchment, network effects, untapped pricing power, high fixed-cost operating leverage) remain intact, supporting mid-single-digit revenue growth and structurally high margins. The important questions are how has such a lucrative franchise resisted competition and what, if anything, could erode it?

The industry's origins lie in the American railroad boom. In 1860,

Henry Varnum Poor published *History of the Railroads and Canals of the United States*, the first serious attempt to arm investors with standardised data on these nascent industries. John Moody followed in 1909 with *Analyses of Railroad Investments*, compressing that analysis into letter grades, and Moody's Investors Service was incorporated in 1914. Poor's Publishing merged with the Standard Statistics Bureau in 1941 to form Standard & Poor's (S&P).

For decades the agencies sold their manuals to investors, but in the 1970s the model inverted and agencies began charging issuers of debt rather than buyers. That 'issuer-pays' shift, combined with growing regulatory reliance on ratings, converted two subscription publishers into toll booths embedded in the plumbing of capital markets.

The agencies remain structurally advantaged. S&P and Moody's have consistently rated c80% of global debt issuance and are virtually

Most tellingly, the structure survived its own near-death experience. The agencies were central to the 2008 global financial crisis, having stamped AAA ratings on subprime structured products which later collapsed. Yet the reforms that followed (Dodd-Frank, the removal of hard-wired regulatory references to ratings, the expansion of the NRSRO roster and compulsory methodology disclosure) produced no material shift in market share. Regulation designed to broaden competition instead codified the incumbents' role.

One further property sets the rating agencies apart, even among data businesses: they tend towards a global standard. Investment manager YCG draws an instructive contrast with consumer credit bureaus. Bureaus collect proprietary data on individual borrowers and sell it to banks, but banks lend mostly at home, because local knowledge matters in lending and because implicit and explicit state guarantees produce regulation that discourages cross-border expansion. A US bureau's data is invaluable to a lender in Texas and largely useless to a Chinese bank lending in Shanghai, so bureaus must be rebuilt market by market, and the result is a patchwork of country-specific winners.

Capital markets, by contrast, are global, and most global investors take more comfort from the shared language of the Big Three than from a less familiar domestic challenger. An Indian company issuing bonds into the global pool is therefore strongly incentivised to buy a Big Three rating. The economics of bond ratings should thus produce a handful of global winners rather than national ones – exactly the market structure that has emerged.

What makes this such a great business is the combination of moat and economics. The moat manifests, above all, as untapped pricing power. S&P and Moody's have raised prices 3-4% annually for decades and an underappreciated value gap affords them further headroom: a rating saves an issuer an estimated 30-

50bp in annual interest costs while costing perhaps 7-8bp per agency (14-16bp for both). A Moody's study of a hypothetical five-year \$1.5bn corporate bond found its rating cut the option-adjusted spread 22%, or 39bp. That gap implies pricing could rise c3-4% a year for more than two decades before the price charged matched the value delivered. The economics beneath the pricing is equally rare: capital intensity is minimal, margins sit structurally above 60%, and almost all profit converts into cash.

The medium-term outlook is positive. Two secular drivers endure: rising overall debt issuance (a function of GDP growth, spreads and rates), and the long-run migration from bank financing to capital markets in search of cheaper funding and broader investor bases.

Nearer term, a multi-year refinancing cycle dominates. Four-year cumulative refinancing needs across US and EMEA non-financial corporate debt now total \$5.2trn, from \$4.9trn in 2024 and \$3.8trn in 2019, a c8% CAGR since 2015. Add an estimated 40-45% rebound in M&A-related issuance and the vast, substantially debt-funded build-out of data centre and power infrastructure (an estimated c\$3trn industry-wide), and the issuance backdrop looks unusually supportive once monetary conditions stabilise.

The risks are, generally, cyclical rather than structural. Extended periods of wide credit spreads defer issuance, and a sudden risk-off shock can depress volumes sharply, though neither is permanent. Private credit, now c\$2trn, poses a more structural question, since middle-market loans funded by private vehicles can bypass public markets and the need for a public rating.

Even here the threat is fading, as the asset class becomes steadily more ratings-penetrated driven by demands for transparency from limited partners, the prospect of tighter regulation and rising insurance allocations to alternatives.

Any stress in private credit is more likely to consolidate demand around the two incumbents than to disintermediate them. Reputational and regulatory risks persist, but the post-crisis experience suggests both are survivable. AI, often cast as a threat, is more likely a benefit. The moat rests on the rating's status as a reference standard, not on the content of any individual report, so large language models mostly promise to automate analyst workflows and lift incremental margins across the next cycle.

The more profound risk is macroeconomic, one the moat cannot defend against. The agencies are a toll on the flow (and stock) of debt, so their fortunes ultimately rest on the world's willingness and capacity to keep borrowing. Since 2010, global debt has risen from c\$200trn (286% of GDP) to c\$350trn (308% of GDP). Debt can outpace GDP for decades, but not indefinitely, because the revenues available to service it are finite. History shows that when the burden becomes too heavy it is eventually worked off, through some combination of austerity, restructuring and inflation, as it was across much of the developed world between 1950 and 1980.

Should the macro mix now shift in that direction (labour regaining bargaining power, interest rates grinding structurally higher, trade turning more protectionist), the equilibrium debt burden could move lower, and global debt-to-GDP revert towards a reduced range over the coming decades. Such a deleveraging cycle would not break the duopoly, and refinancing alone would sustain base-level demand.

But a smaller stock of debt growing more slowly would reset the industry's growth algorithm, from mid to high single digits towards mid, or even low to mid single digits, and with it the rate at which these remarkable compounding machines create value. That, rather than any competitor or new technology, is the risk deserving the most attention.

Pigouvian Taxes

The Road to Regulating AI



Melissa Davies
Economics

Artificial Intelligence (AI) displacement is already having an impact on labour markets, visible primarily in slowing wage inflation. As we argued in ‘Rage Against the Machines’ (April 2026), if AI creates a ‘manufacturing-level’ worker displacement event for parts of the services sector (which is quite a cautious scenario given the possible range of outcomes), the political and policy implications could be deep and long-lasting.

The report discussed various public policy responses to a situation where AI creates ‘more producers that don’t consume’ and demographic shifts create ‘more consumers that don’t produce’, leaving the government

balance sheet to bridge the gap in some way. Oft-discussed policy interventions include taxing labour income less and capital more, taxing consumption, wage subsidies and even AI token or compute taxes.

This essay digs into the feasibility of an AI token tax (or, alternatively, an AI-permitting system akin to the EU ETS) and will show the difficulties of implementing such a policy in a world of sharply falling token prices. Perhaps more feasible, and easier to facilitate, are a ‘reverse payroll tax’ or even a floor for token pricing – a ‘minimum wage’ for compute.

The major problem with a token tax is, ironically, that tokens are too cheap to tax and declining in cost by the day. The cost of a million tokens from a frontier model has fallen from around \$30 in early 2023 to roughly \$2.50 today, while the cost of ‘good-enough’ GPT-4-class output has dropped to around \$0.10-0.15 – a fall of roughly 200 to 300 times in three

years. A tax levied per token, or per unit of compute, is therefore a tax on a base that is evaporating.


Even as token volumes explode – Google alone went from processing 9.7 trillion per month in 2024 to over 3.2 quadrillion by 2026 – the dollar value of that compute, while rising, remains small relative to the economic activity it enables. If the point of the exercise is to raise meaningful revenue or to change behaviour, a compute tax fails on both counts.

The real anxiety is not about compute per se though, it is about the impact of AI adoption on workers: large employers using cheap models to displace employees.

Consider a set of large employers – a national retailer automating customer service, a bank automating back-office work, and an outsourcer automating its agents. These examples are situations where AI is likely to prove highly substitutable with labour, rather than value-enhancing. Displacing \$140-450m of annual payroll might require only \$0.1-0.6m of cheap compute. A compute or token tax could perhaps charge such a firm \$0.25-2m, even if set at a very high rate. Token taxes are likely to under-price the actual social harm by many multiples. A permit tied instead to the displacement itself – the reduction in headcount, priced at a share of the wages removed – would link the levy to the externality it is meant to address.

A headcount-based permit measures only one of the two channels through which AI suppresses labour income. Displacement is not only about layoffs. AI also erodes the bargaining power of workers who keep their jobs, allowing employers to freeze or cut real pay. The harm – a falling share

The grapes of wrath... growing heavy for the vintage



“Displacement is not only about layoffs. AI also erodes the bargaining power of workers who keep their jobs”

of income going to labour – shows up as much in the wage rate as in the headcount.

The total wage bill captures both. A firm’s compensation cost is simply headcount multiplied by pay, so a fall in that bill decomposes into a headcount effect (fewer workers) and a wage-rate effect (lower pay). A permit based on the shortfall in the total wage bill – measured against a baseline indexed to the firm’s output growth and to economy-wide pay – encapsulates both.

A firm that quietly freezes wages while holding headcount would owe nothing under a headcount rule but faces a real charge under a wage-bill rule; a firm that grows and pays the going rate owes nothing under either. The wage bill is also the more practical base: total compensation is already reported through payroll data, it is harder to game than a headcount base, and it aligns the charge with the social cost.

If a permit market feels too elaborate, there is a blunter tool: a minimum price for compute, pegged to wages – in effect, a minimum wage for machines. AI displaces labour because compute is cheaper than the workers it replaces; a floor set as a fraction of the human-labour cost of the equivalent task narrows that gap directly, so automation happens where AI is genuinely better, not merely where it is artificially cheap. Because the floor is indexed to wages, it also neutralises the deflation problem – it does not fall as compute costs collapse. It is radically simpler than a permit market: no cap, no auction, no measurement of who displaced whom. A minimum price for compute could be levied upstream on the handful of large cloud providers, clawing back the

gap between the floor and true cost to a worker-transition fund (rather than leaving it as a windfall to the hyperscalers), and also applied to imported compute so buyers cannot simply route offshore.

Broadly speaking, because compute is of the order of 1,000 times cheaper than the labour it substitutes, even a floor set at just 10% of the human-labour cost of a task could imply a compute price of roughly \$100 per million tokens – some 200 times today’s market price. For tokens to be at parity with labour cost, their price would have to be raised over 1,000 times relative to the current market. But that would not be the ultimate objective, given the many benign and socially helpful uses for AI. The point would be to tip the scales in favour of labour to slow displacement and employment pressure.

A price floor cannot tell the difference between a labour-displacing use of AI and a beneficial one. Being a price rather than a quantity instrument, it blunts the incentive to substitute without ever fixing the quantity of displacement or taxing the harm directly. It is best understood not as a substitute for the wage-bill levy but as a simple backstop that sits underneath it, a rule that compute may never be sold far below the labour it replaces.

A possible plan could look like this: the primary policy instrument based on the output- and inflation-indexed shortfall in a firm’s total wage bill as AI is introduced; restrict coverage to large employers, mirroring existing employment-law thresholds; include contractor and offshored labour cost so the base cannot be dodged by moving work off the payroll; and recycle the revenue into a worker-transition fund.

A quantity-capped, auctioned permit market, as in the EU ETS, could be used, which auctions allowances with a rising floor and a cost-containment ceiling. Underneath, a modest wage-linked price floor on compute could serve as a backstop against the raw cost-arbitrage. The quantity-cap design appeals if one believes the social damage of rapid, concentrated job losses is threshold-like – where fixing the quantity and letting the market discover the price beats trying to set the price directly.

The counterargument is that the rapid productivity gains from AI adoption will, ultimately, enrich workers via stronger growth and income acceleration. But the historic record argues against this for workers displaced by technology – while the overall economy grows and profits expand, workers can be ‘left behind’ by tech disruption. It seems unlikely that the unfettered advent of AI is going to reverse the post-war trend of a steadily declining wage share in favour of the capital share.

Inevitably, though, there is little chance of public policy interventions taking place before AI disruption is entrenched, by which point the organisational changes will have taken place and displaced workers will have to seek work elsewhere. If, as we expect, one of the consequences of AI disruption is a prolonged period of wage disinflation and profit share expansion, *ex post* traditional tax interventions seem impossible to avoid as income tax revenues are hollowed out. The labour market externality created by extremely cheap compute and its substitutability with labour is unlikely to be internalised in time.

Three Conditions

A Not-so-‘Super App’



Joseph Barker
*Internet
Research*

The Super App is often touted as next in a long line of Western imports from China. We have taken tea, silk, gunpowder, paper currency and the compass among other things, and each is now woven so deeply into Western life its origin is forgotten. The flow has not stopped in the 21st century, with the most recent arrivals perhaps the most consequential. This underpins the belief a Western Super App is a ‘when’ not an ‘if’. We disagree.

Expectations hinge on the extraordinary success of recent Chinese imports. Alongside cheap quick-commerce, TikTok has reshaped how the West spends its time, potentially rewiring our brains. Vertical video existed, but China pioneered a feed built on predicted ‘interests’ rather than social connections, taking the ‘social’ out of social media. It has become so pervasive so swiftly it is now perhaps the most socially sanctioned drug on earth, given its widespread adoption by children.

Short-form video is not a product, but a utility with universal appeal. China perfected entertainment at home, and despite cultural differences the Western consumer proved just as susceptible to the one-size-fits-all format for maximising enjoyment from time spent watching video. The level of success for a format with both Western and Chinese consumers opened the door for people to expect the next Chinese conceptual

consumer staple, the Super App, another utility-enhancing format that should be culturally agnostic, to be equally pervasive.

The incentive is huge, as Super Apps run China. Despite Communist Party governance, the country’s society and economy are driven by private corporations with culture, commerce, communication and almost everything else flowing through this handful of apps.

WeChat, used by 97% of China’s internet users, began as a plain messaging service, added payments before Apple Wallet existed, and now runs an app store of third-party ‘mini programmes’ offering roughly twice as many apps as Apple. WeChat is now so gravitational that over a third of other Super Apps (e.g. JD.com and Meituan’s) traffic flows via it, as if Amazon drew a third of its usage from WhatsApp.

Zooming out, Super Apps are the starting point of every online consumer journey and the default way to pay in the physical world, so the whole economy is tied to these platforms. With all commerce either disintermediated by or dependent on them, the revenue opportunity and thus incentive to build a Western equivalent has been strong for years, sharpened by TikTok’s proof that cultural difference was no barrier to adoption of Chinese mobile formats.

Indeed, for years the best evidence a Western Super App was coming was that Apple remained determined to stop them. Behind closed doors, its board described Super Apps on iOS as an “undifferentiated user experience on a super platform” with the potential to become a “major headwind” to iPhone sales.

In 2024, the Department of Justice sued Apple over precisely that intent.

But in November 2025, Apple quietly reversed course. Its Mini App Partner Program lets Super Apps run in-app stores at half Apple’s usual app toll – 15% rather than 30%. Thus, the single greatest structural barrier to the Western Super App had fallen.

So why has no-one moved since? This non-event is stranger still against the artificial intelligence (AI) backdrop. AI has made software interfaces, which is all a Super App really is, far cheaper to build, while creating new candidate platforms.

Our view is the Super App expectation has always been, and remains, misguided, and any Western

Legacy shoppers



replica is challenged both structurally and culturally. Looking closely at China in the 2000s, we find three conditions unique to time and place that made the Super App feasible, without which it would not exist. Lacking those conditions in the West, that ship has now sailed.

China's Super Apps hinge on that point in time when the smartphone was born. First, the West came to mobile dragging decades of desktop habits, so dominant apps (Amazon, Yahoo, Facebook, YouTube) were websites ported onto a smaller screen. China had little such baggage, with barely one in ten Chinese people using the internet at that stage. With no habits to change, and consumer habits malleable, whoever designed mobile-first had the right to win.

Second, there was no established brand trust. The West had spent half a century building beloved national

names, and that trust continued into the mobile era. The speed of industrialisation meant this hadn't happened in China, allowing trust to form around the first mobile platforms. Alipay was launched pre-mobile in 2004 for exactly this purpose, to solve "the lack of trust in transactions between strangers". Consumers came to rely on these platforms to do almost anything.

Third, there was no payments precedent. The West's deep credit-card penetration meant mobile commerce simply adopted the existing rails to ease the transition for consumers. China was still largely a cash-and-debit society, so Super App mobile wallets (payment infrastructure designed around the phone itself) had little barrier to becoming the dominant way to pay across China.

Free of legacy consumer habits, these apps could design natively for mobile, install themselves as the trust layer between consumer and brand, and own the nation's payment rails. But without this set-up in the West, it is unlikely to be recreated. We cannot un-know the PC, un-learn brand trust, or stomach such an overhaul of payment infrastructure all at once. The door that was open in China for perhaps five years is bolted shut everywhere else.

A multitude of Western failures bear this out. Over the past decade, from PayPal to Meta's in-app checkout to TikTok Shop, every attempt to create an 'everything-app' has fallen short.

The purest case is Uber, which has folded both its leading ride-hailing and delivery apps into one core space, the most obvious Super App potential on offer in the West, yet consumers continue to use each separately. If convenience were going to move them, it would have; if a shock were needed, the pandemic, where we all lived on ride-hailing and takeaways, would have been it.

It is clear once habits form around specialised point solutions, the bar a new format must clear to alter habits is extraordinarily high. A Super App

offering little convenience over the device-level OS it would replace, without China's favourable set-up, has little chance of clearing it.

How then did short-form video become so popular? Because it asks nothing of the user but to watch and scroll, sweeping the West so completely it forced Instagram and YouTube into Reels and Shorts. It demands little and offers a lot. An ecosystem is a different proposition, requiring the user to forego a dozen familiar tools for one unfamiliar one, a necessary act of relearning for which the average consumer has proven to have near-zero tolerance.

AI optimists will counter that the conversational layer offered is as important a new technology era as the dawn of the smartphone, with potential to capture the start of all consumer journeys. This is a new position from which to branch out into a Super App. In practice, as with the West's first move into mobile, it is likelier to entrench the incumbents than crown new winners, given how dominant each vertical's players have become. Amazon should win 'by default', because AI will want to deliver best price and fastest delivery.

None of which is to say AI changes nothing. It is dismantling businesses in plain sight: US educational technology company Chegg, whose one-size-fits-all product was rendered redundant by a free AI tutor, has lost 99% of its value. Where AI offers a real convenience uplift consumers will adopt it, but for a Super App alone the convenience uplift over an already-dominant point solution remains modest.

Silicon Valley has tried and continues to try to build the West's Super App, but the plethora of failures and starkly advantageous landscape in China in the 2000s indicate it will remain a concept. The Valley over-indexes the willingness of consumers to take action and too easily forgets its past. The future of mobile use in an AI world is likely to look much the same as it does today.





Rocket Men

The Problems with Space Law



Katherine Stamp
Aerospace & Defence Research

The International Energy Agency projects that data centres will account for nearly half of US electricity demand growth between 2025 and 2030. Space presents a way to accommodate this increasing demand, offering abundant solar energy, a way to eliminate the need for terrestrial cooling water and freedom from land constraints. So, companies are taking data centres off-Earth: Nvidia-backed Starcloud has already put an H100 GPU in orbit and plans a 5GW orbital data centre, while Google's project Suncatcher aims to network solar-powered satellite constellations carrying its TPUs. The ambition is international. China launched twelve Xingshidai computing satellites in 2025 while the European Union's ASCEND project has examined industrial-scale orbital data centres.

Despite raising several questions on overlapping jurisdictions and liability law, the real legislative difficulties emerge when we consider how such projects scale. Specifically, on legislation surrounding property and appropriation. This essay argues that the current international framework is not fit for this next stage of commercialisation. In particular, the Outer Space Treaty (OST) does not explain how its prohibition on appropriation can coexist with resource extraction that necessarily requires some degree of exclusive control. Two futures remain: a reprise of 19th-century colonialism or a reformed multilateral framework. The former seems more likely.

The Low Earth Orbit centres are hardware launched from Earth and raise no immediate appropriation question. However, the next step could require the use of lunar resources. NASA, for example, is actively developing technologies for producing oxygen on the Moon so spacecraft do not need to carry all their propellant from Earth. This is

commercially attractive because Earth has a deep gravity well compared with the Moon: launching heavy material from Earth requires far more energy and fuel than moving the same material from the lunar surface.

China is pursuing a comparable strategy. Its Chang'e-8 mission, scheduled for launch in 2029, will conduct *in situ* resource-utilisation experiments near the lunar south pole as groundwork for the International Lunar Research Station. This includes work relevant to manufacturing construction materials from lunar regolith.

A second scenario involves building data centres or industrial sites directly on the Moon. Lonestar is already experimenting with the concept, having demonstrated lunar data storage in 2024 and flown an 8 TB edge-computing payload in 2025. In this camp we also see Elon Musk, whose SpaceX prospectus cites the intention 'to establish lunar-based manufacturing capabilities... to produce large-scale AI compute satellites'.

In both scenarios knotty legal issues arise. The first question is whether a company can legally own the resources it extracts from outer space. To answer this, we must turn to the Outer Space Treaty. Created in 1967 and ratified by 117 states, it is considered the constitution, or magna carta, of space law, stipulating a series of fundamental principles. Regarding property, Article I proclaims that the Moon is 'free for exploration and use by all States'. Arguably 'use' could include using it for its resources. However, Article II creates the central difficulty: 'Outer space, including the moon...is not subject to national appropriation by claim of sovereignty, by means of use or occupation, or by any other means.'

Article VI brings private organisations into this framework. The Treaty outlines that non-governmental entities and their activities must be compliant and remain subject to continuing supervision by the appropriate State party to the Treaty. Yet the Treaty is ambiguous on property rights, leading to the development of two interpretations.

A restrictive interpretation sees private resource rights amount to appropriation and thus renders resource extraction forbidden *in toto*. This stance is adopted, for example, in the 1979 Moon Agreement, which holds that neither lunar resources in place nor parts of the lunar surface or subsurface may become the property of states or international persons. None of the major spacefaring states have ratified this.

In contrast, a permissive interpretation prohibits sovereignty over lunar territory and ownership of resources *in situ*, but permits ownership once they are lawfully extracted. National laws have adopted the latter interpretation. In 2015, Obama signed an Act which lets US citizens (and companies) possess, own, use and sell space resources 'obtained', 'subject to authorisation and continuing supervision' by the

federal government. Meanwhile, Luxembourg in 2017 declared flatly that 'space resources are capable of being owned'.

Even if we accept the permissive interpretation, the thornier problem lies in the matter of accessing the resource. Picture a company discovering economically viable ice deposits at the lunar south pole, building a mine and processing it. Mining could require infrastructure to be in place for decades. A company simultaneously needs to stop another company excavating the same ice deposit and landing beside its machinery. Yet, the right to exclude for long periods of time looks remarkably similar to the 'appropriation...by means of use or occupation' which the OST explicitly condemns.

The challenge of resource access is amplified by scarcity. The south pole of the Moon is viewed as particularly attractive because it has permanently shadowed regions potentially containing ice, highly illuminated ridges useful for solar generation and areas particularly suited to landing. Thus, both the US and China are concentrating their lunar ambitions on the south pole. The US Artemis programme aims to establish a sustained presence there, while China plans a crewed lunar landing by 2030 and is developing the International Lunar Research Station with Russia.

Current workarounds do not present concrete solutions. This is first evident in the Artemis Accords (2020), a non-binding political agreement between participating states setting out principles for how they intend to conduct activities on the Moon and other parts of space. Section II of the Accords talks about the establishment of 'safety zones', defined as areas within the limits of which notification and coordination are used to avoid 'harmful interference' with another actor.

The Accords do not formally confer a right to exclude others and expressly preserve free access to celestial bodies. Yet, the rights they

appear to confer are still difficult to distinguish from property rights. Where a safety zone remains in place for years around a valuable site, coordination may begin to resemble control. A mechanism intended to prevent harmful interference could thereby create *de facto* exclusive possession. This is precisely the result that Article II's prohibition on appropriation by occupation seeks to prevent.

Even if one ignores the contestable interpretation of the OST the Accords impose, they are by no means an international solution. First, this is not a treaty nor law. Therefore, the Accords create no independently enforceable international obligations or enforcement mechanism. Second, while seventy states have signed, these signatories do not include China and Russia, two of the largest spacefaring states. When the Accords were announced, China's Foreign Ministry said the country "supports discussion on a legal system for space resource exploration on the UN platform", implicitly not signing the Accords owing to their US-led creation.

This leads us to potential for resolution via COPUOS, the United Nations Committee on the Peaceful Uses of Outer Space. The body, which contains all major spacefaring states, is in the process of negotiating a set of principles to be voluntarily adhered to by 110 states.

The April 2026 draft comprised two parts: part one covered consensual points while section two covered contentious points. Part two maintains extraction is compatible with the OST only if a series of conditions are met. This includes preserving free access, avoiding harmful interference, respecting due regard, aligning with sustainability expectations and remaining subject to authorisation and continuing supervision. This seems a heavy load to bear for companies wishing to extract some ice. Extraction and free access especially do not seem coterminous.

There also remains international divergence even surrounding the property claim. Russia's and Iran's submissions, among others, have for example emphasised equitable sharing of benefits.

While the case of data centres and the Moon has been used, the same difficulties arise when considering the mining of asteroids or Musk's hope to set up a Martian settlement. However, commercial exploitation is unlikely to wait for international law to reach a perfect solution. Nor is a vision of equitable sharing of benefits from the Moon likely to prevail.

The failure of the 1979 Moon Agreement illustrates this. It said the Moon and its natural resources are the 'common heritage of mankind' and that resources in place cannot become property. But only seventeen parties signed it, and the US, China and Russia were not among them. Countries are not willing to give up profits so hastily.

So, we are left with the conclusion that the current international legal framework for space is outdated. Nor are current attempted workarounds faring well. This is because, much like any founding constitution, the OST was melded to address contemporaneous problems. In 1967, the stage was dominated by state actors hoping to break scientific barriers; now SpaceX's June 2026 valuation of \$1.77trn is more than seventy times NASA's entire FY2026 budget of \$24.4bn.

Yet major reform may prove difficult. States are unlikely to dismantle the basic principles of the OST lightly – it is difficult to make more than 110 states agree to a set of rules. This is also why COPUOS does not offer an obvious route to agreement. Its work on space resources remains consensus-based, while the 2026 draft principles expose continuing disagreement over extraction and *de facto* appropriation. By the time international rules are agreed, long-term occupation of valuable lunar sites may already have created a *fait accompli*.

The US and China are also unlikely to reach a compromise easily. Their lunar programmes form part of a much wider strategic rivalry in space in which China is currently pulling ahead as the only state to have samples from the far side of the Moon. At the same time, the economic prize is potentially enormous and neither side has much incentive to restrict its own activities in space. Moreover, relations are already icy, as reflected in the Wolf Amendment, which restricts NASA's ability to engage in bilateral cooperation with China without congressional approval.

A world with no framework would be likely to see property rights revert to John Locke's formulation: "As much Land as a Man Tills...so much

is his Property". Published in 1690, Locke used such logic to justify the colonisation of indigenous lands in America.

This time, colonisation of the New World is replaced by land-grabbing the worlds beyond Earth. China's militarisation of contested features in the South China Sea, notwithstanding the 2016 UNCLOS award rejecting its wider resource claims, demonstrates how possession may become entrenched even without recognised legal title. The one upside to this legal wild west is that the stakes are high, foretelling increasing investment by both the US and China into space to ensure that they are the first to set up their manufacturing bases and ultimately write the rules.



Flashing amber

Opportunity

AI and Quantum in Drug Discovery



Simon Baker
*Biopharmaceuticals
Research*

Drug development is difficult. Really difficult. If we go back to the 1960s, the failure rate of drugs in clinical development was 89.7%. However, a lot has changed since then. We have uncovered the sequence of the human genome, can readily determine the structure of proteins and peptides and derive detailed insights into the working of individual cells. Yet the failure rate has crept up over sixty years

promoters claimed, but it did result in some incremental benefits.

Similarly, with the human genome, the thinking was that unlocking the code to human life would unlock disease along with it. Did it? No. We found that 30,000 genes, a far smaller number than we expected, were involved in multi-factorial ways in given pathways and diseases. The conquering of disease as a result simply did not materialise, but again it did lead to some improvements.

Now we face the latest technological breakthrough: Artificial Intelligence (AI). Can this crack the problem of drug discovery? Most would argue yes. Some companies have suggested

problem. There are only twenty amino acids and they only connect and self-assemble in a fixed number of ways. Moreover, there is a very large, high-quality database of protein structures on which to train the system. The DeepMind example has arguably been instrumental in driving unrealistic expectations.

AI in drug discovery is not that new. A paper in 2022 listed 24 companies describing themselves as AI-native drug discovery organisations, going back to assets that began life in 2010. When that analysis was done, those 24 companies had disclosed 160 discovery programmes and preclinical assets, with about fifteen in clinical development.

“Some companies have suggested the timelines for drug discovery can be brought down from years to months”

from 89.7% to 90.8%. Despite huge technological advances, we are no closer to reducing the enormous challenges of discovering new drugs.

The low success rate means drug discovery and development has always been amenable (or vulnerable) to new technologies. Thirty years ago, that simply meant higher throughput: if one molecule in ten thousand made it to the market, simply make more molecules and therefore develop more drugs. That led to high-throughput screening and combinatorial chemistry. Did it work? Not as its

the timelines for drug discovery can be brought down from years to months, which would lead to a huge increase in productivity and a massive benefit to society.

An illustration of the potential of machine learning in particular, coupled with AI, has been the work of Google's DeepMind, which has been hugely successful in determining the structure of proteins. The accuracy almost matches conventional techniques and the work led to the 2024 Nobel Prize for Chemistry.

Impressive as this undoubtedly is, it is a relatively defined, finite

By way of context, the combined in-house portfolios of the top twenty pharma companies contain 330 discovery programmes and preclinical assets and 430 in the first phase of clinical development. That means these 24 startups had a discovery programme base half the size of the world's largest pharma companies – an impressive performance. But how has it worked out?

As far as we can tell, five have announced positive phase 1 or 2 data, six have disclosed failures mostly leading to terminations, and 13 seem to have said nothing.

The attraction of AI pioneers to early-stage drug discovery is understandable. It is a data-heavy environment, but also relatively cheap and quick to enter. This is about moving molecules from the idea stage through preclinical development and then handing them off to somebody else to move through into the clinical stage of development and hopefully onto the market. We have accumulated vast amounts of information on the activity of molecules, the structure of proteins, genomics and patient records, all of which suggest if only we could manipulate and interrogate this data more effectively, we could increase the throughput of R&D. But is that really where AI would benefit most?

This goes back to the issue of the hopper and combinatorial chemistry. If we put more in the top we shall get more out at the bottom, a not unreasonable view from the outside. But would it not be better to change the shape of the hopper? If the hopper narrowed less, productivity would improve materially. Are there other points in the drug development story where AI could make a bigger impact?

Phase two, the proof-of-concept stage of clinical development, still sees very high attrition rates. Molecules which appear to have strong preclinical activity and acceptable tolerability and toxicity from their first phase in human studies then fail in phase two, by which time considerable time and resources have been spent.

A paper published this summer suggested a 20% improvement in speed, success or cost at any point in the development journey shows the biggest impact when the phase two failure rate is reduced by 20%. That is not an easy or obvious starting point for an AI company entering the space, and as a result few if any are pursuing it currently. Perhaps, therefore, a better question to pose in respect of the use of AI within drug discovery is not “What can it do?” but rather “What should it do?”

Already, there are impressive real-world examples of how this has led to benefits in drug discovery, but they are a far cry from prompting Claude to suggest the structure of a new antibiotic.

For example, some years ago AstraZeneca described a process by which it would use a combination of robotic chemical synthesis and AI in a loop to create, essentially, a robotic medicinal chemist. That would synthesise a molecule, screen it, look at the test data and use AI to generate the next wave of molecules and repeat the process. This has the potential to cut costs (or more realistically, cut unit costs and increase activity) and reduce development time, none of which is to be dismissed, but the benefits are modest compared to improvements in clinical development.

Training is the central challenge for AI drug discovery. If you give a

large language model a selection of photographs of animals, it will rapidly distinguish between a Siamese cat and a Yorkshire Terrier with high reliability. What about discovering a drug to treat Alzheimer’s? Where is the data? Where would we start, given that the central problem is we do not definitively know what causes the disease? This is a problem across much of disease biology. For all the accumulation of data we can access, for all the advances we have made in comprehending the biology of disease, there are still enormous gaps in our understanding. Even worse, some of our previous assumptions have been subsequently invalidated as the state of the art has advanced.

In thinking about where AI should be deployed, we hit a problem. If an external startup is seeking venture capital funding for an exciting new AI idea, it is much more likely to be funded if those proposing a particular

Building the hub of the lab



“Incremental benefits have the greatest potential to make a meaningful contribution to R&D productivity”

solution claim drug discovery can be reduced from years to months rather than an idea which will make things a little bit quicker, a little bit cheaper and a little more certain. That is not the sort of pitch that sets the pulse racing for most venture capital investors. But it is precisely those incremental benefits that have the greatest potential to make a meaningful contribution to R&D productivity and drug discovery. This creates a gulf in vision for AI between those inside and outside the industry.

One such mundane but important application of AI is in document production. Drug discovery and development is extremely documentation-intensive – one

good example would be in clinical trial design and management. A protocol for a clinical trial can run to hundreds, if not thousands, of pages and take large teams months to produce. Some companies have reported that can be reduced to minutes, and the saving in terms of time and money is substantial. It may not be very exciting, but it makes a meaningful contribution to R&D productivity.

Another commonly discussed application of AI is in clinical trial enrolment. Often, it takes longer to enrol patients in a clinical trial than the trial itself is designed to last. A two-year study may take four years to recruit. If recruitment can be speeded up, that is a material benefit to the company. Reducing two years to one year is an extra year on the patent life of the drug, which effectively increases its time on market by 7-10%. Using AI to interrogate patient records allows investigators to identify hospital trial sites more effectively than the old method, which is, “I’ve known Dr X at Y for years and she is great”, with access to the specifically required patients almost a secondary consideration.

There are areas of drug discovery which really are a numbers game, lending themselves to the sheer processing power afforded by quantum computing.

An evolving field over the past few decades has been computer-aided drug discovery, where molecules dock in computer memory against a given target. That is much quicker than making molecules and testing them in a living system. However, it suffers from two problems: the accuracy of the model and the vast number of permutations that need to be tried. Screening ten billion molecules seems a very impressive number and is certainly well beyond

the capability of any testing process in real life. However, that is a tiny fraction of the supposed 10^{63} drug-like molecules in existence. For context, there are estimated to be 10^{80} atoms in the observable universe. If we could screen all of them in a model, that could be transformative. Plenty would still fail down the line, but we would at least begin from a better starting point.

With current quantum computing, that figure of 10^{10} has moved up to 10^{30} , achieved in thirty minutes. That’s still a long way from 10^{63} , but as quantum computing power improves, we would expect the gap to close. This has considerable advantages because computer-aided drug design’s exploration of chemical space is so tiny it effectively optimises around a microscopic part of the overall picture – like evaluating a football pitch based on the evidence of one blade of grass.


Even if we do reach the situation where we can test every drug molecule permutation against a given target, the problem is not solved, because success comes down to the applicability of the model itself. It would certainly improve, as a starting point, the current approach whereby only an infinitesimal segment of the opportunity set has truly been sampled.

Unlike many industries, AI is purely a question of opportunity (albeit one not equally distributed) rather than a threat in drug discovery and development, but its benefit probably lies outside the areas many think are obvious. Quantum computing has the potential to boost the effectiveness of existing approaches and fill the hopper, but AI’s role should be aimed further downstream. Measuring the effects will not be easy, but over time, the process becomes a little quicker, a little cheaper and a little easier.



All That Glitters

Gold's Fluctuating Monetary History



As hard to get into as...



Michael Taylor
Economics

Gold has been used as a form of money and store of wealth for millennia. The first gold coins were minted around 600BC in modern-day Turkey. Here we focus merely on the past 150 years or so, examining gold's evolving role in the global financial system. Four discrete periods, interrupted by two world wars, warrant study: the nineteenth-century gold standard, the inter-war period, the post-war Bretton Woods system and today.

The classical gold standard was the international monetary system that prevailed from the 1870s until

the outbreak of the First World War. Each currency was defined as a fixed weight of gold and freely convertible into it on demand. Hence, exchange rates were effectively fixed, for example the sterling to US dollar rate was \$4.86. Exchange rates could vary only within the narrow 'gold points' set by the cost of shipping bullion.

In theory, the system was self-correcting: a deficit country shipped gold abroad, shrinking its money supply and pushing prices down until its exports regained competitiveness and gold returned; a surplus worked in reverse. In practice, central banks managed the process through interest rates – raising them to attract gold, lowering them to do the opposite.

The discipline imposed by the gold standard was designed to achieve long-run price stability. By 1913, both UK and US wholesale prices were

close to the levels prevailing in 1870. But this masked two distinct periods: one of price deflation, followed by one of rising prices. From 1870 to 1896, the gold-based money stock grew more slowly than the supply of goods, forcing falling prices. But this was a price depression, not a slump. Real income per head still grew. After the mid-1890s, new gold supplies from South Africa and Western Australia allowed prices to reflate back towards 1870 levels.

Much economic pain was caused by post-war efforts to return to the pre-war gold standard. The UK returned in 1925 at the pre-war rate of \$4.86 under Chancellor Winston Churchill. But relative inflation rates in the intervening period meant the UK went back at a substantially overvalued exchange rate. The move was heavily criticised,

including by John Maynard Keynes in his lengthy polemical essay, 'The Economic Consequences of Mr Churchill'. Keynes was proved right as the deflationary policies needed to sustain the exchange rate became intolerable with declining exports, falling wages and rising unemployment. Britain left the gold standard in September 1931.

Two years later, the US effectively left too as it battled the Great Depression. Export of gold was banned, and it was made illegal for US citizens to own gold other than jewellery. By boosting the government's holdings of gold, these measures allowed an increase in the money supply, which had to be backed 40% by gold. Early in 1934, gold was revalued to \$35 an ounce from \$20.67, allowing a further boost to the money stock. (The ban on gold ownership lasted for forty years until President Ford ended it in late 1974.) A few countries, including France and Belgium, maintained a fixed exchange rate against gold for a while after 1933, but by 1936 the gold standard was no more.

During the Second World War, several countries went to extraordinary lengths to keep their gold holdings out of the hands of the Nazis. In Operation Fish in 1940, the UK implemented the largest movement of physical wealth in history, by shipping over one thousand tonnes of gold to Fort Knox in the US. In Norway, as the Nazi invasion spread, desperate attempts were made to remove gold to safety; eventually around fifty tonnes were shipped first to the UK and then on to the US. Polish gold had an even narrower escape. As Germany invaded in September 1939, the entire stock of Polish gold reserves, about eighty tonnes, was moved first to Romania and then to Turkey, France and West Africa. Finally, it was split between London and New York.

After the war, not only were most gold reserves held in the US, but they were owned by the US. Allies such as Britain and France had paid

for war supplies by drawing on their gold reserves. The Bretton Woods conference in 1944 reflected both this and the dominance of the US dollar in the post-war world.

The Bretton Woods system involved fixing the US dollar at \$35 an ounce, with dollars freely exchangeable for gold at this rate. Other countries then fixed their exchange rate against the US dollar. Gold was once again the anchor to the global financial system. But currencies could be devalued or – less often – revalued against the US dollar (and gold) if large imbalances such as in the balance of payments developed. The UK initially joined at \$4.03 but as early as September 1949 was forced into a 30% devaluation, to \$2.80. In November 1967, a further 14% devaluation took sterling to \$2.40. Gold as a global monetary anchor, and the Bretton Woods system, lasted until August 1971, when President Nixon announced the end of dollar convertibility. Increased spending by the US government, including on the war in Vietnam, resulted in large overseas holdings of US dollars. Foreign central banks began converting dollars into gold on a scale such that the Federal Reserve was in danger of exhausting its stock. The link between the US dollar and gold had become unsustainable.

Freed from its \$35 fix, gold came into its own as a store of wealth during the 1970s. This was a period of rising inflation, two oil price shocks, and at the end of the decade the Soviet invasion of Afghanistan, raising Cold War tensions. In January 1980, gold reached the then all-time high of \$850 an ounce.

From there, gold began a prolonged period of declining prices, as newly elected governments in the US and UK tackled inflation with high interest rates. The fall of the Berlin Wall in 1989 heralded the end of the Cold War, reducing geopolitical risks. Gold remained under pressure throughout the 1990s, and in 1999 hit its lowest price in over twenty years at \$250 an ounce. At this point, UK Chancellor Gordon Brown began

selling 400 tonnes (60% of total holdings) raising \$3.5bn at an average price of \$275 an ounce. Six years later, the gold price was \$972, more than three times higher. Brown was not the only one to do so. Around this time, authorities in Australia, Switzerland and Canada also sold large quantities of gold reserves.

For most of the 2000s, gold was on a rising trend against a backdrop of 9/11, wars in Afghanistan and Iraq and a weakening US dollar. The Global Financial Crisis (GFC), with the global banking system appearing in danger of collapse, pushed gold through \$1000 for the first time in 2008. But a year later, forced selling as the GFC panic intensified brought the price back to \$700. Once economic recovery was established and the banking system stabilised, gold regained momentum, reaching \$1900 an ounce by 2011. For the next ten years or so it fluctuated between \$1300 and \$1900 an ounce.

From the end of 2022 onwards, gold began a meteoric rise, culminating in an all-time high of \$5589 an ounce in January this year. Various factors contributed. Geopolitical risk rocketed with wars in Ukraine and the Middle East; there was renewed central bank buying; and, perhaps most importantly, inflation soared to multi-decade highs following the 2022 energy price shock. Once again, gold performed its age-old function of acting as a store of wealth.

It is highly improbable that gold will ever again form the anchor to the global financial system. But global government debt is at record levels, concerns grow about central banks' commitment to hitting their inflation targets and geopolitical uncertainty remains elevated. Gold's continued role as a store of wealth and as a 'safe haven' asset looks assured. A recent innovation has been tokenised gold, creating digital tokens that represent ownership rights in physical gold. This should boost liquidity and enable easier and cheaper trading in smaller quantities. Gold's future still looks bright.



Fanfiction

The Rise of the Music Superfan

Easier to find in a crowd



Aleia de Buretel
Software & Cloud
Research

The lights drop and the bass finds my ribs before my eyes adjust to the darkness. Thousands of phones rise like a second skyline as *Que Dieu vous bénisse Paris* flashes across the screen, the crowd pressing towards the barricade. A few days earlier, it read God bless you, London. That night, I'd lost my voice before the second song. Glastonbury, London, Paris. Three Lana Del Rey concerts, far too much of my savings. By the third show, I was wondering: am I a superfan?

Before superfans, there were groupies, Beatlemaniaacs, Deadheads and Stans. They slept outside venues, traded bootlegs, studied B-sides, bought T-shirts.

For the 1970s superfan, there was only so far it could go. For today's obsessives, there's a new realm of possibility. They find each other, at all hours, on TikTok, Tumblr or X, multiplying their 'super' powers.

Fans can use their power to resurrect an old hit. Today's teenagers have put their head on Paul Anka's shoulder and skateboarded to Fleetwood Mac's *Dreams*. Savvy artists have

embraced the opportunity. When @heiratet posted 'I am depressed' on a rainbow seascape of dolphins to Zara Larsson's seven year old song *Symphony*, a viral movement followed. Larsson integrated the meme's aesthetic into her styling, stage set and the rollout of her album *Midnight Sun*. One post became a No. 1 on Billboard's Viral Top 50 Chart.

Fans can unearth talent. TikToker Haley Sharpe choreographed a dance to Doja Cat's *Say So* in her bedroom, posting it move by move – and the world copied. Dua Lipa and TikTok star Charli D'Amelio joined in, turning *Say So* into a global hit. Sharpe earned a cameo in the music video, deservedly so. The promotion cost the label nothing.

Beyond TikTok, on Tumblr, superfans took the unsigned singer-songwriter Ethel Cain from bedroom demos to Billboard artist, one reblog at a time. On Twitter, they made Charli xcx's album so inescapable that even a US presidential campaign rebranded itself as 'brat'. On Discord, fans coordinate streaming to make each play count towards the charts. The rules? Stream for thirty seconds minimum, no track twice in a row and never on mute.

At first, fan involvement was a source of off-beat stories. In 1975, two teenagers declared themselves "President" and "Field Marshal" of

a KISS Army and threatened their local radio station into playing songs by KISS. Their grassroots campaign helped to break the band.

Despite these successes, the recorded music industry has a love-hate relationship with superfans. Why? In a word: control.

In the twentieth century, labels held the keys to success: radio playlists, MTV rotation, the cover of *Rolling Stone*, a nod from the Grammys, all required industry backing. But gradually, internet communities picked the locks. Today, anyone can kickstart a career if audiences connect with them. Discovery is democratised. The fans know how to wield their power.

Streaming drives c70% of global recorded-music revenue. But growth is stalling in Western markets and revenue per user is under pressure. US subscriber gains fell from c29% in 2019 to c7% in 2025. Listeners now spend only c4% of their time on their favourite artist. Here too, attention is fragmenting.

As a result, digital service providers (DSPs) are shifting from tolerance of fan involvement towards pure monetisation. Spotify users' subscriptions have crept up again this year. This is because price rises only travel so far before casual subscribers eye the exit, and there are few new listeners to recruit. Hence, DSPs have set their sights on the superfan,

“As the internet splinters, the superfan splinters with it, into millions of tiny actors, impossible to discern”

or ‘super-premium’ tier. It is a last resort, dressed up with enough extras to feel worth the spend. The unpaid salesman has been recast as just another paying customer.

But to charge the superfan you must identify them, and the industry cannot. The current flat-fee pricing model does not distinguish between the fan playing one song on a loop and the fan who plays a thousand. Luminate, whose numbers drive the charts, counts a superfan as anyone engaging with an artist in five of thirteen ways (streaming consumption, advocacy, community behaviour...), no spending required, then sells that superfan as a listener who spends 80% more on music each month than average. Devotion is hard to measure, so money counts, and the superfan is redefined as a superspender.

Spending is the wrong signal to follow. On Luminate’s own Venn diagram, fewer than four in ten superfans fit both the ‘purchase’ and the ‘engagement’ criteria. Commitment does not always equal deeper pockets. A generation raised on abundance considers music to be free. Thus, James Blake’s paid fan tier stopped charging members within two years, while on Patreon, free memberships now outnumber paid four to one.

Until now, monetisation has failed because labels and DSPs, under pressure, are unsure how to harness fan engagement. Their solution? Fake it.

“Everything on the internet is fake” anyway, at least according to Andrew Spelman, co-founder of Chaotic Good (CG). CG is an agency that artificially creates hype, building networks of ‘fan pages’ and threading its clients’ songs through them. Spelman calls this “trend simulation”. You want to believe your favourite artist earned

their place, but can you be sure? Indie-rock band Geese seemed to be having an organic breakout in 2025, until it emerged Chaotic Good had helped to engineer the viral attention.

None of this is new. As SoundCloud marketer Kalesha Madlani, who turned fangirlism into an 800,000-follower account for One Direction, explains, clipping – scattering song snippets across fan accounts to seed a trend – “is such an old strategy”.

Major labels have always staged enthusiasm; the change is where it happens. Girls were allegedly hired to scream for Frank Sinatra and placed in the audience. Today, the trick is to make these orchestrated swells of excitement so well disseminated as to be practically impossible to detect as artificial.

Our only clues come when someone slips up. ‘Superfan’ dailysombr gave herself away with improbable backstage access to pop singer Sombr, while R&B star Normani’s “biggest fan” accidentally revealed herself to be a member of her team.

For the label, it is a risk worth taking. Like a Trojan Horse, a sleeper superfan is planted and does its work before anyone notices it is there. Listeners may eventually realise and turn on the artist, but it is probably too late. Real fans will have arrived. The majors may now be paying to manufacture what fans once gave them for free, but they are farming momentary attention, not true fandom.

So why are labels and DSPs in the West still trying to capture the superfan? Because they can see Tencent Music’s Super VIP (SVIP) tier in China. It works somewhere.

SVIP subscribers passed twenty million by December 2025, paying five times more for perks including exclusive music, priority tickets

and collectible artist cards. Tencent credits the tier with helping to drive its 7% rise in revenue per user, even as free listening shrinks. But the model doesn’t travel. Tencent owns the artists and the platforms, allowing it to set the price. The US cannot replicate this. Kanye’s *The Life of Pablo* was released as a Tidal exclusive in 2016, with West vowing it would never appear elsewhere. Until fans pirated it 500,000 times in a day. Within weeks, it was everywhere.

That is the key. Online, superfans are everywhere. New platforms have allowed their role to transform. The superfan has always sought some kind of symbiotic relationship with an artist. Over time, however, they have given themselves a backstage pass. This no longer means shimmying through a stage door, past security, for banter and flattery, but indirectly working for the artist, becoming amateur marketers, tastemakers, even gatekeepers. As the internet splinters, the superfan splinters with it, into millions of tiny actors, impossible to discern.

Although the industry attempts to identify the superfan in the data, they are easier to find in a crowd. Live music has what recorded music lacks: scarcity. There are only so many seats at Wembley, and Taylor Swift fans will pay anything to be in them. While recorded music has stalled, Live Nation’s revenues grew nearly 10% last year, with premium and VIP spending growing even faster. It’s no mystery. I have listened to Lana Del Rey through headphones hundreds of times, but nothing matches the feeling in my ribs hearing her live.

As Jerry Garcia put it, when asked where the Deadheads at his Grateful Dead shows kept coming from: “Maybe we’re one of the last adventures in America.”

Tidal Wave

The Risks of AI



Omar Sheikh
*Software & Cloud
Research*

I am an above-average professional AI user, maybe \$200 of subscription spend plus token overage per month. In my research process, most of my work involves using frontier models – everything from testing hypotheses and tracking product pricing trends, to collecting data and building first-pass models.

As part of a research project over the summer, I used Anthropic's Claude Code and OpenAI's Codex to build systems which connect to the latest AI and video generation models and use them to complete complex multi-step content creation processes, then scour thousands of pages of developer documentation and user feedback to judge how this compares to what can be achieved in a traditional software tool. What would have taken hundreds of hours without AI models was completed in a fraction of the time by typing (and sometimes speaking) natural language instructions into a text box.

The capability of these models, and the agentic harnesses around them, is awesome and exciting. Overall, my productivity is probably up fivefold this year. But my experience of spending a considerable amount of time using these tools has made me think more carefully about the broader societal impact of AI use. As the infrastructure and models become more capable, and especially as the frontier labs train models using recursive self-improvement


techniques, the path to general intelligence and potentially super intelligence is starting to become clearer. I think we all have to be prepared for what might come next.

First, my experience. The mental fatigue I felt from using these models was off the charts. A big project starts as a series of simple prompts, but Claude Code and Codex spin up sub-agents to go down research rabbit-holes in pursuit of the set goals. Very quickly, I had thirty agents beaver away to dig out, read, digest and analyse documentation and changelogs, to learn how to use software tools from scratch and measure performance and write code.

But the work isn't truly autonomous, i.e. completed end-to-end automatically with no human

intervention. Teams of sub-agents create a substantial volume of work that must be reviewed, and it is often not quite what you wanted. In addition, the starting scope tends to expand, quickly, as the sub-agents uncover new wrinkles in the process you hadn't considered. This means the initial simple prompts become a sprawling, complex branch system. You have to process and approve all this content at each step. It is exhausting.

As an aside on mental fatigue, it is worth highlighting that, as many Claude users have been saying for months, Anthropic's model performance appears to have degraded. Opus 5 in Claude Code has become intolerable for complex multi-step tasks – answers have got



longer, language harder to decipher and workflow gaps have started to appear. When the model starts to misfire, the pressure on your own brain function skyrockets.

Second, you start to think the system has 'god-like' capabilities, and this becomes a problem. Models have so much information at their fingertips, they respond immediately and are so pleasant to talk to. You get to know them, trust them, respect them. They make you smarter in many ways, because you can push the boundaries of your knowledge. But it's also easy to slip into letting the models think for you. Especially as I got more tired, I found myself not reading responses in full, only skimming the huge markdown files with feedback on each turn of the process and therefore slowly ceding some of the thinking to the model.

As these models become more capable and widely adopted by the general population, my experience suggests we should prepare for four big shifts.

First, many jobs will change radically because the cost of finding, processing and presenting information has collapsed. This

means a large proportion of what knowledge workers do today can be devolved to machines, and our roles will have to up-level: we will have to learn how to design and manage systems that do much of the work we do today. Those who learn this now will become hugely productive and valuable. By contrast, workers who do not up-level, and remain focused on work that can be automated – building a model, summarising basic information – will look noticeably less productive and will ultimately become less valuable.

Second, AI model usage is taxing on brain function in ways we don't fully understand, so unfettered access by children is probably not a good idea. AI models provide children with tools that can process information for them on demand, without having to use their own brains to read and understand the material. If I slowly handed over thinking during my project – when I was consciously trying not to – it's likely many children may take the short cut. I would not be optimistic about how this will impact the development of their reasoning ability, critical thinking and memory retention. As with social media, I worry we will only prove harmful effects of a new technology many years down the line, and hope that policymakers are alive to the risks.

Next, I worry about the impact of models' confirmation bias on how people think about the world around us. Frontier labs cannot fully explain why but, on default settings, models generally want to be helpful and validate your point of view. The problem with becoming reliant on systems that rarely disagree with you or challenge you is you can become more entrenched in your beliefs and world view. Like an extreme version of social media echo chambers, the risk is that we become more polarised and less likely to agree with those who hold different opinions. After all, if 'god in your pocket' agrees with you, everyone else must be wrong. Right?

If this is true, it is also potentially a problem if most people use models from two or three closed-source frontier model providers, all privately owned. How can society be sure the model training process hasn't been tilted to favour the priorities of one segment of the population at the expense of another? And is there any way other than direct government regulation to verify this isn't the case?

There is active investor debate about whether closed-source models will be used more than open-source and open-weight models. We should probably also have a debate about what is best for society as a whole.

Finally, one of the silent side-effects of AI usage in the general population may be the loss of ownership and self-worth many will feel. It sounds enticing to be freed from drudge work, but what if you love doing all that you do? What if doing those tasks fills you with pride and leaves you fulfilled? Maybe you don't get as much personal satisfaction from orchestrating a team of AI agents as from updating a spreadsheet.

If that's the case, the question becomes what these people will do with their time. One of the first areas of discussion appears to be that universal basic income might be a solution, but this would do nothing to solve the self-worth problem. Instead, it may be that society will have to construct more profound solutions to fill people's time productively.

Of course, every ground-breaking technology – printing press, steam engine, electricity, aircraft, the personal computer, the internet – has trailed societal risks. They have also been the drivers of enormous gains in human prosperity. There is no question AI models will bring both too.

Our sense is that most of the public discourse is currently focused on the potentially huge opportunity for human prosperity. Much thought will also need to go into how society should prepare for the tidal wave of changes that are coming.



The Global Retail Bank

From Empire to Platform



Harry Bartlett
*FinTech & Banks
Research*

There was a moment, around 2007, when the global bank believed it could be all things to all people, everywhere. Citigroup was its avatar: 375,000 staff and roughly \$2.2trn of assets, running a consumer bank on almost every continent. The logic felt compelling: more markets, more products and more balance sheet meant more profit.

It was, we now know, a category error. The defining lesson of the past two decades is that in banking scale and strategy are not the same thing. The institutions that prospered were not necessarily the largest, but those most disciplined about where they chose to compete.

The old model had a certain elegance. Deregulation had torn down the walls between commercial and investment banking; cheap wholesale funding made leverage look almost free; and globalisation promised that a bank that could sell a mortgage in Manchester could surely sell one in Mumbai. Banks planted flags, exporting entire retail and trading franchises across the world. The hubris peaked in 2007 with the record €71bn dismemberment of ABN AMRO by a consortium led by RBS – a deal with disastrous consequences for all three of its acquirers.

The crisis did not merely bruise this model; it also discredited it. Capital became expensive, leverage costly and growth for its own sake increasingly difficult to justify. Citigroup hived off unwanted assets into a ‘bad bank’

and by the end of 2009 had shrunk its balance sheet by about half a trillion dollars from its 2007 peak. The age of the flag-planter was over.

There followed an era of subtraction. The survivors worked out where they held a durable advantage and jettisoned everything else. Mass-market retail banking abroad, once the crown of empire, proved to be its weakest component: capital-hungry, often sub-scale and easily out-competed by local rivals. Under Jane Fraser, Citi has exited fourteen overseas consumer markets, while HSBC has sold numerous retail operations outside its core markets in the UK and Asia.

The global bank did not disappear. It changed shape. The prizes were the capital-light franchises a local rival could not easily replicate: payments, foreign exchange, transaction banking, cash management, trade finance and corporate and investment banking. Citi’s Services business, for example, now spans more than ninety markets and sits at the centre of its international strategy. The common theme is a shift from balance sheet to network, from owning risk to charging for expertise and reach.

The distinction between wholesale and retail banking is important. Corporate clients often have an inherently international financial life; consumers generally do not. A multinational may genuinely need a bank in twenty countries. A household typically needs a current account that works within one regulatory and payments system. Without a meaningful source of differentiation, an international retail bank could spend heavily to recreate what a domestic bank already did well. Yet technology is beginning to change the economics

that made international retail banking so difficult.

The first generation of digital banks changed the economics of international retail banking. Instead of rebuilding branches, staff and technology in every market, digital banks could reuse a common platform and customer proposition, adding only the licences, payment infrastructure and local capabilities that regulation and market conditions required. Technology has not made banking borderless, but it has reduced the amount of a bank that needs to be local – and therefore the cost of crossing a border.

Two players stand out. Revolut now has more than seventy million customers, operates in more than forty countries and has ambitions to become the “first truly global bank”. Nubank serves more than one hundred and thirty million customers across Latin America. Their paths are revealing because they started from very different propositions. Revolut’s initial wedge was international by nature: foreign exchange, multi-currency accounts and international transfers. Nubank began with a distinctly domestic problem – access to more affordable credit in Brazil – before taking its technology and operating model into Mexico and Colombia.

Both used their initial foothold to broaden the customer relationship into a wider range of financial services. But neither is recreating the old universal bank. The difference is the old global bank tried to transplant a complete institution into another country, whereas a digital bank exports those parts customers notice – the products, app and service – while rebuilding only what regulation and local market conditions require.

The old global bank exported branches, people and balance sheets; the new one exports a platform.

What is striking now is how incumbents are adopting the same model. The smart money is re-scaling differently, with a narrower target, a digital proposition and a technology platform that can travel.

JPMorgan is the clearest example. Launched in the UK in 2021, Chase has reached around three million customers and £30bn of deposits. It is now moving into Germany, initially with a savings-led proposition, with further European markets in its sights.

Santander is pursuing a similar strategy through Openbank. The platform has expanded from Spain into Germany, Portugal, the Netherlands, the US and Mexico. In the US, it passed \$2bn of deposits within its first three months, and now exceeds \$7bn. Santander describes Openbank as its global digital consumer bank, with further expansion planned.

The attraction is obvious. Unlike a standalone neobank, an incumbent does not have to build the entire financial machine from scratch. It already has a balance sheet, funding, licences, risk infrastructure, brand recognition and, often, millions of customers. The digital bank becomes a new distribution layer over an existing financial institution.

The technology architecture is the key. These propositions are not simply mobile interfaces bolted onto legacy systems. Chase UK leverages Thought Machine's cloud-native core banking platform, Vault, while Santander has developed its own platform, Gravity. The result is an intriguing reversal: the technology which once allowed fintechs to attack banks is now allowing banks to reproduce some of the advantages of fintechs.

For domestic neobanks, this raises the stakes. Their original advantage – being faster, cheaper and easier to use than the incumbent – is no longer sufficient if a global bank

can reproduce the same experience while offering the safety of a household name and the funding capacity of a trillion-dollar balance sheet. As domestic markets mature and customer acquisition becomes more expensive, the ability to take the proposition abroad becomes increasingly important.

For domestic incumbents, the threat is different. Neobanks have demonstrated that customers can be persuaded to switch, but customer numbers alone do not determine banking economics. Their average balances and revenues per customer have often remained below those of established banks. Armed with trust, funding and a lower cost-to-serve, incumbent digital banks may be better positioned to compete for higher-value, more established customers. Chase UK's roughly £10,000 average deposit balance, for example, is broadly in line with UK incumbents and materially above that of some domestic neobanks.


The likeliest future is convergence. Neobanks are becoming more like banks as they add deposits,

lending and licences; incumbents are becoming more like fintechs as they adopt cloud-native technology, digital distribution and lower-cost operating models.

The result is intensifying global banking competition. Domestic neobanks are increasingly looking overseas for their next leg of growth, while domestic incumbents face pressure from both international fintechs and the digital arms of foreign banks. Geography, once one of the strongest protections in retail banking, is a shallower moat.

The global bank is therefore making a comeback. But it is not returning as an empire. The old global bank planted flags; the new one deploys platforms.

The lesson of the past twenty years was not that banking could not scale across borders. It was that branches and balance sheets were the wrong things to export. The next generation of global banks may finally have found the right answer: not a bank in every country, but one banking proposition that can travel.



The market

Screen Time

Detecting Cancer



Jamie Clark
MedTech
Research

Molecular diagnostics (genetic sequencing) are already well established in cancer treatment. New approved cancer therapies are increasingly precision therapies, which work on patients whose tumours carry a specific genetic mutation. As the number of genes for which precision therapies are available for a given type of cancer has increased, molecular diagnostics have expanded from single-gene polymerase chain reaction to tests that sequence hundreds of genes simultaneously. In addition to identifying specific genetic mutations, they can detect other signals that help in selecting the right treatment.

Today, new molecular diagnostics products have the potential to change how cancer is detected and treated. Within three years, it is possible that multiple cancers will be detected by a single blood test, and the current practice of giving all patients chemotherapy after surgery will become personalised.

The most important factor in determining cancer survival is when it is detected. If caught early, surgery can remove it. Once it has started to metastasise, it is much harder to treat, and the chances of survival are lower. For that reason, screening programmes for various cancers try to improve early detection, and so survival.

Breast, lung, colorectal and prostate cancers account for around half of cancer diagnoses annually. There are

established screening programmes for colorectal and breast cancer in most countries. Breast cancer screening compliance in the US is 78% and the data shows 92% of these patients are alive after five years. Colorectal cancer screening compliance is lower, at 69%, and so is the five-year survival, at 65%. Lung cancer survival is less than half this.

Lung cancer screening is typically only available to smokers, which is not much help if you are an unlucky non-smoker. Nearly half of lung cancer cases are only diagnosed after it has spread, and five-year survival is only 30%. For cancers outside the big four, there are no screening programmes currently available (beyond cervical); this is visible in five-year survival rates – 40% for stomach cancer, 22% for ovarian cancer, 14% for pancreatic cancer.

Colorectal cancer screening typically involves a faecal immunochemical test (FIT), which detects blood in a stool sample, an indication of cancer. Its ability to detect cancer is 90%. Since 2013, a product called Cologuard has been available in the US, which combines the FIT test with DNA sequencing to identify cancer signatures. Its ability to detect cancer is the same as the FIT test, but its ability to detect early-stage abnormalities that will develop into cancer is twice as good.

In 2024, the first blood-based colorectal cancer test was approved, Guardant Health's Shield, which examines DNA methylation patterns to determine cancer signatures. In 2026, the US Food and Drug Administration (FDA) approved the second, Freenome's SimpleScreen. More are likely.

First blood



These blood tests offer the same sensitivity as a FIT test or Cologuard but are less sensitive to early-stage abnormalities. However, we expect blood-based testing for colorectal cancer to grow rapidly because providing a stool sample puts people off; if a patient has an annual physical, they will probably already be giving a blood sample; and because a physician would prefer to know the patient has completed the screening rather than relying on them doing the FIT test.

Shield is designed for multi-cancer early detection (MCED). If you take a Shield test and agree to give Guardant Health access to your medical records, it will include additional screening results on nine further cancer types. Most patients are opting in.

The results beyond colorectal cancer have yet to be FDA approved, but a large study is running and could provide substantive data in late 2026 or early 2027. Grail filed its Galleri MCED test with the FDA in January 2026. It aims to identify fifty different types of cancer. It has not provided a detailed breakdown by type but disclosed the sensitivity towards twelve cancers accounting for two-thirds of cancer deaths is 70%. This could be sufficient for approval given the majority of these have no screening programme.

Several companies offer MCED tests and are working towards FDA approval. All the tests are only currently available on a cash-paying basis but a law passed in 2026 allows for Medicare reimbursement of FDA-approved MCED tests from 2028.

Prostate cancer screening is controversial, with the UK National Screening Committee rejecting a population-level screening programme in early 2026. In the US, commercial insurance coverage of prostate cancer screening is not mandated. The rationale behind these decisions is that the only currently available screening test, measuring levels of prostate specific antigen (PSA) in the blood, is simply

not very good. Raised levels are taken to be indicative of cancer, with an MRI-guided prostate biopsy required for firm diagnosis.

However, not all cancers raise PSA levels, and levels may also rise for reasons other than cancer. In its rejection, the UK cited 15% of men with prostate cancer having a normal PSA level, 72% of men with raised PSA levels not having cancer and 80% of those who were subsequently diagnosed with prostate cancer not requiring treatment. Spending a lot of money to scare people unnecessarily would make a poor screening programme.

Many MCED tests seem to have low sensitivity to prostate cancer, but the Stockholm3 test, available privately in some Western and Northern European markets, including the UK, showed 90% sensitivity in a large prostate cancer screening trial. It uses an algorithm that integrates the level of certain proteins in the blood (including PSA), genetic mutations associated with prostate cancer and a patient's age and medical history. The UK plans to include its data in modelling that will inform the next prostate cancer screening recommendation.

For patients who have been treated for cancer, another new use of molecular diagnostics is in Minimal Residual Disease (MRD) testing. This sequences any DNA found within a patient's blood sample to identify signals associated with cancer, known as circulating tumour DNA (ctDNA). If no ctDNA is found, it suggests there is no cancer.

This is subject to the caveat that some cancer types are less likely to shed ctDNA than others, but breast, colorectal and lung are all high shedders. The availability of MRD testing offers a number of ways in which cancer care can be dramatically improved, all dependent on the generation of substantive clinical data demonstrating it can improve patient outcomes.

Firstly, it is standard practice after a tumour has been removed

via surgery for the patient to undergo a course of chemotherapy to ensure any missed cancer cells are eliminated. There are costs to chemo, both economically and for quality of life.

Data from a Japanese study showed that patients after surgery who had no ctDNA in their blood (ctDNA negative) had the same survival chances without chemo as ctDNA negative patients who underwent chemo. Japan became the first country in the world where medical guidelines now recommend colorectal cancer patients undergo MRD testing after surgery to determine if they need chemo or not. Data needs to be generated in a Western population before we can enjoy the same.

Secondly, MRD can be used to monitor the effects of a patient's treatment plan, potentially allowing cessation if they test ctDNA negative, or switching therapy if they continue to remain ctDNA positive.

Finally, cancer survivors typically undergo follow-up imaging to determine whether any new tumours are visible that suggest cancer had returned. MRD testing can detect the return up to nearly a year earlier than imaging, potentially allowing treatment to begin substantially earlier.

Clinical trials are needed to answer the question of what to do with this new information. Treatment should probably begin as soon as possible, but whether it should be treated as it would if a tumour was detected through imaging, if dosing should be different or whether an entirely different treatment should be used, is as yet unknown.

Blood-based screening for colorectal cancer is already here, and successfully generating data in other cancer types could see MCED become standard. MRD offers great prospects for changing cancer care if data is generated showing improved outcomes. It seems only a matter of time.

One Signature or Five?

The Reality of Enterprise AI Adoption



Rowan Adley
*TMT Specialist
Sales*

OpenAI launched 'Instant Checkout' in September 2025, enabling shoppers to buy from Etsy and Shopify merchants without leaving ChatGPT. By April 2026, it had been wound down following limited uptake.

Of Shopify's millions of merchants, "only about a dozen" had gone live, according to Shopify's president, compared with the million-plus OpenAI had forecast. OpenAI was candid: "The initial version of Instant Checkout did not offer the level of flexibility that we aspire to provide, so we're allowing merchants to use their own checkout experiences while we focus our efforts on product discovery".

That flexibility meant multi-item carts, promotion codes, loyalty schemes, shipping guarantees and the remittance of VAT and sales tax. This was a failure of integration, not model intelligence. For a transaction to occur in ChatGPT, OpenAI, Shopify, the merchant, the payment processor and a consumer willing to change their habits all had to agree. The merchant, watching the customer relationship and data migrate into somebody else's walled garden, had every reason to refuse.

The demand was real. Adobe, analysing over a trillion visits

to US retail sites, measured AI-referred traffic up roughly five-fold year on year in the first quarter of 2026. By March, its conversion rate was 42% higher than non-AI traffic, having been 38% lower a year earlier. Those purchases were completed on merchants' websites, using existing payment systems. Shoppers were willing to let a model find the product; buying still relied on familiar plumbing and established trust.

Developer tools tell a different story. Cursor reached \$100m in annualised revenue in January 2025, faster than any software company on record at the time, and around \$4bn by June 2026. Anthropic's Claude Code, generally available from May 2025, reached a \$2.5bn run rate by February 2026, its last disclosed figure.

The models underneath are broadly the same. What differs is the shape of the sale. A developer buying a coding assistant can be buyer, user, beneficiary and judge in one, and, whether paying on a company card or not, is able to assess quickly whether it was worth it.

For all the technological momentum, the most important variable in enterprise AI adoption may be how many people must say "yes". One signature, and revenue can arrive almost immediately; five or more, and you need pilot projects and committees.

Even the developer's advantage has limits. A compiler verifies that code runs; a test suite verifies that it passes the tests somebody

chose. Neither establishes that the code is right. Once one developer becomes a large engineering organisation, human review capacity becomes a constraint.

Still, many coding errors can be spotted and corrected quickly by the person who paid for the tool. A sofa bought by an agent has no equivalent verdict. Whether it was the right sofa, at the right price, delivered to the right address with the right taxes remitted depends on more than an intelligent choice. Contracts, liability rules, fraud handling and tax systems underpin the transaction. Lawyers



and insurers have spent decades building those arrangements. New model releases do not necessarily accelerate them.

The economics complicates matters further. Cursor initially ran negative gross margins and by April 2026 was only just gross-margin positive, making money on large enterprise accounts while losing it on individual developers. Enterprise's share of revenue grew from a little over half in March to three-quarters by June.

The revenue that arrived first, from the developer with a credit card, was less profitable. The revenue paying for itself came through the enterprise route, requiring more people's agreement. Stakeholder count seems to govern the quality of the money as well as its speed: fast revenue is the hare, slower revenue the tortoise.

Widespread use, meanwhile, need not mean widespread productivity gains. A National Bureau of Economic Research paper by researchers from the Bank of England, Deutsche Bundesbank and Atlanta Fed surveyed six thousand senior executives across the US,

Britain, Germany and Australia between November 2025 and January 2026. While 69% of firms reported using AI, 89% reported no notable change to labour productivity over the previous three years.

The behavioural detail deserves as much attention as the headlines. Executives who use these tools do so for only an hour and a half a week. These are perceptions rather than measured output, so they can lag. A developer can adopt in the morning and be adding value by the afternoon. Encouraging a whole organisation to work differently takes longer.

Industrial electrification offers a precedent. Stanford economist Paul David's 'The Dynamo and the Computer' records that electric motors were on sale from the 1890s, yet powered under 5% of American factory machinery in 1899 and took until 1919 to pass half.

The reason was a building. Steam factories ran every machine off one overhead shaft by leather belt, so they were built tall, cramped and in the shaft's order rather than the work's. Early adopters simply bolted a motor to the shaft, saving coal but leaving the factory much as it was. The gains came when each machine had its own motor and the factory was rebuilt around the work. That required architects, electrical engineers, insurers and, usually, a new building. Many signatures, gathered over two decades.

Of course, the models must also be good enough. In February 2024, Klarna deployed AI customer service agents, displacing around 700 jobs, but by May 2025 found quality suffering and began rehiring. A deployment can clear every organisational hurdle and still disappoint.

That example is dated on an AI timeframe, but data from METR suggests capability remains a constraint. Its benchmark measures the length of a task, timed by a human expert, that a model can complete at a given success rate.

Performance falls sharply as tasks lengthen. The tasks derive from software engineering, so extrapolation is risky, but genuinely autonomous agents still face technical limits.

A sharper objection turns the argument on itself. If a second model can reliably check the first model's work, perhaps it can provide the immediate verdict that makes coding tools so useful. The ability to judge results would then depend on capability too. Perhaps. But that does not reduce the signature count. No model, however capable, settles who pays when an agent buys the wrong product.

In the UK, that question has an answer of sorts. Section 75 of the Consumer Credit Act (1974) makes the card issuer jointly liable with the retailer on credit purchases above £100. However, it has yet to be tested against a purchase not initiated by a human. That will be a matter for the courts, working at their own speed.

None of this tells us when, or whether, an agent will routinely spend money on a person's behalf or do a day's work inside a large company. The problem is much the same. A purchase in a chat window needed OpenAI, Shopify, the merchant, the processor and the customer. Serious enterprise deployment needs procurement, security, legal, the budget holder and the team whose job will change. Somewhere in that chain there is usually a turkey which will refuse to vote for Christmas, possibly for reasons unconnected to the technology.

So watch the paperwork and the signoffs. Will merchants adopt the agentic payment protocols being drafted? Will software vendors indemnify customers against their agents' mistakes? Will they price on outcomes rather than seats? Who will insure an agent's mistake, and what will they charge?

Each decision requires a signature. No model release will supply it. Count signatures, not benchmarks.





They might have been right

Regulatory Friction

The Road Ahead for Light Vehicles



Tobias Beith
Automotive
Research

The expanding cohorts of middle-class consumers in countries such as China, and broader economic growth elsewhere, have increased demand for all forms of transportation over the past few decades. However, that growth was not carbon-free, and it is unsustainable. Vision blurs as the writer recalls the destructive effects of human action-derived greenhouse gas (GHG) emissions that are evident today, such as deadly wildfires in Europe and floods in Pakistan. The Paris Agreement perhaps best describes global climate efforts currently. Nearly two hundred countries and the European Union signed this wide-ranging treaty. The US appears undecided

about compliance, withdrawing, readmitting and re-withdrawing. The Paris Agreement's most important themes are: (i) pursue efforts to limit rises in global temperatures versus the pre-industrial era; (ii) achieve net-zero GHG emissions in the second half of this century; (iii) a pledge from developed countries to provide finance, technology and capacity-building support to developing nations for climate projects.

This, and the carbon intensity of road transport, explains the existence of longstanding policies on new and in-service vehicle emissions (GHGs and NOx) in all major countries and trade blocs measured by volumes. Indeed, transport contributes c25% of global CO₂ emissions today; of which, light automobiles (cars, vans and trucks weighing less than 3,500kg) and medium- and heavy-duty vehicles (tractor units, other commercial) generate three-quarters.

The absolute CO₂ emissions of these subcategories are comparable, whereas the size of their productive bases is quite different, c1.5 billion and c100 million units respectively. The average age of each fleet is roughly fifteen years.

Combusting hydrocarbon-based fuels in a compact environment and transmitting the resultant forces through componentry, first the pistons then the crankshaft, is the primary mechanism that has moved vehicles for more than one century – the Otto internal combustion (IC) engine in the abstract. Regarding climate action, the issue is that the IC engine exhausts gases, such as GHGs, alongside mechanical power. Over the typical lifecycle of an 'old world' light vehicle, International Energy Agency (IEA) analysis indicates c90% of its aggregate emissions are generated from usage, chiefly in the tank-to-wheel stage. Legislation on vehicle emissions does not

advocate for a particular compliant technology. That is the automobile value chain's challenge.

IC engine capacity downsizing, hybridisation and lightweighting *inter alia* have decreased fleet average emissions across China, Europe and the US by c3% pa since 2000. A short saga with diesel fuel is intermixed here; we learnt ex-post about its non-GHG harmful byproducts. That rate of productivity is not fast enough for forthcoming targets. Legislation in the EU, for example, asks that fleet average tailpipe emissions fall at c9% pa between 2021 and 2030. We also posit that the industry has reached the point of diminishing returns on R&D in IC engines, given independent research over the past century by multiple organisations. While the case for IC engines using synthetic fuels is sometimes discussed, they are presently too expensive and so serve as a niche bridge.

The all-electric propulsion system (motor, inverter, transmission, batteries) is empirically an effective decarbonisation solution. Differently, the exploitation of magnetism, powered by a local electrical energy store, is responsible for the torque applied to its wheels. The only direct byproducts of motion in this system are heat and noise. And it is more efficient – the ratio of useful output to input energy – at that task: c80-90% for an AC motor versus c30-40% for an IC engine. Critics frequently remark that its unique production processes, specifically building batteries, are dirtier, so emissions are merely being shifted upstream rather than eliminated. That statement is true, but it overlooks how clean this method of motion could be. The IEA estimates the incremental carbon 'debt' from an all-electric light vehicle's manufacturing stage is repaid after two years of normal usage, i.e. 35 kilometres per day. At this point, the asset is typically good for another thirteen years, over which its total emissions could be as much as 70%

lower, assuming global power systems evolve per stated policies, such as coal to solar, nuclear etc.

The all-electric propulsion system is the consensus technology to comply with current and future fleet emissions standards. In 2025, 14.2 million new light vehicles of this type were sold. That was 27% higher than in 2024 and represents 17% of global sales volumes. The sales proportions in China, Europe and the US were 32%, 17%, 7% respectively.

Management at certain automakers proclaim that the demand for this propulsion system variant is lower than necessary to comply with present policy. In Europe, for example, a share of new vehicle sales between 80% and 90% is calculated by the International Council on Clean Transportation to be required by 2035. "If I had asked people what they wanted, they would have said faster horses," as Henry Ford is alleged to have remarked. This conviction – that the market is blind to innovation – was echoed by former Volvo Cars Group CEO Jim Rowan, who neatly listed the benefits of all-electric propulsion at its 2024 Capital Markets Day: "Less noise, less parts, less vibration, less servicing, zero tailpipe emissions and eventually less expensive than cars. It also enables the bigger shift around technology."

The downtime of a vehicle with an IC engine can be relatively short. This is all-electric propulsion's largest source of friction, amplified by subscale recharging infrastructure away from home and additional variable costs that automakers cannot absorb. It is conceivably liveable for most: average daily driving distances are c35 kilometres, equivalent to one-tenth of the average all-electric range of models sold in 2025. But further innovation is needed in electrochemistry (battery C-rate, volumetric energy density) and power electronics.

In isolation, hydrogen fuel cells seem complementary to traditional charging networks. They are clean and could act like

an on-board charger, reducing downtime; electrical power is siphoned from a controlled reaction between ambient air and hydrogen (compressed, stored somewhere in the vehicle), which ejects just water vapour. The technology's primary drawback is that hydrogen storage infrastructure is not readily accessible or cheap to build out. It also reduces the vehicle's well-to-wheel efficiency. Nikola Corporation, the now-defunct heavy-duty truck manufacturer, and its providers of capital learnt these lessons the hard way.

We contend that access to the latest digital features and sophisticated driver assistance systems, certain of which are approximating driving robotisation, is a growing organic incentive to go all-electric. Their critical requirement is a consolidated electronic and electrical architecture; the foundational structure for in-vehicle componentry (computers, microcontrollers, actuators, wire harnesses, networking technologies etc), it is like the nervous system in an animal. Technically, it is feasible for production platforms compatible and/or optimised for IC engines – see Volkswagen's brand-new China Electronic Architecture – but is scarcely available because engineering attention has been diverted elsewhere and the lifecycles of 'old world' assets were unexpectedly extended.

Recent corporate actions reveal tens of billions of euros of impaired capital targeting the decarbonisation of road transportation via electrification. Stellantis, for example, in February 2026, announced non-cash losses on related non-financial services assets of €22bn. The calculations behind these amounts reevaluate the duration and scale of the path to net-zero GHG emissions. All-electric propulsion still seems to be the uncontested end point. The task for us, and market observers, is to determine where the value chain's original estimation error resides, whether it was knowable and where the final answer lies.

Uphill Climb

The Democrats' Senate Problem



Fred Hill

*Energy
Research*

The conventional wisdom in US politics is that midterm elections, which determine control of Congress and can give the opposition an effective veto over the president's agenda, are invariably challenging for the president's party. And while one chamber of Congress, the House of Representatives, has been widely expected to flip to the Democrats in the 2026 midterms, Republicans had every reason to be confident about maintaining control of the Senate.

While Democrats 'only' need to take four Republican seats to gain control of the upper chamber, there are few among the 35 up for election in the Senate this November that are realistically winnable.

The current Senate is made up of 53 Republicans and 47 Democrats, leaving Democrats four seats short of control. To build a majority then, the Democrats must first defend two seats in states Donald Trump won in 2024, Georgia and Michigan. They then must win not only in North Carolina, a swing state Trump won by only 3% in 2024, but also in Maine, which Kamala Harris won in 2024 but has a strong Republican incumbent in Susan Collins. Collins has defied political gravity before, beating her Democratic opponent in 2020 even as Joe Biden took the White House. Democrats are hoping even she cannot withstand a more Democratic-leaning year.

If the Democrats hold their two vulnerable seats, flip North Carolina

and finally take down Collins, the two remaining pick-ups required are even more challenging. The remaining Republican seats up for election are in states which were won by Trump in 2024 by a margin of at least 10% and that voted Trump in 2016 and 2020 – hardly alluring territory. The states widely considered to be the Democrats' best three prospects for these two pick-ups are Alaska (Trump +13% in 2024), Texas (Trump +14% in 2024) and Ohio (Trump +11% in 2024).

And this uphill battle is unlikely to be a one-off 'bad year' for Democrats. Rather, it is the result of two trends which are combining in an increasingly brutal way for the party.

Firstly, swing voters and swing states are increasingly scarce. For several reasons, including the divisive power of social media, the way Americans consume news and President Trump's categorical reshaping of the Republican party and its message, there are far fewer 'crossover' voters than before. Today, voters are more likely to be loyal to their preferred party no matter the election year, and no matter the individual candidates. As a result, most House and Senate seats are less competitive than they were historically. In 2026, only eleven Senators represent states won by the other party in the most recent presidential election, only one of whom is a Republican. By comparison, 25 Senators were in this position in 2006, including nine Republicans.

Secondly, small states have a disproportionate share of Senate seats as every state is represented by two senators, whether it is California (population: 39 million) or Wyoming (population: 600,000). Troublingly

for the Democrats, the Republicans have become increasingly dominant electorally in smaller states: even as Donald Trump lost the electoral college and popular vote in 2020, he still managed to win 25 out of 50 states.

Combined, these two trends are beginning to give the Republicans a near-insurmountable advantage in the competition for control of the upper chamber. In theory, Republicans can now win and maintain a Senate majority without gaining any state lost by Donald Trump in 2020.

Yet despite the above, there is still a real possibility the Democrats might win back the Senate in November. Control is now seen as a genuine toss-up by election forecasters – including Nate Silver, whose highly respected statistical model has the Democrats narrowly favoured. Meanwhile, in the betting markets, Democratic chances of success have risen from less than 20% in early 2025 to around 50% today. It is some transformation.

So, what has gone wrong for Republicans? Firstly, Donald Trump has promoted poor candidates. The most prominent example was his endorsement of State Attorney General Ken Paxton in Texas, which led to his victory in the Republican primary over four-term incumbent Senator John Cornyn, who was seen as a far safer bet. Paxton faced a securities fraud indictment in 2015, allegations of adultery and accusations of bribery while serving as Texas' attorney general, with his own deputies reporting him to the FBI in 2020. This resulted in a rare same-party impeachment, with the Republican-controlled Texas House of Representatives voting to remove

him from his role. At the time of writing, Paxton is trailing in the polls to Democrat James Talarico, putting a strongly Republican-leaning state into play.

While there are several examples of Trump endorsing Republican candidates who are not widely seen as the most electable option, including in states where it matters, the most important example (beyond Paxton) is North Carolina. There, Trump successfully pressured incumbent Senator Thom Tillis into dropping his re-election bid, seemingly in revenge for Tillis' 'no' vote on the 'Big Beautiful Bill'. The candidate is the untested Michael Whatley, who, with no electoral experience, is facing the former Democratic Governor Roy Cooper in a state with only a small Republican lean.

Next, President Trump's dominance of the Republican party has made it harder for candidates to put space between themselves and him. Distancing yourself from the incumbent president and your party is an age-old tactic in more competitive states or those which tend to lean against your party. Traditionally, the president and other party leaders have quietly tolerated this for electoral purposes, but Trump is not a man known for quiet toleration – though he is quietly tolerating Collins, knowing how important her seat is.

In return, she doesn't mention him, but most Republican swing-state candidates are more vocal in their support for the president, despite his approval rating being underwater.

Conversely, Democratic candidates have been willing to criticise their leadership, for electoral advantage or otherwise, which of course is much easier to do if you are not holding the presidency. In this cycle, Roy Cooper (North Carolina) has run adverts boasting of 'bucking his party', while Sherrod Brown (Ohio) has consistently criticised his party as 'elite' and having a 'toxic reputation'.

Beyond Republicans nominating poor electoral performers, the



Spectre at the feast

Democrats have benefited from other unexpected slices of luck. The most important initially seemed like bad news for their Senate prospects. A credible allegation of sexual assault was made against their candidate in Maine, Graham Platner. Platner had already scraped through several other scandals, including the discovery he had a tattoo based on Nazi imagery (he claimed not to have appreciated the connotations at the time). While this appeared to be terrible timing for the Democrats, who had just confirmed Platner as their candidate, the allegation surfaced early enough for the state party to hold a nominating convention and replace him with Troy Jackson, a former state Senate President untarnished by scandal. Had the allegation been weeks later, the party would have been stuck with Platner and would either have had to cede the race to the Republicans or stick by their candidate, suffer the brickbats and almost certainly lose.

In Alaska, alongside Democrat Mary Peltola who won the open primary,

former teacher Dan J. Sullivan will appear on the November ballot alongside the two-term Republican incumbent, Dan S. Sullivan. Despite Republicans attempting to have the 'other' Dan Sullivan removed from the ballot, they have failed. While Alaska uses ranked-choice voting, many voters do not fill in more than one box, and it is possible that in a close election enough voters could vote for the 'wrong' Dan Sullivan to alter the outcome. It is unclear to what extent this is luck, or the result of political manoeuvring.

Strikingly, despite these developments, the race for the Senate remains too close to call. This is important, as it is clearly not a sustainable situation for the Democrats that, even in a favourable midterm year, everything needs to break their way for the party to even have a shot at winning the Senate. Even if control of the Senate does flip in November, the party must adapt. That is politics, and in US politics there is no prize for winning the popular vote.

The People's Currency

Liberalising the Renminbi



James Invine

*Banks
Research*

China has spent the summer tightening various ways in which its citizens invest. First came the securities regulator's promise to "resolutely crack down" on online brokers Tiger, Futu and Longbridge, saying it would "confiscate all illegal gains... while imposing severe penalties". Subsequently, it started enforcing a 20% tax on offshore trusts and offshore insurance policies. Shares in HSBC, Standard Chartered, AIA and Prudential all fell substantially on the days these headlines appeared, albeit the banks have since recovered. Given the (gradual) internationalisation of the renminbi (RMB) has long been touted as the next step in China's development, why is the country now apparently leaning against this process? In this essay, we show why China has capital controls, what the recent regulatory actions mean and the way forward.

While China's capital controls have supported its growth agenda, they have also left it reliant on the US dollar. China accounts for roughly 15% of world trade but less than 5% of global trade is settled in RMB. Even of China's own trade, only 30% is settled in the local currency. Increasing these numbers would allow China to become less dollar reliant and open a new source of demand for Chinese assets, if China were to pay for more of its imports with RMB and foreign exporters were to hold on to the currency. Liberalisation of the currency regime

would allow Chinese households to invest in international assets.

There is also a political angle. For trade denominated in US dollars, final settlement must run through a bank with access to the US's currency clearing system, which brings the transaction into the US regulatory (and sanctions) purview. Even if the exporter and importer are both not using it, the dollar is often an intermediate currency, given most currencies' most liquid exchange is with the dollar. Following the invasion of Ukraine, Western allies applied sanctions partly by cutting off much of Russian banks' dollar/euro clearing; this tactic would not work if the trade in question did not touch those currencies.

China wants to enjoy some of the benefits of a liberalised currency, but a complete liberalisation is not on the agenda because it also wants to retain the protections of capital controls given the experience of some of its Asian neighbours in the late 1990s. A full liberalisation of the currency would also risk unintended consequences given the size of the numbers involved. For example, Chinese households have \$25trn of bank deposits, equivalent to 120% of Chinese GDP. Both these numbers are much higher than the US equivalents of \$13trn and 45%. Even a modest international diversification by Chinese households would drive meaningful flows and could be a popular option given the other main domestic investment (property) has slumped.

The compromise is to use Hong Kong as a bridge to the rest of the world. More than 60% of China's inbound and outbound direct investment already flows through the territory and roughly three-quarters

of offshore RMB payments clear there. The various 'Connect' schemes let money pass to and from the mainland under quotas that can be widened (or narrowed) at will.

Hong Kong is a natural location for Chinese companies' international treasury centres. Its use helps to create trust on both sides of treasury activities, especially given the large role of HSBC and Standard Chartered. None of this is new, but the internationalisation effort has moved up the political agenda: previous five-year plans called for 'prudent' RMB internationalisation, but now, in the 2026-30 plan, the 'prudent' has been dropped, suggesting greater urgency.

This brings us back to the summer crackdowns. China is not backtracking on its commitment to RMB internationalisation, but is insistent that it happens on its terms and under its control. As the chair of China's securities regulator commented, "We will open the front door while blocking the side doors".

The international rise of the RMB is clear but, if anything, has been slower than expected. Back in 2011, HSBC talked of reforms and progress by 2015-17 that would pave the way for full convertibility. Clearly, we are behind that timetable but there is some acceleration: 30% of Chinese trade is now denominated in RMB, against 12% in 2017.

In a recent presentation, Standard Chartered benchmarked the current position of the RMB with the position of the US dollar in 1890. Over the following sixty years, the dollar rose to account for 50% of global reserves (mostly at the expense of sterling). We do not expect the RMB to take over as the world's reserve currency, even on a sixty-year timeline. But we are likely



Life imitating art

to see continuing fragmentation, with a dollar-centric core alongside a China-adjacent bloc.

Half a decade of tariffs, wars and sanctions has presumably increased a number of countries' interest in an alternative to the dollar. At the same time, companies are reconfiguring their supply chains to cope with these factors and implement the lessons learned during COVID about overreliance on a single country within a manufacturing supply chain. No doubt they will also consider the same lessons in relation to financial supply chains. These trends will only be reinforced by Chinese companies' international expansion, with outbound foreign direct investment rising from virtually zero in 2000 to over \$150bn in 2025.

Beyond political considerations, corporate and institutional users

want the plumbing to work. The faster, cheaper and more reliably a payments system works, the better. This is why China is investing so much in digital, for example mBridge (a multi central bank digital currency platform designed to enable real-time peer-to-peer cross-border payments and FX transactions) and the digital RMB, which has already processed more than \$2.4trn worth of transactions and is the world's largest live central bank digital currency trial.

While these projects involve state-sponsored initiatives, Hong Kong has become a test bed for private digital cash. Earlier this year, its Monetary Authority approved the first two licences under a new stablecoin regime for use by HSBC and Anchorpoint, a venture led by Standard Chartered. Both banks are tokenising deposits, bonds and other financial assets. While the overall

amounts are not particularly large today, growth should continue, with Standard Chartered expecting \$2trn of tokenised real-world assets by end-2028. Agentic commerce will require money on digital rails and will therefore probably reinforce these trends. No doubt, the western financial institutions are investing heavily too but, at least, there is no sign that Hong Kong and China will be left behind.

China is not retreating from RMB internationalisation but rather tweaking it to ensure it proceeds steadily and under government control. The summer crackdown and increased urgency towards internationalisation are different sides of the same coin. While the pace will probably remain gradual, this is a long-term shift within the global financial system that is only likely to go one way over the coming decades.

Tech Talk

AI in Legal Workflows



Henry Hayden
*Capital Markets
Research*

Arriving at my first LegalTechTalk in London in June, I couldn't help noticing a palpably bull market atmosphere. Entry to the pavilion outside the O2 was flanked by a LegalAI-sponsored DJ booth, a CLM-backed 'rage room' where contestants could win prizes, and an Aperol Spritz bar with its own garden.

Inside, the tone was equally positive. The pavilion was overflowing and the number of exhibitors had soared compared with previous years, fuelled by an influx of venture capital into the sector and a proliferation of AI-native new entrants. However, my first discussion with a representative of a start-up opened with an unexpected question: when will the legal technology market consolidate?

Growth in technology spending at law firms has accelerated over the past three years. This has lifted from the long-term average of c6% to almost 11% in 2025. While an absolute measurement of investment, it also holds true on a relative basis, with IT expenditure as a percentage of law firm revenues rising from c2% in 2020 to nearly 4% in 2025.

These dynamics, taken together, have created an environment conducive to innovation, with the legal industry embarking on a period of experimentation and broad-based AI workflow adoption. To date, however, many of these solutions have focused on a flat-fee subscription model, simultaneously offering a relatively small cost to

early adopters and carrying negative incremental margins predicated on fixed revenues and highly variable, usage-dependent costs.

Over the long term, therefore, we expect new entrants to move towards usage-based pricing in order to establish a business model that will persist at scale. Legora has been the first mover, introducing consumption-based pricing in June 2026. As this strategy bleeds into the rest of the sector, the cost of AI adoption to law firms multiplies.

Hence, firms exiting an exploratory period will seek to rationalise the number of AI-native workflow-provider relationships they maintain. Meanwhile, technology budgets will begin to mean-revert, albeit to a structurally higher level than in the past, reflecting strategically transformed, more tech-centric firm construction. Both trends will underpin consolidation in the industry.

In the near term, the principal concern from new entrants is likely to continue to be generating growth at all costs. This serves two functions. It allows them to become embedded into client operations and gain sufficient scale to ensure they fall on the right side of inevitable industry consolidation.

To that end, price competition has become a feature of the industry. Many AI-native businesses cited consistent undercutting by competitors through the contract negotiation process alongside steep initial discounts to lock in longer-term contracts.

Legal technology M&A is already underway, mostly relating to talent-focused acquisitions and both horizontal and vertical technological

expansion. The two largest players, Legora and Harvey, have closed eight deals between them since March 2026; Wexler, Cadastral, Graceview, Qura and Walter for the former, and Benchmark, Lume and Hexus for the latter.

For those businesses with critical mass, longevity is a principal consideration in discussions with clients and a competitive differentiator. Law firms and corporate legal teams expect their workflow providers to remain standalone to ensure there is no risk of the loss of integrated content or financial and operational disruption to everyday business processes.

This helps the incumbents, who have history, track records and robust financial profiles in their favour. The incumbents have also been acquisitive, demonstrated by RELX's acquisition of Doctrine, Thomson Reuters' buyout of Noetica and Wolters Kluwer's integration of Brightflag – each adding selective data and functionality to the existing platform.

We note three issues disconnected from both public market valuation volatility and our own discussions with investors.

Firstly, increasing competition in the legal sector continues to weigh on the share prices of RELX, Thomson Reuters and Wolters Kluwer, with many concerns stemming from the potential replicability of content stores over time. However, new major AI entrants have shown minimal appetite to attempt to replicate the data assets of incumbents. They consider this impossible to achieve in its entirety given many datasets are outside the public domain or otherwise scarce and non-

scrapeable; it would be prohibitively expensive; and it is non-essential to their core capabilities, which are positioned outside the research workflow landscape. Yet integration and interoperability with the data providers was viewed as essential to everyday workflows and reliable performance – both Harvey and Legora have integrations with RELX or Wolters Kluwer in some capacity. As Legora has shifted into legal research, its integration with Wolters Kluwer underpins the majority of its US content sets. Furthermore, sales discussions with clients did not entail displacement, but complementarity between AI workflow solutions and research data and analytics providers.

Secondly, Legora and Harvey each highlighted business differentiation and defensibility against the charge of being superficial ‘model wrappers’. This squares with commentary from incumbents that model agnosticism enables AI tools to delegate subtasks across various model providers, optimising for cost and performance quality.

More recently, Thomson Reuters departed from this approach, developing its own proprietary model, Thomson, which it claims can deliver optimisations as a function of data tagging and storage sophistication. Meanwhile, Harvey, after initially scrapping plans to develop its own model, has reignited the strategic push to lower its own cost base, publicly stating it views the frontier labs as its greatest competition.

Thirdly, despite market focus on model context protocol (MCP) monetisation across data and workflow providers, this dynamic remains largely immaterial in the legal technology space, particularly when compared with financial services terminals.

This partially reflects a new technological sophistication within the legal industry, which historically lagged in its quantum of investment, and the more relationship- and text-focused nature of legal workflows. While Legora, Harvey and Thomson



You know things are going well

Reuters all offer MCP integrations with Claude for Legal Services, adoption remains modest.

Furthermore, the performance of AI capabilities in external uncontrolled environments remains a critical concern for specialised providers. Hence, LexisNexis opted not to offer an integration with Claude on the basis of lower-than-desired precision, instead incorporating Anthropic’s workflows into its own solution. While underlying data scarcity, rather than ownership of the user interface, defines pricing power, we view this approach favourably. It simultaneously offers comprehensive functionality to clients and preserves direct customer relationships in parallel with enhanced output control and visibility.

As consolidation within legal technology accelerates, we expect to continue to see diverging approaches. Incumbents, after a period of AI-enabled acquisitions, have reached sufficient capability and have turned their focus towards complementary assets such as content breadth expansion.

Among startups, Harvey has pursued horizontal functionality, seeking to become a ‘Swiss Army knife’ covering fragmented vertical

use cases and embedding itself broadly into lawyer workflows. Conversely, Legora has adopted a sharper verticalisation approach, focusing on agentic operating system refinement and deep service-led implementation. Both aim for a differentiated edge that embeds permanence even as funding tapers and clients rationalise vendor relationships.

The consensus at LegalTechTalk was the market will condense around the incumbents and the handful of new entrants who have amassed critical scale. However, despite early signs of acquisitive momentum, underlying technology budget and law firm profitability growth continue to fuel adoption via experimentation and the entry of new players into the legal technology space.

We therefore assess the current wave of deal flow as representative of opportunistic talent and capability buildout, rather than a systemic and broad-based industry consolidation. In the near term, as long as funding persists, fragmentation will continue to be a feature of legal technology and AI investment. However, the notion of industry consolidation is increasingly a question of ‘when’, not ‘if’.

Research Matters

A Selection of Recent Analysis



APPLE

Open Source to the Rescue

Technology Hardware
Timm Schulze-Melander

63

17 August 2026



OHB

A Pure-Play Space Prime

Aerospace & Defence
Joe Orchard

79

5 August 2026



One Prompt Entrepreneur

Meta, Shopify

Banks (FinTech)
Dominic Ball

135

21 July 2026



Emily Cooledge, Aisha Austin

Burberry, Hermès, Kering, LVMH, Moncler,
Prada, Richemont

7 September 2026

Luxury Goods Accentuation

Successful brands are dynamic and responsive. Yet change takes time. The Rothschild & Co Redburn Luxury Product Survey is a unique dataset, simultaneously vast and granular. It provides insight into the intricacies at play across companies, examining category, region, price point and customer group. Changes we identified a year ago are maturing and being accentuated. This reinforces confidence in industry growth inflecting positively if gradually, and with material implied divergence in performance.






Michael Taylor

27 August 2026


Economics

A Volatile Monetary Future


Meeting inflation targets is becoming more difficult. AI, demographics and digital money are set to be disruptive. AI will affect aggregate demand and supply, which are subject to uncertainty, implying an evolving neutral rate of interest. Ageing populations will undermine the drivers of credit growth and money creation, while digital money – stablecoins and central bank digital currencies – can affect not only monetary policy but also the banking system and financial stability.



13 July 2026



18 June 2026



9 June 2026

Letter from America

Hidden Heroes of Gameday



Connor Cheynet
US Equity
Sales

Catch the game this weekend? If so, the chances are you probably remember the game-winning goal, that last-ditch tackle or a controversial call from the officials. What you almost certainly don't remember is the field itself. Nor will you know the crew who were out there hours before kickoff, coaxing it into looking perfect, as though it had always been that way. Every groundskeeper and grounds crew knows if you do your job perfectly, few will realize you do it at all. The only time your work receives

for next to nothing, I found myself preparing for a job that would teach me much more than how to mow grass.

My first day began not long after a five AM alarm, as I hustled to be through the shop doors by six. The only person who beat me there was a guy who looked like he'd been doing the job longer than I'd been alive. He was the senior man on the crew, and after asking my name, he handed me a nickname no HR department would ever sign off on. It was an experience that prepared me perfectly for the rough and tumble of a dealing desk.

The first lesson I learned was that 'groundskeeper' is a word that only hints at everything the job entails. Fields are cut to a precise height and in deliberate patterns; the stripes you see on television are the result

events and everyday play chewed them to pieces.

It is skilled work, and it is relentless. Greenskeepers work in silence year-round in all weathers, against a clock counting down to a kickoff that will not move for anyone. Rain is no excuse. It merely makes the job trickier. Working by feel matters more than the people in the stands realize. A poor surface is not just an eyesore; it can be a twisted ankle waiting to happen and an edge quietly handed to one team over another. Done right, nobody notices. Done wrong, it can cost you.

If that sounds like I'm overselling, blame my job in sales. But I wouldn't be a stockbroker if I didn't have some numbers to back my call. Turfgrass is the largest irrigated crop in America, covering more ground than

“Converting the AT&T Stadium in Dallas took roughly 45,000 hours of labor and 15,000 tons of materials”

attention is when something has gone wrong.

My grandfather, after thirty years of coaching collegiate sports, began his second career as a groundskeeper for the athletic department at Virginia Tech. He loved the work but made it clear to me there was little glory to be had, at least, it wouldn't be like the glory of his coaching days. Given I was an athlete myself, he thought it would be character-building for me to join him in the summers.

After promising the head groundskeeper I'd work full-time

of which way the blades of grass reflect light, not how they are dyed, as some assume. We painted every line and logo ourselves, usually with nothing more than a length of string pulled tight and a walk-behind paint machine. Those lines are not decoration. A ball an inch out of bounds, a foot on or off the line, a fourth-and-one measured to the blade can decide a game, a tournament, even a season. We laid sod, aerated, top-dressed pitches, patched divots and put entire fields back together after weather, fan

irrigated corn. There are between 300,000 and 750,000 sports fields in America, spread across 300-plus professional facilities, 3,500 colleges and universities, 7,800 parks and recreation districts, and 20,000-plus school districts.

Nowhere was the craft more visible than at this summer's World Cup. Eight of the sixteen host stadiums normally play on artificial turf, forcing crews to think outside the box and transform acres of plastic into living, championship-grade grass for the world's largest sporting event.

To put this into perspective, FIFA spent \$5m and the better part of eight years on the problem, leaning on groundskeepers and their crews across the country to turn plastic into gold. In domed stadiums, where sunlight doesn't reach the grass, crews kept surfaces alive under LED grow lights burning twelve to fourteen hours a day. Converting the AT&T Stadium, host of the semi-final in Dallas, took roughly 45,000 hours of labor and 15,000 tons of materials. As ever, the devil lies in the detail, and the precision of this feat is measured in millimeters.

Those details are laid out in FIFA's Stadium Guidelines, which measure everything from grass height to ball behavior. Not unreasonably, the world's best players expect a flawless surface on which to perform in front of millions. Greenskeepers are held to a ruler-strict standard. A window of roughly 20-28 millimeters to be exact.

It doesn't stop at FIFA. College football, the NFL, Premier League, the PGA, all have targets that require grounds crews to be nearly perfect. My team and I at Virginia Tech chased this level of meticulous care every day. The pressure is intense but, as with stockbroking, you grow to love it.

The years I spent as a sports-field groundskeeper are a big part of why I'm in broking today. I wrote my college essay on my summers mowing grass, moving dirt and painting fields in the sun, which, via the butterfly effect, landed me at Rothschild & Co Redburn.

Early on, I thought groundskeeping might be the career for me. I held out as long as I could but eventually left for a seat in the markets. Funny thing is, I keep finding it's the same job in a different uniform. Both run against a remorselessly ticking clock and a market open is another form of kickoff. But the parallels aren't all so serious. Stay in either seat long enough and you're bound to meet some characters who teach you things that were never in the job description, whether it's how to drive a stick shift



in the shop truck you 'borrow' over lunch, or learning old-school broking from Simon Baker over a few airport beers. The characters are half the reason you stay.

The experience of groundskeeping taught me what a trading floor would demand, long before I knew I'd need it: the early mornings and long days, anticipating change and a genuine passion for the parts of the job nobody sees. The Virginia Tech grounds crew probably have no idea they trained a stockbroker, but I hold myself to the same standard every morning, whether the medium is

grass or my Bloomberg screen.

So, here's my one ask: this Saturday, when you settle in to watch the Premier League or college football and the camera sweeps across a pitch so flawless it looks as if it's been painted, don't just admire the game. Spare a thought for the surface underneath, and the people who were out there for months to make it look that way.

They won't merit a line in the match report or a word from the commentators, but tip your cap to the groundskeeper and his crew. The hidden heroes of gameday.

United

The FIFA World Cup



Daniel Sepahi
*Software & Cloud
Research*

Few sporting events have achieved the cultural reach and emotional heft of the FIFA World Cup. Every four years, football ceases to be merely a game and becomes a global spectacle through which nations express pride, communities reaffirm their identity and millions share a common experience.

Its influence extends far beyond the ninety minutes of a match. Since its creation in 1930, the World Cup has reflected political and social changes, from nationalism and war to migration and globalisation.

For diaspora communities, it carries a deep significance. It can transform a distant homeland into something emotionally present, allowing people to reconnect with their heritage and pass cultural belonging from one generation to the next.

The World Cup is therefore not simply the history of a sporting tournament. It is also a history of identity and belonging – something the current FIFA president, Gianni Infantino, appeared to forget during his recent attempt to sell a stake to private equity.

The idea of a global football tournament emerged in the early twentieth century, as the sport expanded rapidly beyond its British origins. Football had already become hugely popular throughout Europe and South America, but international competition was largely confined to the Olympic Games. However, Jules

Rimet, president of FIFA, believed football could unite nations through peaceful competition and advocated an independent international tournament. His vision became reality in 1928, when FIFA approved the establishment of the World Cup. Two years later, Uruguay hosted the first tournament.

The inaugural competition was remarkably non-global. Only thirteen nations participated, and every match was played in Montevideo. Yet the tournament established a tradition whose reach and importance would eventually extend far beyond the expectations of its founders. Uruguay defeated Argentina 4-2 in the final, becoming the first world champions. The victory produced enormous national celebration, illustrating one of football's enduring qualities: its ability to transform sporting success into a shared expression of national identity.

The subsequent history of the World Cup unfolded in parallel to the turbulence of the twentieth century. The tournaments of 1934 and 1938 took place amid the rise of Nazism and growing international tensions, before the Second World War brought the competition to an abrupt halt. When it returned in Brazil in 1950, the world had been fundamentally altered. Its return was therefore more than a sporting revival. It represented a tentative restoration of international contact after years of division.

As the World Cup developed, its relationship with national identity became increasingly powerful. A national football team offers a uniquely visible representation of a country. During a tournament,

people who might otherwise be separated by age, class, wealth, politics or religion can gather behind the same colours. The national team becomes a symbol through which millions experience a collective belonging. The World Cup consequently gives nationalism – even patriotism, famously ‘the last refuge of the scoundrel’ – an emotional dimension few other institutions can replicate.

Yet national identity has become increasingly complicated by migration. Millions now live in countries different from those of their parents or grandparents, creating communities whose identities cannot always be contained

Even unites the neutral



within a single nation. A person might be born in Germany, move to Britain, be raised within a Nigerian family and maintain the cultural traditions of their homeland while participating fully in British society. Another might be born in France to Algerian parents, growing up simultaneously within French society and an Algerian cultural heritage. For such individuals, identity is not necessarily a choice between two alternatives. It can be plural, layered, constantly evolving.

The World Cup provides one of the clearest public expressions of this complexity. For diaspora communities, supporting a national team can be an intensely emotional way of reconnecting with a homeland. Someone who has never lived in the country represented on the pitch may nevertheless feel a profound attachment to it. A flag displayed from a window, a football shirt worn through the streets or a family gathering around a television can become an act of remembrance

and belonging. The World Cup allows people to participate in a national culture that may otherwise feel distant. Thus, football becomes a form of cultural inheritance.

In this way, the World Cup challenges the assumption that nationality is determined solely by birthplace or citizenship. Football supporters demonstrate that belonging can be emotional as well as legal. A person can feel deeply British while also feeling a profound connection to Ghana; they can be French and Algerian, German and Turkish or American and Mexican.

These identities do not necessarily compete with one another. Instead, they reflect the realities of multicultural societies shaped by migration. The World Cup provides a rare public space in increasingly challenging times when multiple identities can be celebrated rather than concealed.

The players themselves often embody the complexity of contemporary identity. National teams are shaped by histories of migration or colonialism. Players may have parents from different countries, have been born in one nation and raised in another or have developed their careers across several continents. The apparent simplicity of one nation playing another therefore conceals a more complicated reality. Beneath the flags and national anthems are individuals whose histories often cross the very borders the competition appears to reinforce.

The tournament also creates temporary communities within the host cities and countries that watch it. Public squares, restaurants, homes and pubs become places of collective celebration.

A city can, for a brief period, contain multitudes, and countless versions of home. This is one of the World Cup's most fascinating contradictions: it is a competition between nations that simultaneously reveals how interconnected those nations have become.

Its social significance extends into politics and economics. Hosting the tournament can bring in tourism, investment, infrastructure and international recognition, but it can also generate controversy concerning public expenditure, workers' rights, illegal foreign labour, inequality and environmental impact. These debates demonstrate that football is no longer simply entertainment. The World Cup commands enough attention and resources to influence entire societies.

Ultimately, the enduring power of the World Cup comes from its ability to transform something as simple as a football match into a profound expression of human identity. For one person, it may represent patriotism; for another, family; for another, memories of a country they left behind; for members of a diaspora, it can make a distant homeland feel unexpectedly close. A national shirt can carry the memory of a grandparent. A song sung in a crowded stadium can preserve language and tradition. A single goal can unite people separated by thousands of miles and several generations.

The World Cup's growth has mirrored the expansion of global communication, the movement of people across borders and the increasing complexity of national identity. What began in Uruguay in 1930 as a tournament involving thirteen teams has become a global cultural phenomenon capable of touching billions.

Its importance lies not simply in determining which country is strongest on the football pitch, but in the emotions and identities that the competition brings to the surface. For a few weeks every four years, borders remain visible but are freely crossed. The World Cup allows nations to celebrate their histories, diaspora communities to rediscover their heritage and individuals to belong to something larger than themselves.



Octopus Farming

Our Diet, Their Intelligence



Neil Tyler
*Business Services
Research*

I can trace the origin of one of my family's more frequent dinner-table disagreements to a summer holiday snorkelling excursion. My daughter and I lost track of time as we watched an octopus going about its daily chores. The creature was constantly changing shape and appearance depending on the task at hand. It soon became clear it was also simultaneously monitoring its two spectators, changing colour and seemingly vanishing if we came too close before reappearing in the same spot when the perceived threat lessened.

By the time we returned to the beach, my daughter was incensed. Recalling that one of us, on the same holiday, had ordered an octopus salad at a restaurant, the recent underwater encounter had showcased enough of the octopus's skills to convince her the species was much too intelligent to occupy a spot on any menu. Anger turned to outrage as some hasty Googling of 'octopus intelligence' led to her ongoing (albeit so far unsuccessful) opposition to plans to build the world's first commercial octopus farm in Gran Canaria.

Ever since its unveiling, several years before, campaigners, scientists and commentators have condemned this project. The objection most frequently raised was the same: octopuses are too intelligent to be farmed. Our brief encounter suggests this argument is compelling. Inevitably, it raises a much harder question: where

should we draw the line? And are we drawing it consistently? I doubt our family situation is unique, but this has become a topic of considerable debate.

With respective positions becoming entrenched, we appear no closer to agreement. On one side, my snorkelling companion is increasingly questioning established livestock farming practices more broadly. In the opposite corner, the trainee vet is presenting the commercial pressure on livestock farmers to meet consumers' demands and still make money. The youngest of the three siblings has armed herself by choosing to research the topic for an A-level project.

The challenges begin as soon as we try to measure intelligence. In practice, the definition is built around human cognition. The tools we use to assess animal intelligence – problem-solving tests, mirror self-recognition, social learning experiments – are designed by humans and measure performance against human criteria. An octopus that unscrews a jar to reach food scores well. An animal that simply huddles with its group to stay warm might score poorly, even if that collective behaviour is, in its own ecological context, as sophisticated.

There is also a basic issue of research volume. We know a great deal about the cognitive lives of pigs and cattle because we have studied those species most extensively. Poultry, despite being farmed in colossal numbers, has received far less scientific attention. The result is that our perception of a species' intelligence partly reflects how much effort we've put into understanding it, not necessarily how intelligent it is. An

animal we have not studied closely is one we can more easily dismiss.

Octopuses are a case in point. The famous example of an octopus in a University of Otago laboratory learning to short-circuit its tank light by squirting jets of water at the bulb made headlines because it was so unexpected. But as philosopher Peter Godfrey-Smith noted in his brilliant study of octopus intelligence, the researchers would never have discovered this ability if the light hadn't been left on in the first place. Intelligence often only becomes visible when the conditions reveal it.

It perhaps therefore makes sense that animal welfare science has largely moved away from intelligence as its primary measure and towards sentience. The capacity to feel pain, pleasure, fear or contentment is harder to fake and easier to test. For this reason, formal objections to the Nueva Pescanova octopus farm have focused less on the animals' problem-solving abilities than on welfare – specifically the absence of any accepted humane slaughter method for octopuses.

Most livestock farmers are at pains to point out that animal welfare and commercial farming are generally seeking the same goals. And the commercial farming industry itself has quietly begun to integrate AI-based sentience monitoring tools into its operations. These can detect lameness in cattle far earlier than a human observer. Facial recognition software is being developed to track emotional states in livestock, not purely out of concern for the animals, but because welfare and productivity are closely connected. A stressed pig, or uncomfortable hen, produces less pork, or fewer eggs.

Not everyone will agree these objectives suffice. A series of emotive posters, funded by The Animal Justice Project, recently appeared on the London Underground, attempting to draw commuters' attention to this so-called 'meat paradox' – the cognitive dissonance between caring about animals and continuing to eat them. Faced with one this week while contemplating this essay, it is difficult to justify why the prospect of an octopus farm provokes international outrage, while pigs are farmed on an enormous scale. Pigs are widely understood to be intelligent, possibly more so than dogs, to whom they're routinely compared. It would be surprising, however, if this latest poster campaign made genuine headway. A 2023 study found that people, consistently presented with evidence for animal sentience, consistently fail to internalise it.

Nevertheless, meat consumption in Europe has fallen nearly 20% since 2010. This decline may be a response to rising prices driven by stricter welfare laws. But it is also possible consumers have become more aware of the purpose of those laws. In the US, where standards have tightened less, consumption has risen over the same period. Correlation is not proof of causation, but it might suggest the London Underground poster campaign is tapping into a society slowly reconfiguring its relationship with meat.

Tradition seems to have played a huge role in drawing the line between what is and is not acceptable. In the UK, eating horse meat is broadly considered unacceptable. The 2013 incident in which horse DNA was found in processed beef products was considered a serious transgression. Yet in France, although consumption has fallen substantially, around 18,000 tonnes of horse meat are still consumed annually. The practice carries no stigma, and horse meat is widely available. Guinea pigs are household pets in Britain; in parts of Peru, they're a traditional delicacy known as *cuy*.



Chameleon

Research into public attitudes towards the octopus farm reveals how much geography shapes moral response. A large-scale survey conducted by Colorado State University across several European countries found widely varying levels of support for the project. In Spain, where octopus has long been a culinary staple, support was higher than elsewhere. In countries with less tradition of eating octopus, opposition was far stronger.

My daughter's A-level survey on commercial farming was perhaps even more telling. Respondents who accepted the principle of large-scale livestock farming still drew the line at certain species. Octopuses and dogs attracted similar levels of disapproval. Pigs, despite their well-documented intelligence, provoked the least. Intelligence, it seems, barely registers in these responses. What matters more is familiarity and the degree to which an animal is already embedded in the food system.

The species that are not farmed remain, for now, outside that process. Although, as the octopus debate shows, not necessarily indefinitely, and here the octopus may have recently scored an own goal.

Earlier this summer, crab and lobster fishermen in Port Isaac were enjoying something of a windfall as they frequently discovered octopus in their pots. Several were pondering whether to switch

permanently to catching octopus after an unprecedented 'bloom' in UK waters. One of the motivations for establishing the octopus farm in Gran Canaria was a dwindling supply from local fishermen with rising sea temperatures having caused the population around the Canary Islands to decline. But these same climate effects are behind the recent abundance around the UK.

The new arrivals had begun feasting on the traditional crab and lobster catch. And in a further illustration of their intelligence, even using the pots as larders to store their own catch until they were ready for their next meal. Initially, these imposters had been more valuable to local restaurants than their traditional catch. But their abundance has caused prices to fall and, with local demand for now at least limited, most of the excess is being exported to Spain.

The octopus farm controversy suggests our dinner-table disagreement is unlikely to be settled. The honest answer to the question of why we eat some animals and not others seems to involve very little science and quite a lot of history, and habit.

As well as being smart, the octopus has novelty going for it – the novelty that engaged the two snorkelling spectators for so long and provoked the dinner-table debate in the first place.

The Serpent in the Grove

AI in Fiction



Archie Cotterell
Editor

The literary world loves a spat. The arguments, which rarely make it beyond the book pages, so far as those that still exist, recall Wallace Sayre's memorable observation about university in-fighting, "Academic politics is the most vicious and bitter form of politics, because the stakes are so low".

While declining numbers of readers ease the literary world towards the academic, a recent row which briefly escaped its cloisters has wider resonance. Jamir Nazir's *The Serpent in the Grove* won the 2026 Commonwealth Short Story Prize, but almost immediately was accused of having been written by AI.

Ethan Mollick, a professor at Wharton School at the University of Pennsylvania, wrote on Bluesky, "Pangram flags at one hundred percent but also, come on, if you know, you know". He was joined by Nabeel Qureshi, an 'entrepreneur writer and researcher', who declared on X, "Well, this is a first: a ChatGPT-generated story won a prestigious literary prize".

Notwithstanding Nazir's claim the story emanated from his childhood in Trinidad, an online pile-on ensued. Many considered sentences such as, 'Her hair is midnight rain; her laugh is bright as zinc', or 'Sun on galvanise is a cruel instrument', when coupled with well-known AI tropes such as antithesis, tricolons, tall-poppy adjectives and earnest poeticism, to be a surefire indication of guilt.

Setting aside the truism that there is plenty of bad writing generated without AI, many 'AI indicators' have been in use since Cicero. Moreover, literary prize judges and critics through the ages have mistaken laboured lyricism for mellifluous prose. Nevertheless, the Nazir controversy highlights the looking-glass position of writers and publishers in an AI world.

Without a confession, which is unlikely to be forthcoming, it is almost impossible to prove an author has used AI (this is not to suggest Jamir Nazir did – he has been cleared by the prize judges and their judgement stands). The problem is exacerbated by the fact a not insignificant proportion of novelists admit to using it, not only for research but occasionally to find synonyms or to help rework a troublesome sentence.

This raises the question: how much use of AI is permissible in a novel? If it is used to change a single word, possibly even a sentence, most readers would have no problem. The idea, the execution, the literary choices remain those of the author. But what if it is the first sentence? Or first paragraph? What if the novelist, having reached the final chapter, asks ChatGPT to suggest five possible endings from which he or she chooses one? Where do readers stand if the author selects one of the endings but then writes it in his or her own words without further recourse to AI?

There have been attempts to define the parameters of what is acceptable. The pressure group, Books by People, argues a novel needs to be, "Conceived and written by humans, where structure, voice, and character development originate

from the author. If there is limited AI use involved, it is for spelling or formatting, not for generating or rewriting meaningful content". In short, a 'human-written' book can benefit from AI editorial support but cannot include 'meaningful content'. But what constitutes 'meaningful content'? Surely every sentence in a book is 'meaningful', or it wouldn't be there?

'Meaningful' clearly means different things to different people. To my mind, as a spare-time novelist, it is an absolute. If a sentence or paragraph isn't 'meaningful' by definition it must be superfluous. While there is obvious irony in a writer of fiction complaining about a piece of work being made up, it is subservient to the idea that, at its heart, a novel is the communion of two human beings, author and reader. Notwithstanding fiction is chaotic, primary-coloured and twanging with the neuroses of its creator, the communion demands an element of trust and integrity that cannot be forfeited or replicated by AI, or non-fiction which trembles with the slant of opinion.

Publishers are caught. Given the near impossibility of identifying offending authors, and the increasing numbers of fully AI-generated novels appearing in the online marketplace, they have resorted to 'badging' novels as 'Human Written'. The Society of Authors has weighed in, creating a logo for non-AI work, while Books by People has partnered with publishers to deliver an 'Organic Literature' certificate.

It is hardly a solution. Any certification relies on the author being truthful to their publisher. But can a purveyor of fiction be relied upon to tell the truth? Or will

“Novels are messy and unreliable, made in the image of their creators. Therein lies their irreplaceability by a robot”

it be – must it be – the truth as they see it? Attempts to watermark, use AI detectors or trace the genesis of a work through the early drafts are a start, but are yet to become foolproof. The diffusion of large language models (LLMs) makes detection harder.

There is a secondary problem, albeit for authors rather than publishers. In the same way AI-generated music on Spotify is a boon to the platform, which doesn't have to pay publishing royalties to a robot, so AI-written books benefit publishers who do not have to deal with temperamental creatives failing to deliver on time, demanding more money, drinking themselves to death or displaying any of the other peccadilloes with which they are popularly (and not always unreasonably) identified. If novels are messy and unreliable, it is because they are made in the image of their creators. Therein lies their seduction, their value and their irreplaceability by a robot.

The problem for the industry is acute. Yes, Bloomsbury, the home of Harry Potter, won a judgement in July

against Anthropic for scraping and retaining books without permission or payment in order to train its LLMs. But the 14,000 books the judgement covered is a bagatelle relative to the scale of LLM theft, as is the size of the publishers' balance sheets relative to the AI heavies.

There is more the industry can do. It should litigate more aggressively *pour encourager les autres*, require authors to sign legally binding commitments that their books are 'human written' – taken as an absolute not in part – and use technology to trace every step of the creative route. Ultimately, it will always require a degree of trust.

This is the nub. Because although the death of the author has been proclaimed for more than fifty years – initially by Roland Barthes in 1967, and Michel Foucault two years later – fiction, like all art, relies upon the suspension of disbelief and therefore absolute trust in the integrity of the creator.

It is symptomatic of a wider problem in society. Deep fakes, altered images, lack of transparency and the tumult of social media are increasingly

blurring the distinction between the real and bogus. In so doing, they have undermined consumers' trust. This impedes the empathy that is the basis of art, and of our understanding and appreciation of our fellow man.

Does this matter to the tech behemoths? No, but one day hopefully it may do. In August, 29 states led by California, Colorado, Kentucky and New Jersey succeeded in their suit against Meta, alleging it knowingly hid the truth about, and exploited, social media's effect on youngsters' mental health. The fine, \$18bn, is barely a slap on the wrist, but it may be the first small step, in parallel with Meta's promise to tighten safeguards.

Ultimately, there is a question about the balance of EQ and IQ. For hundreds of years, countries and institutions have been led by men and women hailing from a broad mix of scientific and humanities backgrounds. Today, the world is being shaped by a tiny cadre of people, all of whom have scientific or mathematical backgrounds. It is hard to believe something is not being lost.



Sir Garfield Sobers

28/07/1936 – 17/07/2026

Gary, or Garry, Sobers was best known for two extraordinary feats. He was the first batter to hit six sixes in an over in first-class cricket, and for 36 years he held the record for the highest individual score in a Test match, 365 not out.

Captaining Nottinghamshire, for whom he was the overseas professional, against Glamorgan at Swansea in 1968, Sobers decided he needed to score quickly to engineer a declaration. The unfortunate bowler, Malcolm Nash, traditionally a seamer, was experimenting with left-arm spin. Successive balls disappeared over the boundary, the last – speared in by the humiliated Welshman – out of the ground. Sobers had clearly chosen between the advice of two former West Indian greats: “Everton Weekes used to tell me if I kept the ball on the ground no one could catch me, while Sir Learie Constantine said if you hit it out of the ground no one could catch it”.

A decade earlier, Sobers had flayed Pakistan for 365 not out at Sabina Park in Kingston, Jamaica. It was a record that stood until his fellow West Indian Brian Lara scored 375 against England in Antigua in 1994. Remarkably, given he’d been playing Test cricket for almost five years, it was Sobers’ first Test century. He quickly made up for lost time, scoring 125 and 109 not out in the next game, and 142 not out, 198 and 106 not out in successive Tests in India the following year.

Although those two were the Barbadian maestro’s most celebrated feats, his finest innings was often considered to be the 254 he scored for the Rest of the World in Melbourne against an Australian attack spearheaded by the young Dennis Lillee. Sir Donald Bradman called it “the finest innings ever seen in Australia”.

But Gary Sobers’ qualities lay beyond statistics, hugely impressive though

they were. An outstanding captain who led by example, never looked down on lesser mortals and achieved series wins in England and India, and at home to Australia, Sobers was renowned for his sportsmanship. He never sledged or tried to distract an opponent, was chivalrous towards umpires and eschewed defensive play. In 1968, he was condemned by the media for declaring against England at Port of Spain to revive a Test that was dying, only to lose it by seven wickets.


Widely considered to be the finest all-rounder in the history of the game, Sobers was five cricketers rolled into one. A batsman of power and beauty, a bowler capable of extreme pace, unplayable swing, orthodox slow left-arm and wrist spin, he was also a fielder of feline grace, whether patrolling the covers with a rocket arm or in the slips, pouching the uncatchable with deceptive ease.

It was a remarkable achievement for a boy raised by a single mother in St Michael, Bridgetown, who learnt his

cricket on the streets and beaches, and made his debut for his island aged 16 and for his country a year later.

But he was not without demons. These were summoned after he was involved in a fatal car accident in 1959, which killed his friend and fellow West Indian cricketer, Collie Smith. During his final Test in England, at Lord’s in 1973, he celebrated being 31 not out at the close of play on the first day by carousing all night. The following morning, after a double brandy and port administered by his captain, Rohan Kanhai, he blasted the English bowlers for 150 before retiring hurt complaining of “stomach cramps”.

Sobers was a committed gambler who lost whatever money he made on the horses and ultimately needed to be bailed out by the Barbadian authorities. A scratch golfer, in retirement he claimed he wished he’d become a professional, the financial rewards – particularly in those days – being far greater on the greens. His loss was cricket’s gain.



Rothschild & Co Redburn Research

Directors of Research

Fahed Kunwar
Alex Brignall

Steve Chesney
Qize Ding
Joseph Simon
Michael Woodthorpe

BLT: Business Services

Neil Tyler
Oliver Davies
William Blunt

BLT: Transport & Leisure

Alex Brignall
James Goodall
Edward Lewis
Andrew Tam
Oliver Holmes
Nikunj Kaushal
Vera Lisitsina
Sam Stewart
Gavin Tennant

Consumer: Beverages and Food & HPC

Chris Pitcher
Edward Lewis
Charlie Higgs
Bingqing Zhu
Artem Chubarov
Lindsey du Toit

Consumer: Retail

Geoff Lowery
Sam Hudson
Ellen Lundstrom
Haoqing Zhu
Ravi Saroya

Consumer: Luxury & Beauty

Emily Cooledge
Aisha Austin
Kyriaki Koutta

Energy: Oil & Gas / Power & Utilities

Tony Jones
Simon Toyne
Skye Landon
Sebastian Erskine
Fergus Neve
Fred Hill

Financials: Banks and FinTech

Fahed Kunwar
Nicholas Watts
Dominic Ball
Charles Bendit
Harry Bartlett
Gonzalo López
Ellie Hall
James Invine
Michael Donian
Zeshan Shiekh

Financials: Capital Markets

Nicholas Watts
Simon Clinch
Charles Bendit
James Invine
Henry Hayden
Cara Thomson

Healthcare: Biopharmaceuticals

Simon Baker

Global Economics

Melissa Davies Michael Taylor

Strategy

George Karamanos

Healthcare: MedTech

Ed Ridley-Day
Issie Kirby
Jamie Clark
Jackson Philips
Natalya Davies

Industrials: Aerospace & Defence

Olivier Brochet
Joe Orchard
Maggie Schooley
Edward Hawney

Industrials: Automotive

Tobias Beith
Edward Hawney

Industrials: Capital Goods

James Moore
Richard Radbourne
Lush Mahendrarajah
Adam Parr
Oliver Henry
Edward Hawney

Materials: Chemicals

Tony Jones
Artem Chubarov
Mazahir Mammadli

Materials: Construction & Building Mats

Will Jones
Harry Dow
Edward Hawney

TMT: Internet

Dominic Ball
James Goodall
Joseph Barker
Zeshan Shiekh

TMT: Media

Bianca Dallal
Hamilton Faber
Ed Vyvyan
Freddie Corrin

TMT: Software & Cloud

Alex Haissl
Omar Sheikh
Nina Marques
Lachlan Brown
Daniel Sepahi
Luke Han
Aleia de Buretel

TMT: Tech Hardware

Timm Schulze-Melander
Richard Radbourne
Mike Harrison
Nabeel Aziz

TMT: Digital Infrastructure

Hamilton Faber
Max Findlay

Technical Analysis

Nick Glydon Hugo Bromell

IDEAS Systematic Analysis

Chris Bryant Devanté Wallace Joseph Worsley

Rothschild & Co Sales

Andy Preston	Max Georgiou (Financials)
Anil Sharma (Financials)	Michelle Lodge
Anna Dietrich	Millie Edwards
Brandon Henderson (Consumer)	Niroshan Ganeshan (Chemicals)
Cedric Stiegler	Oliver Prayon
Christopher Middleton	Philip Burton
Edward Hawney (Industrials)	Phil Gestrin (Healthcare)
Edmund Fitzalan Howard	Pierra Morgan
Elena Scola	Pooja Patel
Fredi Keller	Rahul Patel
Gaia Dean	Rayan Blanchard
George Shaw	Rowan Adley (TMT)
Isabella Blyth Currie	Stuart Joyner (Energy)
Josh Chapman (BLT)	Jonathan de Pracontal
Markus McNally	Janiel Nakrani

Rothschild & Co Sales Trading

+44 20 7000 2000

Andy Willson	Lisa Christou
Alexander Laux	Luc Demoulin
Andrew Quick	Reto Meroni
Blanca Segimont de la Concha	Oisin Doyle
James Kebell	Tony Atkins
KC O'Rourke	Eli Webb

Rothschild & Co Execution

+44 20 7000 2010

Benjamin Rose	Johanna von Armin
Gavin Mitchell	Nick Watson
Kathryn Stonebank	Phil Risley
John Coughlan	Will Winzor-Saile

Rothschild & Co Access

Ivana Tomić	Daisy Joseph
Jessica Perloff	Megan McNally
Katie Hodulik	Sonica Kearney
Keely Oliver	Sofia Kalogeropoulou
Lisa Traeger	William Bristowe
Jennifer Markham	Caoive Kelsey

Rothschild & Co Global Markets Solutions Limited ("Rothschild & Co Redburn")/
Registered in England No. 4710622
Registered Office: Capital House, 85 King William Street, London EC4N 7BL
Authorised and regulated by The Financial Conduct Authority

Rothschild & Co Sales Trading and Execution – US Clients

Andrew Chader	Jessica Marino
Casey McElligott	Tom Dixey
Cameron Ulreich-Power	Kevin Galvin
Emily Henderson	Steven Priestley
Finn Maloney	Timothy Graham
Jared Lechner	

Rothschild & Co Sales – US Clients

Alexa LaRue	Madeleine Falkiner
Bryan Reed (IDEAS)	Max Boudreaux
Carson Glazier	Patrick Dolan
Chris Angioletti	Peyton Miller
Dylan Levine	Rajat Chaudhury
Hunter Brown	Rowan Joseph
Ian Marinelli	Stephen Sousa
Julian Lipinski	Tim Irons
Mia Day	Will Rice
Quinn Wiseman	Cole Berglund
Conor Cheyet	Adam Levy
Ethan Yetman	

Rothschild & Co Redburn/ Member FINRA/SIPC
1251 Avenue of the Americas, 34th Floor, New York NY 10020
131 Dartmouth Street, 3rd Floor, Boston MA 02116

London

Capital House
85 King William Street
London, EC4N 7BL
Tel: +44 20 7000 2020

Frankfurt

Börsenstraße 2-4
60313 Frankfurt am Main
+49 69 667 799 82

Madrid

Paseo de la Castellana, 35
28046 Madrid
Tel: +34 919 26 85 66

Paris

11 Rue de Téhéran
75008 Paris
Tel: +33 1 73 05 40 30

New York

1251 Avenue of the Americas, 34th Floor
New York NY 10020
Tel: +1 212 803 7300

Dallas

6060 N Central Expressway, Suite 561
Dallas TX 75206
Tel: +1 212 803 7321

Boston

1 International Place, 14th Floor
Boston MA 02110
Tel: +1 617 273 8263