

TIER ONE PEOPLE · WHITEPAPER

Agentic workforces in fintech

The operators rewriting the model, and why the establishment is a step behind.

Drawn from Fintech Chatter interviews with Send Payments, Zip and Lendi. October 2026.



Executive summary

A small group of Australian fintechs is already doing what the rest of financial services is only beginning to discuss. They are not using artificial intelligence to trim headcount at the edges. They are rebuilding how their businesses run, building agentic workforces that sit alongside their people, and reinventing the operating model.

This paper draws on three 2026 interviews on the Fintech Chatter podcast, with Matt Barr, CEO of Send Payments; Soraya Alali, CEO of Zip for Australia and New Zealand; and David Hyman, co-founder and former CEO of Lendi Group, now building the AI company Mono AI. It sets their accounts against the current body of external research on AI in financial services.

Three conclusions follow. First, at the frontier AI is not simply cutting jobs. It adds capacity where people meet customers and removes the management and back-office layers that exist to move information, a split the external data supports. Second, operators are setting the frontier, not analysts or consultants. The best practice most enterprises buy is already a generation behind what the leading fintechs are doing today. Third, the constraint is not technology or capital. It is people, and in particular the small number of leaders who have run an agentic operating model. They sit at the frontier, not in the open talent market.

The frontier and the followers

The useful distinction in financial services right now is not between firms that use AI and firms that do not. Almost everyone uses it. The distinction is between firms automating tasks inside an unchanged operating model and firms redesigning the operating model itself.

Send, Zip and Lendi are in the second group. Send runs roughly 140 software agents as named members of staff. Zip has reorganised around AI-powered process, products and people, governed through its own orchestration layer. Lendi rebuilt core customer journeys around agents with humans in the loop. In each case the question was not "where can a tool save time," but "what should this business look like if agents do the repeatable work and people do the judgment."

The incumbent banks are moving more cautiously, and the public evidence shows it. JPMorgan has said its generative AI is so far augmenting roles rather than replacing them, and analysts quoted by Fortune in December 2025 described the idea of an AI-driven takeover of finance jobs as "smoke and mirrors" to date [1]. Accenture's data points the same way: around 76 percent of banks expect to increase their technology headcount because of agentic AI, even as 73 percent of US banking work time is exposed to generative AI [2]. Banks are investing heavily, but the dominant framing inside large institutions is task automation and cost, not a redesign of the model.

The split is not purely one of size. Some large operators are drawing the same line deliberately. In an essay published in March 2026, Block's Jack Dorsey and Sequoia's Roelof Botha argued that most companies applying AI are "giving everyone a copilot," which makes the existing structure work slightly better without changing it, and set that against Block's own effort to replace what a management hierarchy does with a continuously updated AI model of the business [10]. The distinction they draw, between augmenting the current model and rebuilding it, is the one that separates the frontier from the followers.

There is also a structural reason the frontier has tended to sit with the smaller firms. As Barr put it, Send was capital constrained, so using AI to go faster was a priority rather than an option. Necessity, not a slide deck, drove the redesign. The larger and better-capitalised the institution, the less impetus there is to rebuild something that still works.

Best practice moves in one direction. Operators invent a model. Analysts and consultants observe it, codify it, and package it. Enterprises then buy the packaged version. By the time a generic transformation playbook


reaches a bank's steering committee, the operators who wrote the original have moved on. The follower implements a model that is already dated, which may be part of why so many transformations disappoint.

The data behind the shift

The external research confirms three things: productivity gains are real, the move to agents is early and failure-prone, and the sharpest constraint is talent.

On productivity, Bain's 2024 survey of financial-services firms found an average self-reported gain of around 20 percent across generative AI uses [3]. Accenture estimates early adopters in banking could see productivity improvements of 22 to 30 percent over three years [2]. Send's reported lift of more than 40 percent of output per two-week sprint sits at the top of that range, which is consistent with a firm that has redesigned the work rather than bolted AI onto it.


Reported productivity gains from AI



Source: Bain (2024); Accenture; Send Payments (Fintech Chatter, 2026).


On the move to agents, the trajectory is steep but unproven. Gartner forecasts that 40 percent of enterprise applications will embed task-specific AI agents by the end of 2026, up from less than 5 percent in 2025, and that a third of enterprise software will include agentic AI by 2028 [4]. Yet the same body of forecasting is blunt about execution risk: Gartner expects more than 40 percent of agentic AI projects to be cancelled by the end of 2027 on cost, unclear value and weak risk controls, and estimates that only a small fraction of vendors claiming to offer agentic AI actually deliver it, a pattern it calls "agentwashing" [4]. McKinsey's 2025 work describes a similar gap between interest and execution: a large majority of organisations are experimenting with agents, but only around a quarter have scaled one into production [5].

The agentic shift is steep, but early



Source: Gartner forecasts, 2024–2025.

Interest runs ahead of execution




Source: McKinsey (2025); Gartner (2024–2025).

On jobs, the projections are wide and mostly theoretical. Citi's 2024 analysis found banking more exposed to automation than any other sector, with a high share of tasks automatable [1]. The UK's Financial Services Skills Commission estimates AI could automate 30 to 50 percent of tasks in most financial-services roles, and frames the larger risk as a skills shortage: the sector could lose a substantial share of its skilled workforce to turnover and retirement by 2035 regardless of AI [6]. A Zopa and Juniper study put cost savings at £1.8 billion by 2030 and around 27,000 finance jobs at risk, concentrated in back-office compliance and verification work [7], while Bloomberg Intelligence has been cited projecting up to 200,000 banking job cuts over three to five years [1]. These are estimates of automatable tasks, not confirmed losses, and the measured evidence of AI-driven

redundancies so far is thin.

Augmentation, not replacement



Source: Citi (2024); Accenture.

Across these numbers, the technology is advancing faster than the organisations around it. In our many conversations with CEOs over 2026 it appears there has been a quantum shift. Previously CEOs would forecast 12 months out and ask if what they're building will still be relevant. They're now looking out one month and wondering if it's still relevant.

Three operators rewriting the model

Send Payments

AGENTIC WORKFORCES IN FINTECH

Send Payments

CONTEXT

Cross-border payments. Around 70 staff.

AI DEPLOYMENT

Around 140 AI agents, managed as staff with job descriptions and line managers.

REPORTED EFFECT

40%+ more output per two-week sprint. ~\$5m human-equivalent cost.

tieronepeople.com



Send operates in cross-border payments, one of the most heavily regulated areas of fintech, with around 70 people. It runs approximately 140 AI agents and manages them as staff. "Every agent has a name, has an avatar, has a resume, effectively has a job description, has targets, gets coaching and has a line manager," Matt Barr, CEO, said. One agent, Barry, reviews every account-manager conversation for compliance and escalates to a human manager, Charlotte, who initially feared the agent would replace her and now considers it essential. Send describes itself as a human-led business: "We're not going to get replaced by AI," Barr said, "but it's certainly helping us run a highly productive and less risky business." When the firm priced its agents as human roles, a recruiter quoted more than five million dollars a year.

Zip (Australia and New Zealand)

AGENTIC WORKFORCES IN FINTECH

Zip (ANZ)

CONTEXT

Consumer credit and payments.

AI DEPLOYMENT

Agents across process, products and people, via the ZAI orchestration layer.

REPORTED EFFECT

Deployment timeframes cut from months to weeks.

tieronepeople.com



Zip has moved from experimentation a year ago to agents working across the business. Soraya Alali, CEO of Australia and New Zealand, frames the programme in three parts, AI-powered process, products and people, and describes "live agents working side by side... right throughout the business," coordinated through an orchestration layer the company calls ZAI, which controls which role has access to which tools and monitors usage at the token level. Zip began with its people, building AI fluency across staff with learner paths from junior to senior and sponsorship from the top. Deployment timeframes, Alali said, have compressed from months to weeks.

Lendi and Mono AI

AGENTIC WORKFORCES IN FINTECH

Lendi Group

CONTEXT

Mortgage platform. \$107bn loan book, 1,350 brokers.

AI DEPLOYMENT

Guardian runs with humans on the loop and in the loop.

REPORTED EFFECT

Australia's most ambitious financial-services AI transformation.

tieronepeople.com



David Hyman spent thirteen years building Lendi into a mortgage platform with a loan book above \$107 billion, 1,350 brokers and more than \$350 million in annual revenue. Its AI transformation, Project Aurora, is among the most ambitious in Australian financial services. Its product Guardian uses two modes that Hyman is careful to separate. With humans on the loop, the agent runs the customer interaction from end to end and a person steps in only when a guardrail is hit or the customer gets stuck. With humans in the loop, people sit at deliberate points where trust, rapport and judgment are built, such as a product recommendation. Guardian uses both: on the loop to keep the process moving, in the loop where the decision carries weight. Hyman's framing is the clearest statement of the thesis: "AI has moved from being a technology problem to being a human change and business model problem." Having led that transformation inside a lender, he left the CEO role to build Mono AI, a company whose purpose is to help other businesses become AI-native. The frontier operator is now selling the frontier.

What the operators do differently

Across the three businesses the same practices recur, and each has support in the external research.

They manage agents as members of the workforce. Naming, job descriptions, coaching and line management force the organisation to treat an agent as accountable work, not a feature. This is what lets a human work with an agent rather than around it.

They keep people on the decisions that carry risk. Lendi's humans-in-the-loop checkpoints and Send's refusal to let an agent approve a high-value transaction reflect a deliberate line. Agents monitor, check and surface; people decide. In regulated flows, trust is the constraint, not capability.

They govern before they scale. Send defined data controls and agent boundaries from the outset; Zip built governance into its orchestration layer. This is the discipline the market most often lacks: Gartner attributes much of its projected 40 percent cancellation rate to weak risk controls and unclear value [4]. Governance is not the

brake on an agentic workforce. It is what makes one durable.

They enable people before deploying agents. Zip built fluency first; Send develops agents collaboratively, with staff identifying the problem and co-building the agent. Adoption follows capability, not the reverse.

They hire for curiosity and agency. Both Barr and Hyman describe looking for high-agency, high-curiosity people, and Hyman notes founders are holding a higher internal bar because AI now does work that used to justify a hire. The profile shifts from functional experience alone toward adaptability and ownership.

From hierarchy to intelligence

The clearest account of what replaces the old model comes from an operator, not a consultancy. In March 2026, Block's Jack Dorsey and Sequoia's Roelof Botha set out a direct statement of intent: to build a company organised "as an intelligence rather than a hierarchy" [10]. Their argument is that a management hierarchy is, at root, an information-routing system, and that AI can now perform that routing. A continuously updated model of the business carries the context a manager used to relay up and down the chain, which removes the need for a permanent layer of middle management.

What the people do changes accordingly. In Block's design the intelligence sits in the system and the people sit at the edge, where judgment, ethics, novel situations and trust live. The company describes narrowing to three roles: individual contributors who build and run the capabilities, directly responsible individuals who own a cross-cutting problem with the authority to pull resources, and player-coaches who build alongside their teams in place of managers whose main task was routing information [10]. Block is explicit that it is early and that parts of this will break before they work. Early or not, a frontier operator at scale is redesigning the firm around the technology rather than fitting the technology to the firm.

This sharpens where AI exposure actually falls. The roles most exposed are not at the customer edge but in the middle, in coordination, and in back-office routing. The Zopa and Juniper study locates its projected losses in back-office compliance and verification [7]; Block locates the redundancy in the management layer that exists to move information. The pattern holds across the frontier firms in this paper. Send adds agents to expand what a lean team can do; Zip and Lendi put agents on repeatable flows and keep people on the decisions that carry risk. Capacity grows at the edge, coordination shrinks in the middle. The work closest to the customer and the hardest judgment calls become more valuable, because they are the part a model cannot do.

Why the establishment is a step behind

The advisory industry's model is structurally late. Consulting codifies what has already been proven; it is a lagging indicator by design. In a slow-moving market, that works. But when deployment timeframes compress from months to weeks, as Zip describes, what use is advice that codifies what was true a cycle ago?

The most-cited statistic in transformation is that around 70 percent of digital transformations fall short of their objectives, a figure popularised by a 2020 BCG article [8]. On inspection the number is itself a piece of circulated received wisdom: its methodology is thin, and the same range is attributed variously to BCG, McKinsey and change-management literature without a single consistent study behind it [8]. That is the point in miniature. The sector runs on packaged confidence that is rarely traced to the operators who actually did the work.

The failures may not be spread evenly. If a large share of agentic projects are cancelled on unclear value and weak controls [4], and a majority of self-described agentic vendors are not delivering the thing [4], then many enterprises are not implementing the frontier model at all, but a diluted, late, vendor-packaged version of it. So when a transformation fails, is the idea wrong, or was the model being copied already a generation old, adopted without the operating discipline that made it work?

The implication: talent is the scarce resource

If the frontier is set by operators and the codified playbook is late, the way to stay current is not to buy the playbook. It is to bring in the people who wrote it.

This is where the constraint bites. Capital is available. Tools are a commodity, and increasingly cloud and model agnostic. Engineering capacity is being augmented by the agents themselves. What remains scarce is the small number of leaders who have personally designed, governed and run a workforce of people and agents in a regulated business. The UK's skills analysis frames the broader shortage starkly, with a large portion of the skilled financial-services workforce at risk of leaving the sector by 2035 [6]. The specific shortage, of people who can lead an agentic operating model, is tighter still.

Those people are not spread evenly across the market. They are concentrated where the work has been done, in fintechs like Send, Zip and Lendi. For a bank or a scaling fintech, the practical consequence is that the most important hire is no longer the person who can run the existing model efficiently. It is the leader who has already built an AI-native business and can build another.

Conclusion

So far this paper has described AI adding capacity. The louder story over the past year has been the reverse: companies using AI to cut jobs. In an earlier analysis this firm set out how listed companies restructuring around AI, among them Block, WiseTech and Commonwealth Bank, were cutting roles while their share prices held or rose, and argued that the decisive variable in what follows is the quality of the leader at the top [11]. These are two halves of one shift. A cut on its own is cost-cutting; a cut paired with an intelligence layer that does what the management chain used to do is reinvention. Whether a business is removing roles or adding agents, the operating model is being rewritten, and the people who can lead the result are scarce.

The model for financial services is being rewritten, and it is being rewritten by operators. Australian fintechs are ahead of their banking competitors not because they have better tools, but because they treated AI as a reason to redesign the business rather than a way to reduce it. The advisory establishment will codify this eventually, and sell it, but by then the frontier will have moved again. The firms that stay current will be the ones that go to the source. On current evidence, identifying and hiring the operators who have already done it is becoming the most valuable search in fintech.

Dexter Cousins is the founder of Tier One People, Australia's leading executive search firm for fintech, payments and the digital economy. He has completed 200+ executive placements and hosts Fintech Chatter, Australia's leading fintech podcast, now past 380 episodes.

Sources

1. Fortune, "Is AI really killing finance and banking jobs?" December 2025, and coverage of Citi's 2024 analysis and Bloomberg Intelligence projections.
<https://fortune.com/2025/12/21/is-ai-killing-finance-and-banking-jobs-experts-say-wall-street-layoffs-hype-than-takeover/>
2. Accenture, banking and generative AI research, as reported by AI News.
<https://www.artificialintelligence-news.com/news/zopa-ai-automate-banking-threaten-finance-jobs/>
3. Bain & Company, "AI in financial services survey shows productivity gains across the board," 2024.
<https://www.bain.com/insights/ai-in-financial-services-survey-shows-productivity-gains-across-the-board/>
4. Gartner forecasts on agentic AI adoption, project cancellation and "agentwashing," 2024–2025, as compiled in industry analyses. <https://www.raftlabs.com/blog/ai-agents-statistics>
5. McKinsey, State of AI research on agent experimentation versus production, 2025.

6. Financial Services Skills Commission / UK report on AI task automation and sector skills, as reported by Scottish Financial News.
<https://www.scottishfinancialnews.com/articles/ai-threatens-to-automate-half-of-financial-services-tasks-report-warns>
7. Zopa and Juniper Research, generative AI cost savings and jobs-at-risk study, 2025.
<https://www.artificialintelligence-news.com/news/zopa-ai-automate-banking-threaten-finance-jobs/>
8. BCG, "If Failure Is Not an Option, Why Is Success So Rare?" 2020.
<https://www.bcg.com/publications/2020/increasing-odds-of-success-in-digital-transformation>
9. Primary interviews: Matt Barr (Send Payments), Soraya Alali (Zip ANZ) and David Hyman (Lendi / Mono AI), Fintech Chatter, 2026.
10. Jack Dorsey and Roelof Botha, "From Hierarchy to Intelligence," Block, March 2026.
<https://block.xyz/inside/from-hierarchy-to-intelligence>
11. Tier One People, "Who leads what's left: AI restructuring," February 2026.
<https://tieronepeople.com/2026/02/28/ai-restructuring/>