

P A R I · W H I T E P A P E R · S E P T E M B E R 2 0 2 6

The Speed Trap

Why “launch fast and iterate” is frontier advice, and what replaces it in a market with a graveyard.

Yoshi Mua

Founder & CEO, Pari

Co-authored & edited with Claude

34 min read

8 parts + conclusion

Published September 2026

Updated September 2026

CONTENTS

Abstract

Part I: The doctrine and where it came from

Part II: Iteration as a search algorithm

Part III: The evidence problem

Part IV: The pivot tax

Part V: Two kinds of unknown

Part VI: What replaces iteration

Part VII: The method has already won, repeatedly

Part VIII: Application to Pari

Conclusion

References

Abstract

The dominant operating doctrine in venture-backed entrepreneurship, which is to ship a minimum viable product, gather feedback, iterate rapidly, and pivot when the data says to, gets treated as a universal law of company building. It isn't one. It's a search algorithm, and like every search algorithm its efficiency depends entirely on the properties of the space being searched.

That doctrine was forged in a specific environment: consumer internet and mobile software between roughly 2005 and 2012, where distribution was nearly free, feedback arrived in hours, mistakes were reversible, no regulator gated entry, and demand itself was unknown. In that environment, rapid iteration is close to optimal. The fastest way to learn what people want is to put things in front of them.

Almost none of those conditions hold in a mature, regulated market littered with the wreckage of prior attempts. There, demand is not the unknown. Mechanism is. And mechanism is not discoverable by iteration, because the iterations have already been run by the companies that died running them. Their post-mortems are the dataset.

This paper makes four arguments.

1. Iteration speed is a conditional tactic rather than a universal best practice, and the conditions that justify it can be enumerated.
2. The startup best-practices canon is survivorship-biased hindsight. It describes winners after the fact and has almost no predictive power going forward. A practice followed by nearly the entire population cannot explain the outcomes of the small fraction that survives.
3. Rapid iteration carries an under-priced liability. Pivots don't just change the product; they silently rewrite the contract under which every early employee joined, and mission drift erodes teams in ways that stay invisible on a metrics dashboard until the institutional knowledge is already gone.
4. In markets like ours, the correct posture is methodical, constraint-first, and sequenced by irreversibility. This is slower in appearance and faster in reality, because it does not repeatedly discard the only asset that compounds.

The final sections apply this to Pari Finance directly, including the criteria under which we would conclude our own approach is wrong.

Part I: The doctrine and where it came from

1.1 What conventional wisdom actually says

Strip away the variations and the modern startup doctrine reduces to a handful of imperatives. Ship something embarrassing, fast. Get it in front of users immediately. Measure everything and let the data decide. Don't fall in love with your idea, fall in love with the problem. Pivot without sentiment when the numbers say to. Speed is the only sustainable advantage a small company has.

This is good advice in the right context. It was a corrective to a real pathology: the multi-year, waterfall-planned, market-research-driven product launch that arrived to an audience that did not exist. The doctrine solved that problem, and solved it well.

Its authority is reinforced by research that appears to validate it. The Startup Genome Project, which analyzed more than 3,000 high-growth technology startups, reported that founders who learn actively raise 7x more money and achieve 3.5x better user growth, and that companies following recognized startup thought leaders were 80% more likely to raise capital.^[6] That is a real finding. It is also, as Part III argues, a finding about a population in which almost everyone follows the same doctrine, which limits what it can tell us about differentiation.

1.2 The environment that produced it

The doctrine did not descend from first principles. It was extracted from a particular set of companies operating under a particular set of conditions.

CONDITION	CONSUMER WEB, 2005 TO 2012
Cost per iteration	Near zero. Deploy on a Tuesday afternoon.
Feedback latency	Hours to days.
Signal fidelity	High. Millions of users, clean behavioral data.
Reversibility of error	Near total. Roll back the commit.
Regulatory gate	None.
Counterparty dependency	None. You needed a server, not a bank.

CONDITION	CONSUMER WEB, 2005 TO 2012
Prior information	Sparse. Nobody had built photo-sharing at scale before.
Nature of the unknown	<i>Does anyone want this?</i>
Cost of public failure	Low. Nobody remembers your failed feature.

Under these conditions, iteration is not merely permitted. It's mandatory, and anything else is malpractice. When each experiment costs nothing and returns clean signal within a week, running a thousand experiments beats thinking hard about which experiment to run.

The error isn't the doctrine. The error is generalizing from that environment to every environment.

Part II: Iteration as a search algorithm

2.1 The reframe

Treat "launch fast and iterate" as what it mechanically is: a local search procedure. You occupy a position in a possibility space, sample nearby positions, move toward improvement, and repeat.

Local search works extremely well under three conditions. Sampling has to be cheap. The signal returned has to be informative. And the landscape has to be smooth enough that local improvement leads somewhere globally good.

It fails badly when any of those break. If sampling is expensive, you can't afford enough samples. If signal is noisy or delayed, you move in the wrong direction with confidence. And if the landscape is rugged, meaning the good answer sits in a different valley separated from you by a region where everything looks worse, local search will climb the nearest small hill and stop there. That hill is where the graveyard is.

2.2 The five variables that determine whether iteration is rational

Cost per iteration. In consumer software, an experiment costs an engineer-week. In a regulated financial product, an experiment can cost a legal review, a partner re-approval, a compliance policy amendment, and six

weeks of calendar time. When iteration costs rise by two orders of magnitude, the number of iterations you can afford drops by two orders of magnitude, and the value of thinking before acting rises to match.

Feedback latency and fidelity. A consumer app learns whether a feature works in 48 hours. An intra-family loan product cannot learn whether its underwriting assumptions were sound until loans season, which takes quarters. When feedback is slow, "let the data decide" stops being a strategy, because the data arrives after the decision window has closed. You are forced to reason from structure instead of from measurement.

Reversibility. Some decisions are doors you can walk back through and some are not. Choosing a bank partner, a compliance posture, a regulatory classification, a distribution channel, or a data-handling architecture are all one-way doors in financial infrastructure. Applying a "we'll fix it next sprint" mentality to a one-way door isn't agility. It's an unfunded liability.

Density of prior information. This is the variable the doctrine ignores most completely. In a novel market the map is blank, and the only way to fill it in is to walk. In a mature market with two decades of failed attempts, the map is already substantially drawn, in the handwriting of the dead. Continuing to walk it yourself, at speed, while ignoring the markings, isn't empiricism. It's redundant empiricism at your own expense.

Cost of public failure. In consumer software, a failed launch costs nothing but attention. In a market where your product is trust infrastructure between family members and your distribution runs through fiduciaries, a visible failure isn't a data point. It's a permanent write-down of the only asset you have.

How favourable each variable is to rapid iteration

scored 1 to 10; 10 means the variable actively rewards shipping fast

ILLUSTRATIVE · NOT MEASURED DATA



Figure 1. The five variables, scored by market type. Illustrative. The scores are the author's qualitative assessment and are offered as a framework rather than as measured data. The doctrine's home environment scores high on all five; ours scores low on all five.

2.3 The two regimes

	FRONTIER MARKET	MATURE MARKET WITH A GRAVEYARD
The central unknown	Does demand exist?	Why has proven demand never been served profitably?
Best instrument	Experimentation	Structural analysis and sequenced execution
Source of advantage	Speed of learning	Correctness of the initial structural insight
What kills you	Building something nobody wants	Building the thing everyone before you built
Right pace	Maximum	Deliberate
Cost of a wrong turn	A sprint	A company
Prior art	Little	Extensive, and free

The claim here isn't that methodical beats fast. It's that the market chooses the pace, and the founder's job is to correctly identify which market they're in. Applying frontier tactics to a mature market belongs to the same

category of error as applying enterprise waterfall planning to a consumer app. Both are the right tool in the wrong regime.

Which pace the market chooses

the doctrine is optimal in one quadrant of four

ILLUSTRATIVE · NOT MEASURED DATA




Figure 2. Which pace the market chooses. Conceptual. The doctrine is optimal in one quadrant of four and actively harmful in the quadrant Pari occupies. The example companies are placed illustratively.

Part III: The evidence problem

3.1 The base rate that nobody sits with

The most-repeated statistic in entrepreneurship is that roughly 90% of startups fail.^[2] It gets repeated so often that it has stopped functioning as information. It deserves to be treated as the central fact of the discipline.

The supporting figures are consistent across sources. U.S. Bureau of Labor Statistics data on all new businesses shows approximately 20.4% failing in year one, 49.4% by year five, and 65.3% by year ten.^[1] For venture-scale technology startups the numbers are worse. Roughly 90% fail over their lifetime, and tech startups close within five years at 63% against the all-industry rate of 49.4% over the same window.^{[1][2]} Harvard Business School research by Shikhar Ghosh found that about 75% of venture-backed startups never return capital to investors.^[4] CB Insights, analyzing hundreds of startup post-mortems, identifies no market need (42%) and running out of cash (29%) as the leading causes, together accounting for 71% of shutdowns, with wrong team appearing in roughly 23%.^[3]




Figure 3. Business survival, all firms versus venture-scale startups. Roughly nine in ten venture-scale startups fail over their lifetime, and the entire population is executing the same playbook. Dashed segments interpolate between reported points rather than reporting them. Sources:^{[1][2]}.

Why startups shut down

share of 483 post-mortems citing each cause




Figure 4. Why startups shut down. Causes overlap, so the figures do not sum to 100. "Ran out of cash" is reported at 29% in the cut used here and at 41% in later vintages of the same dataset. "Wrong team" is the category Part IV argues is most often mislabelled. Source:^[3].

Now hold that against a second fact. Essentially the entire venture-backed population has, for fifteen years, been executing the same doctrine. Lean methodology, MVPs, rapid iteration, data-driven pivots, and growth-hacking are not a differentiated strategy. They are the ambient default. Nearly everyone in the graveyard did all of it, competently, with good mentors and clean dashboards.

A practice that is universally adopted cannot be the explanation for a 10% success rate. If 100% of the field runs the playbook and 90% of the field dies, the playbook is not what separates the outcomes. At best it's table stakes. At worst it's a shared behavior producing some of the failures while being credited with all of the successes.

3.2 The hindsight machine

There's a deeper problem with the canon, and it's methodological.

The business-advice genre works backward from outcomes. It selects companies that succeeded, examines what they did, and reports the commonalities as principles. Phil Rosenzweig's *The Halo Effect* dismantles this genre directly, targeting *In Search of Excellence*, *Built to Last*, and *Good to Great* by name.^[7]^[8] Three of his objections apply with full force to startup doctrine.

The halo effect itself comes first. When a company is winning, observers attribute to it a brilliant strategy, a visionary leader, a superb culture, and disciplined execution. When it falters, the same observers describe the

same behaviors as arrogance, drift, and complacency. Very little may have actually changed. Performance contaminates the perception of every input.^[7]

Second is what Rosenzweig calls the delusion of rigorous research. Studies that sample on the dependent variable, selecting winners and then hunting for shared traits, cannot establish that those traits caused the winning. Rosenzweig's specific objection to *In Search of Excellence*, which analyzed 43 companies using archival data, press accounts, and interviews, is that the sample was constructed on performance itself and then contaminated further by the halo effect.^{[7][8]} The losers who shared the same traits are structurally invisible.

Third, and most damaging to any playbook, is the delusion of absolute performance. Company performance is relative rather than absolute. Success comes from doing things differently and better than rivals, which means managers have to take risks.^[7] It follows that a formula, once widely adopted, cannot deliver above-market results by construction. Everyone executing the same playbook competes on execution quality alone, in a field where the median outcome is death.

Rosenzweig's summary judgment on the genre is that it operates at the level of storytelling while wearing the costume of science.^[7] Startup best practices are the same genre with a shorter feedback loop and a larger audience.

3.3 There is no formula, and the founder data says so

If a formula existed, the paths to outlier outcomes would rhyme. They don't.

Studies of unicorn founders keep failing to produce a usable template. A study of unicorn founders found that most fit no traditional archetype, that 62% had at least one first- or second-generation immigrant founder, and that the funds backing them were extraordinarily fragmented: outside of Y Combinator (10%) and SV Angel (6.4%), no single fund appeared in more than 2.8% of unicorns, which was Sequoia's share.^[9] The researcher's own reading of that fragmentation is that the earliest-stage market for unicorn investment is effectively wide open, meaning even the most celebrated funds miss the large majority of them.^[9]

Academic work on founding-team heterogeneity and time-to-unicorn status has produced results that contradict prior findings and each other. Variables the literature predicted would matter came back statistically insignificant, larger founding teams correlated with longer paths to scale (the opposite of conventional wisdom), and the authors describe their own results as challenging assumptions about the universally positive effects of founding-team heterogeneity.^[10] Separate work examining 171 unicorn founders across 159 measured traits found meaningful differences on only 18 of them, meaning 88.7% of measured characteristics were effectively identical regardless of founder background.^[11]

The paths themselves are irreconcilable. Some companies launched in weeks and some spent years in stealth. Some pivoted repeatedly and some never deviated. Some were founded by domain experts and some by outsiders who didn't know the rules. Some raised enormous rounds early and some stayed lean for years. Team compositions, founder ages, geographies, distribution motions, and capital strategies show no convergence.

Which leaves an uncomfortable conclusion. Best practices are available only in hindsight. They are a narrative constructed after the outcome is known, and they are not a map forward.

None of this means everything is luck and nothing can be reasoned about. It means the correct posture is not "follow the playbook faster than the next company." It's to understand the structure of your own market, identify the constraints that killed the specific companies before you, and design against them. That's analysis rather than imitation, and analysis has a pace of its own.

Part IV: The pivot tax

The doctrine treats pivoting as a costless product decision. It's an organizational one, and it's the most under-priced liability in early-stage company building.

4.1 What the data actually supports

The evidence here supports a narrower claim than the doctrine does, and the difference matters.

The Startup Genome research, which is the most-cited empirical support for pivoting, found that startups pivoting once or twice raise 2.5x more money, show 3.6x better user growth, and are 52% less likely to scale prematurely than startups that pivot more than twice **or not at all.** ^[5]^[6]

The finding is not "pivot freely." The finding is that an optimum exists, that it's small, and that both extremes underperform it. Rigidity and serial pivoting are both punished. The same body of research identifies premature scaling as the most common reason startups underperform, with companies losing the battle early by getting ahead of themselves. ^[5]

The doctrine took a finding about bounded, deliberate course correction and converted it into a cultural license for perpetual repositioning. Those are different things.

Relative performance by number of pivots

multiples reported by Startup Genome, against a shared baseline of 1.0




Figure 5. The pivot optimum is small. Startup Genome reports the 2.5x and 3.6x advantages for one-to-two-pivot companies against a single pooled comparison group of companies that pivot more than twice *or* not at all. The outer bars are that shared baseline, not two independently measured populations. Sources:^[5]^[6].

4.2 The contract nobody wrote down

Founders don't recruit early employees with compensation. They can't, because the compensation is worse than the alternative. They recruit with mission.

An early employee at a pre-seed company accepts below-market cash, illiquid equity with a high probability of being worth nothing, longer hours, and career risk. What clears the market is belief: that this specific problem is worth solving, that solving it matters, and that they'll be one of the small number of people who did it.

That's a real contract. It just isn't on paper.

When a company pivots, meaning it moves away from the problem rather than toward a better mechanism for the same problem, that contract gets rewritten unilaterally. The engineer who joined to fix how families transfer wealth is now working on a small-business invoicing tool. Nothing in their employment agreement was violated. Everything they actually signed up for was.

4.3 The erosion sequence

The damage isn't immediate, which is why it gets underestimated so consistently. It runs on a delay.

Month zero. The pivot is announced with conviction and framed as learning. The reasoning is usually sound. The team nods.

Months one through three. Discretionary effort quietly declines. The unpaid overtime that early-stage companies run on is a function of belief, and belief just took a write-down. Nobody reports this and it appears in no metric.

Months three through six. The people with the most options leave first, because they always do. The strongest engineer takes a call they wouldn't have taken a year ago. Attrition gets attributed to fit or compensation.

Months six through twelve. Domain knowledge is now gone. The person who understood why the third architectural decision was made no longer works here. The company begins re-deriving conclusions it already reached, at full cost.

Month twelve onward. A second pivot is now more likely rather than less. The remaining team has learned that direction is provisional, so it stops investing in any direction deeply. The company acquires an internal culture of tentativeness, which reads externally as a lack of conviction to candidates, partners, and investors alike.

Terminal. The company appears in a post-mortem database under "wrong team," a category showing up in roughly 23% of failures.^[3] The label is usually wrong. The team wasn't wrong. The team was correct for a mission that no longer exists.

The pivot tax, eighteen months out

indexed to 100 at the moment of the pivot

ILLUSTRATIVE · NOT MEASURED DATA




Figure 6. The pivot tax, eighteen months out. Illustrative, not measured. The argument is about the shape and the lag: the organizational cost of a pivot is incurred immediately and observed much later.

4.4 Why the tax is higher in a regulated market

In a frontier market, a pivot moves you to an adjacent square and most of what your team knows still applies. Consumer product intuition transfers. Growth mechanics transfer.

In a regulated financial market, the assets are non-transferable and slow to accumulate.

A compliance program is written against a specific product and a specific risk profile. Change the product and you rewrite the program. A bank or money-movement partner has underwritten you, doing that thing, with that risk; change materially and the relationship reopens from the beginning, or doesn't reopen. A distribution channel built on fiduciary trust, with advisors putting their client relationships behind you, is earned over months and destroyed in a single repositioning. Regulatory posture, audit history, and operating history are literally accumulations of elapsed time. They can't be bought, accelerated, or restarted.

Every one of these compounds with continuity and resets with discontinuity. In a market where the binding constraints are trust and time, a pivot doesn't cost you a quarter of product work. It costs you the compounding.

Which produces the central asymmetry of this paper. In a frontier market, iteration compounds and continuity is cheap. In a mature regulated market, continuity compounds and iteration is expensive. Running frontier tactics in a mature market means paying a premium to destroy your only appreciating asset.

Part V: Two kinds of unknown

Everything above converges on a distinction the doctrine never makes.

The first kind of unknown is whether demand exists. That's an empirical question about human preference and it has no theoretical answer. Nobody could have reasoned their way to short-form video or ride-sharing. You have to build the thing and observe. Iteration is the correct instrument and the only one.

The second kind is why proven demand has never been served profitably. That's a structural question. Demand isn't in doubt, since it's centuries old and documented. What's unknown is the mechanism: the specific configuration of unit economics, regulatory posture, distribution, trust, and incentive alignment under which the thing can actually be delivered at scale without dying.

Iteration is a weak instrument for the second kind, for a reason that's easy to state and hard to internalize. The iterations have already been run. A dozen well-funded, well-staffed, well-advised companies have each already spent a decade running experiments in your space. They shipped, measured, iterated, and pivoted. They died anyway. Repeating their search at higher speed produces no new information. It's the same experiment with a different founder.

New information comes from the structural insight into why the search failed, meaning the constraint all of them hit and none of them cleared. That insight comes from analysis of the failures, not from another lap around the same track.

There's a related trap worth naming. Local search finds local optima. If the viable configuration sits in a different region of the space, incremental improvement from where the industry currently stands will never reach it. It will reach a slightly better version of the thing that already doesn't work. Much of the graveyard consists of well-executed local optima.

The graveyard, read properly, is a controlled experiment with a sample size in the dozens and a ten-year runtime, already paid for by other people's investors. Refusing to read the results because "we learn by shipping" isn't empiricism. It's a preference for expensive data over free data.

Part VI: What replaces iteration

Rejecting the doctrine is only useful if something disciplined replaces it. "Be methodical" isn't a method, and slowness without structure is just slowness. What follows is the operating framework this paper proposes.

6.1 Post-mortem archaeology

Before writing product specification, build a cause-of-death taxonomy for every meaningful prior attempt in the category. Not a competitive landscape slide. A forensic document covering what each company built, who funded it, what its unit economics were, what its distribution motion was, what its regulatory posture was, when it died, and the hard part: the proximate cause of death versus the structural one.

Proximate causes are what the press release said. Structural causes are what actually killed them, and they repeat across companies. When the same structural cause appears in six of eight failures, you've found the constraint that defines your market.

The output is a set of design requirements that are non-negotiable because they were derived from evidence rather than opinion.

6.2 Constraint enumeration before feature work

Enumerate what must be true for the business to exist at all, before deciding what it should look like:

- The regulatory perimeter and how you stay inside it
- The counterparty requirements, meaning what a bank, custodian, or processor must see before saying yes
- The unit economics floor, meaning the acquisition cost above which the model cannot work given frequency and lifetime value

- The trust threshold, meaning what a fiduciary must believe before putting a client relationship at risk
- The operational requirements that only accrue with elapsed time

Product design then happens inside this box. This inverts the normal order, in which teams build first and discover the box when they hit its walls, typically eighteen months in and with a burned runway.

6.3 Sequencing by irreversibility

The default sequence is to do what's fast and visible first. Build the app, get screenshots, show traction, and defer the slow institutional work because it's unglamorous and hard to demo.

That's backward in a regulated market. The long-lead, one-way-door items should be started first, precisely because they're slow:

- Compliance and risk programs before product surface area
- Partner and counterparty conversations opened early, since their clock runs independently of yours
- Legal and tax structure settled before it becomes expensive to change
- Distribution commitments secured before build-out is complete

Sequencing by irreversibility looks slow at month three and is dramatically faster at month twenty-four, because the critical path was walked in parallel instead of discovered serially.

Two sequences, the same twenty-four months

the conventional order looks faster for two quarters and finishes later

ILLUSTRATIVE · NOT MEASURED DATA

Conventional sequence



Sequenced by irreversibility



Figure 7. Two sequences, the same twenty-four months. Illustrative schematic rather than a project plan. The conventional order appears faster for the first two quarters and is materially slower to a live, underwritten pilot.

6.4 Bounded experimentation inside a fixed thesis

This is the distinction that makes the whole framework coherent, and methodical does not mean rigid.

Fix the thesis. Iterate the mechanism.

The thesis covers who we serve, what problem we solve, and why it matters. It's fixed and defended. It's the thing employees signed up for and the thing partners underwrote. It changes only under falsification, never under discomfort.

The mechanism covers interface, onboarding flow, pricing presentation, sales choreography, copy, and sequencing. Those get iterated aggressively and continuously, because they're reversible and cheap.

The operating rule is to move fast where reversible and slow where not. Most companies invert this. They iterate slowly on interface details and fast on identity.

6.5 Proof of demand before proof of product

In a market where demand is the unknown, you have to build to discover whether anyone wants it. Where demand is already established, you can and should obtain commercial commitment before full build-out, in the form of letters of intent, channel commitments, and design partners with signatures.

This isn't a shortcut. It's a harder standard than an MVP with a signup form, because it requires someone with institutional accountability to put their name on a document. It also directly de-risks the constraint most likely to kill you in a mature market, which is distribution rather than product.

6.6 Explicit falsification criteria

This is the safeguard separating a methodical strategy from a rationalized one, and it isn't optional.

The failure mode of the methodical approach is real. Deliberation becomes avoidance. "We're being careful" becomes a story told to defer contact with reality. A company can spend three years perfecting a thing nobody will buy while congratulating itself on rigor. The doctrine of speed exists partly to prevent that, and the concern is legitimate.

The answer is to pre-commit, in writing, to the conditions under which the thesis is considered falsified. Dated, specific, measurable, and defined before the evidence arrives. If the criteria are hit, the conclusion follows without renegotiation. If they aren't, deliberate pace is a strategy rather than an excuse.

A methodical company that can't articulate what would prove it wrong isn't being methodical. It's being slow and calling it strategy.

Part VII: The method has already won, repeatedly

The doctrine's own hall of fame contains counterexamples to itself. Five are examined here. Two are the technical and institutional cases (Figma and Stripe). One is the direct precedent in our own category (CircleLending). One is the closest structural analogue to Pari's regulatory position (Nubank). And one is a rare

natural experiment in which the same founder ran both strategies in the same market and got opposite results (Zenefits and Rippling).

Months from founding to public launch

every company here entered a mature, defended market; none shipped in weeks




Figure 8. Time from founding to public launch. Every company in this chart entered a mature, defended market. None of them shipped in weeks. Sources:^[12]^[13]^[14]^[21]^[23]^[29].

7.1 Figma: four years before the public saw anything

Figma was incorporated on August 16, 2012 and did not launch publicly until September 2016, spending roughly three to four years in stealth before an invite-only beta in December 2015.^[12]^[13] It entered a market defined by Adobe, as mature and defended a software category as existed at the time.

Received wisdom says this should have been fatal. It wasn't, for reasons that map onto the framework above. The binding constraint was technical: building a browser-native vector engine capable of real-time manipulation at 60 frames per second, which required low-level systems programming and careful performance engineering rather than a faster release cycle.^[13] The trust threshold was also extreme. Designers spend eight hours a day inside their primary tool, which makes switching costs enormous and makes a mediocre first impression close to unrecoverable. Shipping an embarrassing MVP into that market would not have generated learning. It would have generated a permanent verdict.

Two complications belong in any fair account of the case. The team did explore early, including a drone project and a meme generator, before converging on the design tool; exploration and public shipping are different activities, and that exploration happened before the thesis was fixed.^[12] And the period was painful.

Employees quit before the product shipped, multiple investors passed, and one investor told Field directly that he did not know what he was doing.^[12] Deliberate pace is not comfortable. It's only correct.

The market validated the patience. Adobe launched XD two months before Figma's beta announcement but did not ship online collaboration until 2019, roughly four years later.^[13] Figma was eventually the subject of a \$20 billion acquisition agreement from Adobe, blocked by regulators in December 2023.^[12]

7.2 Stripe: the most mature market imaginable

Payments in 2010 was not an unexplored frontier. It was an ancient, crowded, heavily intermediated market with entrenched incumbents including PayPal and Authorize.Net, plus the traditional processor stack.^[14] By the doctrine's logic, entering it as two young founders was hopeless.

Stripe was founded in Palo Alto in 2010 and launched publicly on September 29, 2011, after an extensive private beta.^[14] The instructive part is what the constraint turned out to be. Stripe's binding problem wasn't product-market fit and it wasn't developer experience, which the founders understood well. It was institutional: securing a bank partnership.

Banks were not evaluating the elegance of the API. They were asking who reviews merchants, who bears fraud losses, how funds settle, who handles disputes, and whether a company of a few people could discharge those obligations over the long term.^[15] Wells Fargo initially declined outright. What changed the answer was a hire who could understand what the bank was actually worried about and translate those concerns into processes and commitments Stripe could execute against; the partnership came together roughly two months after that person joined.^[15]

No amount of iteration speed produces that outcome. Institutional credibility accrues on a calendar rather than a sprint cadence, and it is bought with organizational capability rather than product velocity.

7.3 CircleLending, Virgin Money USA, and National Family Mortgage

This case sits in our own lineage and it's the most instructive one in this paper.

CircleLending was founded in Cambridge, Massachusetts in May 2000 by Asheesh Advani, with a proposition nearly identical to ours: formalizing and servicing loans between family members and friends.^[17] It was explicitly differentiated from the stranger-lending platforms of the era, including Prosper, Lending Club, Zopa, and Kiva, by focusing on parties who already had an existing relationship.^{[17][19]} By 2006 it had raised roughly \$10 million from Venrock, Bezos Expeditions, and Omidyar Network. In 2007, Richard Branson's Virgin Group


acquired a majority stake and rebranded it Virgin Money USA, intending it as Virgin's entry into U.S. consumer financial services, with plans to expand beyond family loans into direct mortgages and student lending by blending peer-to-peer funds with institutional capital.^[19]

In June 2008 it acquired a wholesale mortgage lender, Lendia, and sold it back to its founder the following year.^{[16][17]} Advani departed in 2009.^[16] In January 2010 the company was still telling reporters it was exploring partnerships with a large U.S. bank and expected to name a new CEO within weeks; the announcement never came.^[16] Virgin Money US ceased operations in November 2010 and entered dissolution on November 1, with servicing of its existing loans transferred to a third party.^{[17][18]} The exit was never officially announced. The website simply went dark.^{[16][18]}

The category's most important lesson is in what happened next. Timothy Burke, a former CircleLending employee, founded National Family Mortgage in 2010 to fill the vacuum left by Virgin's exit.^[20] It was narrow in scope, focused specifically on intra-family real estate loans, and built without venture-scale capital or venture-scale expectations. It has operated continuously ever since and reports having supported the documentation of billions of dollars in mortgage loans between family members, in every state.^[20]

Two paths out of the same insight

same category, same core idea, opposite pace, opposite outcome



The venture-paced version lasted three years past acquisition. The deliberate version has run for sixteen.

Figure 9. Two paths out of the same insight. Same category, same core insight, opposite pace, opposite outcome. The venture-paced version lasted three years past acquisition; the deliberate version has run for sixteen. Sources: [16] [17] [18] [19] [20].

Set the two outcomes against each other. The well-capitalized, brand-backed, rapidly expanding version, which added a mortgage lender, planned adjacent product lines, and moved at the pace its capital structure demanded, was dissolved within three years of the acquisition. The narrow, deliberate, structurally conservative version, same category and same core insight, run by one former employee, is still operating sixteen years later.

The conventional reading is bad timing, and the financial crisis was undeniably a factor; contemporary coverage attributed the exit to the crisis damaging both peer-to-peer and mortgage lending, and noted that regulatory scrutiny also pressured the model. [16] [18] But timing is a proximate cause. The structural causes are legible: consumer acquisition economics for a low-frequency, high-consideration product; expansion into balance-sheet-adjacent businesses before the core was durable; regulatory exposure inherited from the surrounding peer-to-peer category; and dependence on a corporate parent whose strategic priorities were

exogenous. None of those were problems that faster iteration would have solved. Several were problems that faster movement created.

This is the graveyard doing its job. It tells us, at no cost to us, which configurations don't survive contact with this market.

7.4 Nubank: a regulatory permit that took years, in a market five banks owned

Nubank is the closest structural analogue to Pari's position, and it is now one of the largest financial institutions in the world.

David Vélez founded Nubank in São Paulo on May 6, 2013, after concluding that Brazilian retail banking was an anomaly rather than an inconvenience.^[23] The market he entered was not merely mature. It was an oligopoly: five institutions controlled roughly 80% of financial assets, and credit card interest rates ran in the hundreds of percent annually.^[22] The industry consensus was that entry was impossible. Vélez has recounted being told directly that nobody competes with the five banks that own the market, that they would crush him, and that as a foreigner he did not understand Brazil.^[21]

The relevant fact for this paper is what he did about the regulatory gate rather than how fast he shipped. As a foreign national, Vélez needed a presidential decree to obtain a permit to operate in Brazil, a process he understood in advance would take years, and which did.^[21] He built the company anyway, in the correct sequence.

The consequences are visible in the product timeline. Nubank launched a single product in 2014: a no-annual-fee Mastercard credit card managed entirely through an app, released invitation-only.^{[21][23]} The digital account product, NuConta, did not arrive until 2018, and only after Nubank finally received its special banking permit from the Brazilian government.^[21] Five years elapsed between founding and the banking permit that unlocked the second product.

Three elements of that sequence map directly onto Part VI. The company obtained proof of demand before proof of product: the invitation-only waitlist reached into the hundreds of thousands before national launch, driven by word of mouth rather than marketing spend.^[22] It ran a deliberately narrow initial risk surface, adopting a "low-and-grow" credit strategy of issuing small limits and expanding them only as behavioral data proved creditworthiness. And it sequenced by irreversibility, treating the multi-year regulatory permit as a workstream started early and run in parallel rather than a blocker discovered late.

The result is not a consolation prize for patience. Nu Holdings went public on the New York Stock Exchange in December 2021.^[22] By 2026 the company reported more than 135 million customers across Brazil, Mexico, and Colombia, annualized revenue of roughly \$20 billion, and conditional approval to operate as a national bank in the United States.^[23]

7.5 Zenefits and Rippling: the same founder, both strategies, one market

This is the closest thing the sector has to a controlled experiment. The variable that changed is pace and posture. The founder, the market, and the core commercial insight held constant.

The fast run. Parker Conrad founded Zenefits in 2013, offering free HR software to small businesses and monetizing as their health insurance broker. It became, by its own account, the fastest-growing SaaS company ever, reaching roughly \$50 million in ARR in under two years and a \$4.5 billion valuation on a \$500 million round led by Fidelity and TPG.^{[27][29]} It was the doctrine executed flawlessly.

The compliance layer did not keep up with the growth rate. BuzzFeed News reported in November 2015 that the company had allowed unlicensed employees to sell insurance; a follow-up found that 83% of Zenefits' insurance deals in Washington State through August 2015 had been done by employees without licenses.^[25] Conrad had also created and distributed an internal software macro that let employees circumvent California's pre-licensing education requirements.^{[26][28]}

The unwinding was fast and total. Conrad resigned in February 2016 and was removed from the board.^{[24][25]} Incoming CEO David Sacks told employees that the company's internal processes, controls, and actions around compliance had been inadequate, and framed the stakes in a line worth quoting exactly: "For us, compliance is like oxygen. Without it, we die."^[31] Fidelity marked its investment down 48%, and the company later cut its own valuation by more than half, to roughly \$2 billion, in a settlement with investors who agreed not to sue.^[24]^[28] The California Department of Insurance fined Zenefits \$7 million, among the largest licensing fines in that department's history, and Conrad surrendered his insurance license in 2018.^[26] The SEC charged both the company and Conrad with failing to adequately disclose the compliance lapses while raising capital in 2014 and 2015; Zenefits paid \$450,000 and Conrad personally paid just over \$533,000.^{[27][28]}

The methodical run. Conrad founded Rippling in 2016 to attack essentially the same market with the opposite operating posture.

The most telling number is the pre-revenue one. Rippling spent approximately \$10 million and 18 months building before generating revenue, drawing on the founders' Zenefits experience to map a multi-year product roadmap from day one.^[29] In a category where the received wisdom is to ship a narrow MVP in weeks, that is a deliberate and expensive inversion.

Conrad also rejected the focus doctrine explicitly, coining the term "compound startup" to describe building many deeply integrated products in parallel off a shared platform rather than doing one thing well and expanding later.^[30] And he has been direct about the compliance lesson, stating publicly that Zenefits had a series of compliance failures they are careful not to repeat at Rippling.^[27]

The outcome inverts the first run. Rippling raised a \$45 million Series A at a \$270 million valuation in 2019, crossed \$1 billion in ARR in 2025, and reached a \$16.8 billion valuation in a \$450 million Series G in May 2025, with more than ten product lines each generating over \$1 million in ARR.^[29]^[30]

Same founder, same market, two operating postures

valuation, logarithmic scale; the two series are separate companies on a shared axis




Figure 10. Same founder, same market, two operating postures. Logarithmic scale. The two series are separate companies on a shared axis and are deliberately not connected. The variable that changed between the runs was operating posture in a regulated market. Sources:^[24]^[26]^[27]^[28]^[29]^[30].

The lesson is not that Rippling is slow. It plainly isn't; it ships more products in parallel than most companies ship sequentially. The lesson is where the speed is applied. Zenefits ran fast through the regulatory perimeter and died on contact with it. Rippling spent 18 months and eight figures building the platform underneath before selling anything, then ran fast on top of it. Fast where reversible, slow where not, executed by the same person who learned the distinction the expensive way.

Part VIII: Application to Pari

8.1 We are in the second regime

Intra-family lending is not a novel behavior awaiting validation. Families have lent to one another for as long as families and money have coexisted. Demand is not a hypothesis. The Bank of Mom and Dad operates at enormous scale today, informally, undocumented, and taxed incorrectly.

The category also has a graveyard. CircleLending and Virgin Money USA are the direct precedent. The broader peer-to-peer lending cohort supplies an adjacent set of structural failures we've examined in detail elsewhere. These attempts weren't under-resourced or badly run. They were backed by Venrock, Bezos Expeditions, Omidyar Network, and the Virgin Group, and they failed structurally.

Both diagnostic markers of the second regime are present: proven demand, and a documented record of well-funded failure to serve it. Applying frontier tactics here would be a category error.

8.2 Our binding constraints don't yield to iteration speed

Enumerated honestly, the things that determine whether Pari exists in three years are these.

Money-movement partner risk acceptance. Partners evaluate operating history, compliance maturity, and capitalization. Those are time-accrued and evidence-based, and no release cadence accelerates them. This is Stripe's constraint, not a version of it.

BSA/AML program maturity. A board-approvable compliance program, converted into an actual engineering specification with traceable requirements, is a document read by people whose job is skepticism.

Security and audit posture. SOC 2 is a calendar-bound process with an independent auditor. It cannot be sprinted.

Tax and legal correctness. AFR-compliant promissory note structure is a domain where "we'll iterate on it" isn't an available option, because the counterparty is the IRS and the affected party is a family.

Fiduciary channel trust. An RIA putting client households in front of us is staking a relationship they spent years building. That isn't a conversion-rate optimization problem.

Every one of these is slow by nature and irreversible when wrong. Section 6.3 says to start them first, and we did.

8.3 The chicken-and-egg is a trust problem, not a speed problem

Our central strategic obstacle is well understood. Partners want operating history and compliance maturity before activating rails. Investors want traction before committing capital. Each waits on the other.

The doctrine's instinct is to break the deadlock with velocity: ship something, get numbers, force the issue. In this market that instinct is wrong, because both parties are gating on trust, and trust doesn't respond to speed. It responds to accumulated evidence, meaning documentation that survives scrutiny, commitments from credible counterparties, a compliance program authored by someone whose name carries weight, and elapsed time without incident.

So we invert the sequence.

Proof of demand before proof of product. We've accumulated signed letters of intent from RIA firms representing substantial AUM, with an identified pipeline of client households. That's a harder commitment to obtain than app signups, and it's evidence against exactly the constraint that killed the direct-to-consumer attempts before us. Nubank's pre-launch waitlist served the same function in the same regime.

Compliance before surface area. The BSA/AML program exists as a board-approvable document and as a build specification with stable requirement identifiers, ahead of the product surface that will depend on it. Zenefits is the counterfactual, and it is not a hypothetical one.

Distribution designed against the documented cause of death. The graveyard's clearest structural lesson is that direct consumer acquisition economics don't work for a low-frequency, high-consideration financial product. Our answer isn't to acquire consumers more cheaply. It's to not acquire consumers directly at all, and to reach households through the fiduciary who already advises them. That's a structural response to a structural failure, and it's the kind of answer that iterating inside the old configuration would never have produced.

Deliberately narrow initial risk surface. Our pilot scope is confined to onboarding loans that have already been issued, so facilitated money movement is limited to monthly payments against existing balances. This reduces first-time risk exposure across every dimension a partner or regulator cares about. It is the same logic as Nubank's low-and-grow credit strategy: prove behavior on a small surface, then expand the surface. It's a smaller opening move than the doctrine would recommend, and it's chosen on purpose.

8.4 What we iterate on, aggressively

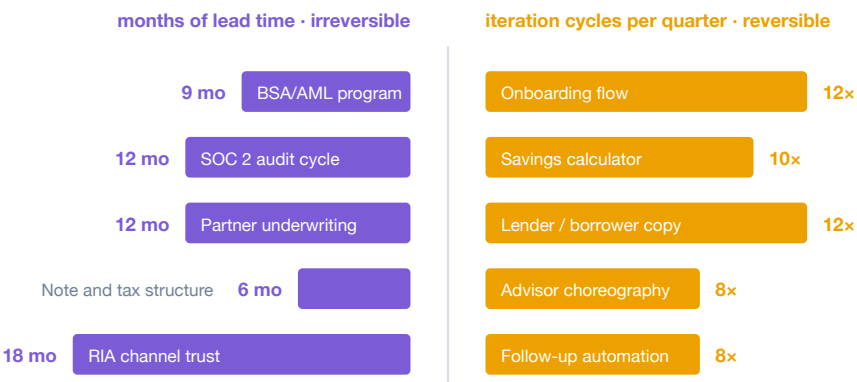
This is where the argument is most often misread, so it's worth being blunt. We aren't a slow-moving company. We're a company that moves fast where speed is free and slowly where speed is expensive.

We iterate continuously on the reversible layer: onboarding flow, the calculation engine that shows a family what a properly structured loan actually saves them over its life, role-specific copy for lenders versus borrowers, advisor-facing sales choreography, prep materials, follow-up automation, and investor communication. All of it is cheap to change, produces fast signal, and carries no structural risk. The right pace there is maximum.

What we don't iterate on is who we serve and why. That's fixed.

Where Pari runs slow and where it runs fast

ILLUSTRATIVE · NOT MEASURED DATA



Fast where mistakes are cheap, slow where they are permanent.

Most companies invert this without noticing.

Representative planning assumptions, not committed schedules.

Figure 11. Where Pari runs fast and where it runs slow. Values are representative of planning assumptions rather than committed schedules. Fast where mistakes are cheap, slow where they are permanent.

8.5 The team consequence

There's a direct organizational return on this, and it's the reason Part IV exists.

The people who joined Pari, including a co-founder and CTO, a co-founder focused on growth, a chief product officer, senior engineering, and a general counsel serving as BSA officer, joined to build infrastructure for intra-

family lending. That's still what we're building. The mission they were recruited into is the mission that exists today.

This isn't sentimentality; it's an operating advantage with measurable components. Domain knowledge in a regulated market takes twelve to eighteen months to build and is largely non-transferable, so retaining it is worth more here than in a category where product intuition ports across pivots. Partner and channel relationships are held by individuals, so continuity of people is continuity of relationship. And discretionary effort, the thing early-stage companies actually run on, is a function of belief that hasn't been repeatedly re-underwritten.

A company that hasn't pivoted out of its mission holds an asset that appears on no dashboard: nobody on the team is quietly deciding whether this is still the thing they signed up for.

8.6 What would prove us wrong

The framework in Part VI is only honest if it's falsifiable, and the failure mode of the methodical posture is deliberation as avoidance. We take that risk seriously. Our commitment is to pre-defined, dated, written criteria across the dimensions that determine whether this thesis holds.

Channel conversion. Signed intent must convert to live, funded activity at a defined rate within a defined window. Letters of intent that never convert are a signal about the channel, not a reason to keep collecting more of them.

Counterparty viability. If no money-movement partner will underwrite the pilot within a defined window despite a mature compliance posture, the constraint is structural rather than temporal, and the architecture has to change instead of the timeline.

Advisor economics. If the advisor-side value proposition doesn't produce repeat placement from firms that have already run one household through, the distribution thesis is wrong regardless of how many firms sign initial agreements.

Household behavior. If families onboarded to formal structure don't sustain payment behavior at rates that make servicing economics work, the core premise requires revision.

Where these criteria are met, deliberate pace was the correct strategy. Where they aren't, we change: bounded, deliberate, once, on evidence, with the thesis re-examined explicitly rather than drifted away from. That's what the pivot research in section 4.1 actually supports, and it's what we hold ourselves to.

Conclusion

The instruction to launch fast and iterate isn't wrong. It's conditional, and its conditions are specific: cheap experiments, fast and clean feedback, reversible mistakes, no regulatory gate, no institutional counterparty, and a blank map. Under those conditions it's close to optimal, and this paper doesn't dispute it.

Those conditions describe a frontier. They don't describe a mature, regulated market with two decades of well-funded failures in it. There, sampling is expensive, feedback is slow, errors are one-way doors, entry is gated by institutions that underwrite trust rather than momentum, and the map is not blank. It has been drawn in detail by companies that ran the fast iterations already and did not survive them.

The best-practices canon offers no help, because it isn't a map forward. It's a story assembled backward from known outcomes, with the losers structurally invisible, in a field where roughly 90% of participants die executing the same universally adopted playbook. A practice that everyone follows can't explain the outcomes of the few who make it.

And the doctrine's central mechanism carries a cost it never prices. Pivoting gets treated as a product decision when it's an organizational one. Founders recruit on mission because mission is the only currency they have, and a pivot voids that contract for every person who accepted it. In a market where the compounding assets are trust, domain knowledge, institutional relationships, and elapsed time, discarding continuity isn't agility. It's the destruction of the only thing that was appreciating.

The counterexamples are not obscure. Figma spent four years in stealth against Adobe. Stripe's binding constraint was a bank's underwriting committee rather than its release cadence. Nubank waited five years for a banking permit while five institutions held 80% of its market, and is now one of the largest financial institutions in the world. National Family Mortgage, built narrow and deliberate by one former employee, has outlived the venture-backed version of the same idea by sixteen years. And Parker Conrad ran both strategies in the same market and got a \$7 million fine the first time and a \$16.8 billion valuation the second.

Pari is building deliberately because the market we entered demands it, and because the companies before us documented what happens when it doesn't. We fix the thesis and iterate the mechanism. We move fast where mistakes are cheap and slowly where they're permanent. We sequence by irreversibility rather than by ease. We obtained commercial commitment before completing build-out. We designed our distribution as a direct structural answer to the documented cause of death in our own category. And we've written down what would prove us wrong.

The graveyard in our market isn't a warning to stay out. It's the most valuable and least-used asset available to anyone entering it: a decade of expensive experiments, already paid for, with the results published. Reading it carefully before moving isn't caution. It's the only rigorous thing to do.

References

FAILURE RATES AND CAUSES

- [1] U.S. Bureau of Labor Statistics, Business Employment Dynamics (2024 data). First-year failure 20.4%, five-year 49.4%, ten-year 65.3%; tech-sector five-year closure 63%. Compiled at <https://www.makerstations.io/startup-failure-rate-statistics/> and <https://www.llc.org/startup-failure-rate-statistics/>. Primary series: <https://www.bls.gov/bdm/>
- [2] Startup Genome. Approximately 90% lifetime failure rate for innovative, venture-scale startups. <https://startupgenome.com/>
- [3] CB Insights, *Startup Failure Post-Mortems* (483 post-mortems). No market need 42%, ran out of cash 29%, wrong team ~23%. <https://www.cbinsights.com/research/startup-failure-reasons-top/>
- [4] Shikhar Ghosh, Harvard Business School. Approximately 75% of venture-backed startups never return capital to investors.

PIVOTING AND PREMATURE SCALING

- [5] Startup Genome, *Why Startups Fail: Premature Scaling* (Report Extra, v2). Pivot findings and premature scaling analysis. https://s3.amazonaws.com/startupcompass-public/StartupGenomeReport2_Why_Startups_Fail_v2.pdf
- [6] "What Makes A Startup Successful? Blackbox Report Aims To Map The Startup Genome," TechCrunch, May 28, 2011. <https://techcrunch.com/2011/05/28/what-makes-a-startup-successful-blackbox-report-aims-to-map-the-startup-genome>

SURVIVORSHIP BIAS AND THE BEST-PRACTICES CANON

- [7] Phil Rosenzweig, *The Halo Effect ... and the Eight Other Business Delusions That Deceive Managers* (Free Press, 2007). Review and summary: "The Halo Effect: Debunking Some Hot Business Books with One of His Own," Knowledge at Wharton, February 2007. <https://knowledge.wharton.upenn.edu/article/the-halo-effect-debunking-some-hot-business-books-with-one-of-his-own/>
- [8] Phil Rosenzweig, "Misunderstanding the Nature of Company Performance: The Halo Effect and Other Business Delusions," *California Management Review* 49(4), 2007, pp. 6-20.

UNICORN FOUNDER STUDIES

- [9] "New study of unicorn founders finds most are 'underdogs,'" TechCrunch, March 27, 2024. <https://techcrunch.com/2024/03/27/unicorn-founders/>
- [10] "The pace of scaling: founding team heterogeneity and the timing of Unicorn status," *Journal of Small Business and Enterprise Development*. <https://www.sciencedirect.com/org/science/article/pii/S1462600425000390>
- [11] "88.7% of unicorn founders share identical traits. Here's why pattern matching fails VCs," Tech Funding News, October 2025. <https://techfundingnews.com/why-pattern-matching-fails-unicorn-traits/>

FIGMA

- [12] "Who is Figma's Dylan Field?" *Fortune*, September 16, 2022. <https://fortune.com/2022/09/16/dylan-field-figma-adobe-acquisition/>
- [13] "Studying Figma's Growth: From \$100K to IPO," Growth Case Studies. <https://growthcasestudies.com/p/figma>

STRIPE

- [14] "Stripe, Inc.," Wikipedia. Founding 2010; public launch September 2011 after extensive private beta. [https://en.wikipedia.org/wiki/Stripe_\(company\)](https://en.wikipedia.org/wiki/Stripe_(company))
- [15] "Stripe 16-Year Chronicle," PANews. Wells Fargo partnership and the institutional-credibility hire. <https://panews.io/articles/01a023a2-294a-76d2-8d2a-f1066aeb119>

CIRCLELENDING, VIRGIN MONEY USA, NATIONAL FAMILY MORTGAGE

- [16] Sara Lepro, "Virgin Money Closes Shop in the U.S., Victim of Bad Timing," *American Banker*, December 1, 2010. <https://www.americanbanker.com/news/virgin-money-closes-shop-in-the-us-victim-of-bad-timing>
- [17] "Virgin Money US," Wikipedia. https://en.wikipedia.org/wiki/Virgin_Money_US
- [18] Peter Renton, "The Demise of Virgin Money US," Lend Academy, November 24, 2010. <https://www.lendacademy.com/the-demise-of-virgin-money-us/>
- [19] "CircleLending Becomes Virgin Money USA," TechCrunch, October 16, 2007. <https://techcrunch.com/2007/10/16/circlelending-becomes-virgin-money-usa-gets-makeover-and-millions-in-funding/>
- [20] "About National Family Mortgage." <https://www.nationalfamilymortgage.com/about-us/>

NUBANK

- [21] "Nubank, a Brazilian startup, is changing the way millions of people bank," CNN Business, December 6, 2019. Presidential decree, permit timeline, NuConta 2018. <https://edition.cnn.com/2019/12/06/business/nubank-david-velez-risk-takers>
- [22] "David Velez & family," Forbes profile. Five banks controlling 80% of the market; NYSE listing December 2021. <https://www.forbes.com/profile/david-velez/>
- [23] "Nubank turns 13 as one of the most influential financial institutions in the world," Nubank, May 19, 2026. Founding date, 135 million customers, ~\$20B annualized revenue, U.S. conditional approval January 2026. <https://international.nubank.com.br/company/nubank-turns-13-as-one-of-the-most-influential-financial-institutions-in-the-world/>

ZENEFITS AND RIPPLING

- [24] "Zenefits CEO steps down over broker compliance problems," HR Dive, February 9, 2016. <https://www.hrdiver.com/news/zenefits-ceo-steps-down-over-broker-compliance-problems/413582/>

- [25] William Alden, "Zenefits Founder Parker Conrad Resigns After Compliance Failures," BuzzFeed News, February 8, 2016. 83% of Washington State deals by unlicensed employees. <https://www.buzzfeednews.com/article/williamalden/zenefits-ceo-parker-conrad-steps-down-after-compliance-failu>
- [26] "Former Zenefits CEO surrenders insurance license following department enforcement action," California Department of Insurance press release, May 9, 2018. \$7 million fine; macro circumventing pre-licensing requirements. <https://www.insurance.ca.gov/0400-news/0100-press-releases/2018/release052-18.cfm>
- [27] Amy Feldman, "The Comeback Of A Fallen Tech Unicorn CEO," *Forbes*, May 28, 2020. SEC settlement amounts; Conrad on not repeating Zenefits' compliance failures at Rippling. <https://www.forbes.com/sites/amyfeldman/2020/05/28/the-comeback-of-a-fallen-tech-unicorn-ceo/>
- [28] "SEC fines Zenefits and former CEO Parker Conrad," Berkeley Law, November 11, 2017. Valuation cut from \$4.5B to \$2B. <https://sites.law.berkeley.edu/thenetwork/2017/11/11/sec-fines-zenefits-and-former-ceo-parker-conrad>
- [29] Rippling company profile, GetLatka. \$10M and 18 months to first revenue; \$45M Series A; funding history. <https://getlatka.com/companies/rippling>
- [30] "Rippling's AI Strategy," Perspective AI, May 2026. Compound startup definition; \$1B ARR; \$16.8B Series G valuation, May 2025. <https://getperspective.ai/blog/rippling-ai-strategy-hr-platform-customer-conversations-product-velocity-2026>
- [31] Kia Kokalitcheva, "Major CEO Shakeup At Zenefits Over Regulatory Problems," *Fortune*, February 8, 2016. Sacks memo to employees. <https://www.fortune.com/2016/02/08/zenefits-parker-conrad-resigns>

Disclosure: This document was drafted with AI assistance and reviewed, edited, and approved by the author.
Content marking is provided in accordance with applicable transparency obligations under the EU AI Act.

Nothing in this paper constitutes an offer to sell or a solicitation of an offer to buy securities, or investment, legal, or tax advice.

Research, thesis, and argument by Yoshi Mua.
Drafting assistance by Claude (Anthropic).

© 2026 Pari Finance Inc. · Written by Yoshi Mua, co-authored and edited with Claude (Anthropic). This paper is for informational purposes only and does not constitute tax, legal, or investment advice.

yoshimua.com · pariapp.com · [Download the PDF](#)