CANTILLION PRESS

POLICY WHITE PAPER · JUNE 2026

PAPER NO. 02 · ENTREPRENEURSHIP ECOSYSTEM POLICY SERIES

FOR GOVERNMENTS WORLDWIDE

AN INTERNATIONAL POLICY WHITE PAPER

The Entrepreneurship Economy

Why every country now needs twice as many companies — and the policy architecture that makes it possible.

PREVIEW CONTENT

From *We're Doing This Wrong: How To Build Entrepreneur Ecosystems That Actually Work*
Available November 2026

DC Cahalane

FOUNDER · INVESTOR
GOVERNMENT ADVISOR

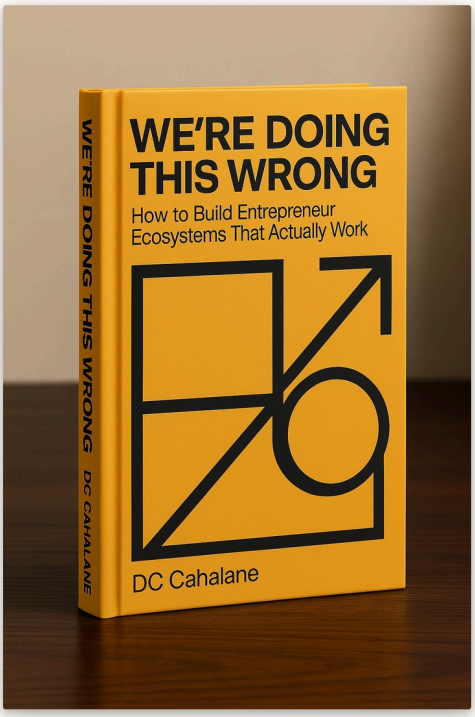

FORTHCOMING · NOVEMBER 2026

BASED ON THE UPCOMING BOOK

We're Doing This Wrong.

How to Build Entrepreneur Ecosystems That Actually Work

This paper distils the policy argument from *We're Doing This Wrong*, the forthcoming book by DC Cahalane, published by Cantillion Press in November 2026.

The book extends this paper into the underlying architecture: six principles of ecosystem design, country case studies, founder testimony, and an implementation manual.

THIS EDITION

The Entrepreneurship Economy · Why every country now needs twice as many companies. Policy white paper, June 2026.

An international policy paper, written to be useful across jurisdictions. The argument is country-neutral by design. It is not aimed at any single government's current portfolio and should not be read as a critique of any specific policy in force; the examples are drawn from a range of OECD countries and emerging ecosystems to illustrate transferable patterns. For political leaders, ministers, senior civil servants, economic-development agencies, university leaders, and public institutions in any country renewing its entrepreneurship economy.

PUBLISHER

Cantillion Press
Cork, Ireland

CONTACT

me@dc.ie

FRONT MATTER · SUMMARY PANEL

At a Glance.

If tomorrow's companies are smaller, countries need more of them. The next jobs strategy is a company-creation strategy.

01

Companies are getting smaller

AI compresses headcount per firm.

+

02

Countries still need the same employment

Labour-market expectations unchanged.

=

03

Twice as many companies

Company-formation rate must rise.

Three forces — productivity gains inside firms, the AI transition, and global concentration risk — are decoupling firm growth from job growth at the same time. The policy response is not to attract more of yesterday's employers. It is to create more of tomorrow's, deliberately and at scale.

> A country that does not create employers will eventually run out of employment policy.

CONTINUED

The six firsts.

A serious entrepreneurship economy fixes the six moments where founders stop. Each is an institutional friction the state can address without picking winners.

FIRST 01 **Permission**	They decide whether to start at all.	FIRST 02 **Customer**	They need someone to pay.
FIRST 03 **Capital**	They need fast, simple early funding.	FIRST 04 **Talent**	They need to build a team.
FIRST 05 **Regulation**	They need a workable path to legality.	FIRST 06 **Recovery**	They need a way back if it does not work.

First twelve months for a government.

A government does not need a five-year plan to start. Each three-month phase has a single deliverable and a named owner.

- **Months 1–3** · Name the strategy and the senior owner.
- **Months 3–6** · Map the founder journey end-to-end.
- **Months 6–9** · Open first-customer pathways and reform spin-out rules.
- **Months 9–12** · Activate local capital and publish the dashboard.

READER'S NOTE

How to use this paper.

A working document for political and policy readers, designed to be read in parts.

This paper is for politicians, ministers, senior civil servants, economic development agencies, regional development bodies, university leaders, and anyone responsible for the economic future of a country. It is not a paper about "startups" as cultural accessory. It is a paper about **entrepreneurship as economic infrastructure.**

SCOPE **International by design**	The argument is country-neutral. The data and case studies are drawn from a range of OECD countries and emerging ecosystems — the United States, Israel, Estonia, the United Kingdom, France, Germany, Sweden, Canada, Chile, the Netherlands, Ireland, and others — to illustrate transferable patterns rather than to grade any single national policy.
TONE **Constructive, not critical**	The paper is not a critique of any government's current portfolio. Every country named in the body has done valuable work; the question this paper asks is what comes next. The argument is forward-looking and intended to be useful to readers building, renewing, or expanding an entrepreneurship economy in their own jurisdiction.

Three reading paths

PATH 01 **Political leaders**	Read the executive summary, Chapters 01–04, and the first twelve months. The political and economic case sits there.
PATH 02 **Senior civil servants & policy designers**	Read Chapters 07–16. The policy architecture, measurement, procurement, capital design, regional renewal, universities, immigration, and implementation sit there.
PATH 03 **Founders, investors, university & ecosystem leaders**	Read the whole paper as a companion to the broader argument of the forthcoming book.

Each chapter opens with a key-takeaways panel. A reader pressed for time can read the panels alone and still brief their staff. The references at the back are auditable. The glossary defines every term used in the body.

TABLE OF CONTENTS

Contents.

FRONT MATTER

Executive summary.

The argument of the paper, in three pages, with the evidence and the policy posture every reader can quote.

KEY TAKEAWAYS

01 Large companies are becoming more productive with fewer people. The relationship between firm growth and job growth has weakened.

02 If tomorrow's companies are half the size, countries need roughly twice as many companies to maintain employment capacity.

03 The policy task is not to pick winners. It is to fix the six moments where founders stop: first permission, customer, capital, talent, regulation, and recovery.

04 The reform agenda is mostly administrative, not legislative. Most of it can be done inside a single political cycle.

Governments are used to thinking about employment in terms of workers, skills, vacancies, employers, sectors, and inward investment. Those categories still matter. But they are incomplete. They assume that the future will contain enough employers to absorb the people who need work, and that assumption is becoming more dangerous.

The World Economic Forum estimates that 170 million roles will be created and 92 million displaced by 2030 across surveyed employers [1]; the IMF estimates that almost 40 per cent of global employment is exposed to artificial intelligence, with advanced economies more exposed because of their concentration of cognitive-intensive roles [2]. The World Bank describes increased trade tensions and policy uncertainty as a substantial headwind to global prospects [4], and the IMF projects global growth slowing through 2026 with risks tilted to the downside [5].

This is a structural shift, not a cycle. Three forces are running together for the first time in the post-war period: productivity-without-headcount inside firms, geopolitical concentration risk in supply chains, and the AI transition in cognitive work. The combined effect is to weaken the link between firm growth and job growth in exactly the firms that used to anchor labour-market absorption.

For most of the late twentieth century the implicit national bargain was: attract large employers, train workers, count jobs. That bargain still functions for some sectors, but it is no longer sufficient at the level of an entire economy. The deficit it leaves — fewer firms creating fewer jobs per unit of revenue — must be made up somewhere. The argument of this paper is that it must be made up domestically, deliberately, and through firm formation.

Countries cannot rely solely on external demand, multinational employment, foreign boardrooms, or global capital cycles. They need domestic engines of renewal. SMEs already represent around 99 per cent of OECD firms and generate 50–60 per cent of value added [6]; about 10–15 per cent of SMEs that scale up account for around 50 per cent of new jobs [7]; growing SMEs create roughly 16 additional jobs for every 10 created by large firms [8]. New and young firms are responsible for most net new job creation [10].

The conclusion is not that every startup matters. Most will not. The argument is that national economies are not built by betting on one perfect company; they are built by raising the rate, quality, diversity, and resilience of company formation. The country that can produce twice as many serious early-stage firms over the next decade — and that can keep their headquarters domestic — will absorb the AI transition. The country that cannot will not.

The political consequence is already visible. The most recent Eurostat data on early-career entrepreneurial activity place the share of 20–29-year-olds engaged in self-employment or new-venture creation at 12.2 per cent in Slovakia and only 5.1 per cent in Ireland, with most other EU member states clustered between 6 and 10 per cent [20]. The countries at the top of that distribution are not the wealthiest. They are the ones whose institutional architecture makes starting a company a normal early-career option rather than an unusual one.

Three structural design choices distinguish the leaders from the rest. First, they treat formation friction — the legal, tax, banking, and regulatory layer that decides how long it takes a founder to legally start a company — as a deliberate national product, owned by a single accountable office. Second, they treat procurement as a demand-side instrument, not a risk-management exercise: the state is willing to be a first customer for credible early companies in defined sectors. Third, they design capital architecture so that public money catalyses private capital rather than substituting for it, and so that exit liquidity recycles into the next generation of founders rather than escaping abroad.

The mindset shift

OLD REFLEX	ENTREPRENEURSHIP-ECONOMY REFLEX
Protect existing jobs	Create future employers
Attract large firms	Build domestic firm density
Fund programmes	Build formation architecture
Count participants	Count companies, employers, exports, repeat founders
Treat founders as grant applicants	Treat founders as economic infrastructure
Use procurement to avoid risk	Use procurement to create strategic first customers

The paper proposes a national policy stack built around **six firsts** — first permission, customer, capital, talent, regulation, and recovery. Each is a moment at which founders most commonly stop. Each requires a named owner inside government, a single measurable target, and a small dedicated team. The aggregate capability is fewer than fifty people, distributed across existing ministries, with cross-cabinet sponsorship from the head of government's office.

The reform agenda is largely administrative. Most of it can be done inside a single political cycle, without primary legislation, without significant new spending, and without picking winners. The state's job is to remove the avoidable obstacles in the founder journey, not to choose between founders. That is a smaller and more defensible role than the standard "industrial policy" frame suggests, and a substantially more effective one than the standard "stand back" frame admits.

This paper is country-neutral by design. The data, case studies, and reform patterns are drawn from a range of OECD countries and emerging ecosystems — the United States, Israel, Estonia, the United Kingdom, France, Germany, Sweden, Canada, Chile, the Netherlands, Ireland, and others — to illustrate transferable architecture rather than to grade any single national policy. Every country named has done valuable work. The question this paper asks is what comes next.

What the rest of the paper does

- **Part I — The case.** Six chapters on why the old development model no longer absorbs enough risk and why home-grown firms have become structural infrastructure.
- **Part II — The architecture.** Seven chapters on the layers underneath programmes — formation friction, procurement, capital design, universities, immigration, regions, and the AI shock.
- **Part III — Implementation.** The policy stack, the habits to retire, the national dashboard, the first twelve months, and named recommendations by office.

> A country that does not create employers will eventually run out of employment policy.

OPENING

A spreadsheet, a press release, and a quiet morning.

DUBLIN · 6:42 AM · A MONDAY

The minister is already awake.

There is a spreadsheet open on one screen and the morning news on another. A large employer has announced a restructuring. The statement is careful, polished, and bloodless. *"Operational efficiency." "A more focused global footprint." "Investment in automation." "Ongoing commitment to the region."*

The region hears something else. Three hundred jobs are gone this year. Another two hundred are under review. The company is not leaving. That is the part everyone is told to be grateful for. But the company that once needed 1,200 people now needs 700, and the press release does not say it out loud.

> The future did not arrive as a factory closure. It arrived as a smaller payroll.

Across the table, the senior civil servant has another document. The global outlook is weaker. Trade policy is uncertain. Investment decisions are slower. Energy costs are unstable. Export forecasts are being revised downward. A multinational board meeting in another country has just changed the economic plan of a town that had no seat at the table.

For decades, countries competed to attract large employers. They built industrial parks, tax regimes, grant schemes, training programmes, investment agencies, and glossy national propositions. Much of that work mattered, and still does. But large companies are becoming more productive with fewer people. Global supply chains are more political. AI lets firms do more with smaller teams. Capital markets reward efficiency. Corporate headquarters treat geography as a variable.

The minister scrolls down the spreadsheet. Employment supports. Redundancy payments. Retraining budgets. Regional response measures. A taskforce. A consultation. A press conference.

All necessary. None sufficient.

Because the real question is not only how to protect the jobs that exist.

> **The real question is who will create the next employers.**

Three kilometres away, on the third floor of a co-working space behind Pearse Street, somebody is writing a pitch deck.

They have not slept much. The product works — barely. The first customer might sign this month, or next, or might pull back to a procurement cycle they cannot afford to lose. The bank account is uncomfortable. A friend has offered ten thousand euros that the founder is not yet sure how to accept. The minister and the founder will never meet. What the minister decides this morning will, regardless, shape whether the founder is still trying in nine months' time.

It is worth being honest about what that founder is signed up for. Entrepreneurship is brutal. The glossy success stories that fill ministerial speeches hide a darker reality: for every visible win, hundreds of ventures collapse under the weight of market indifference, resource constraints, and sheer exhaustion. Founders do not just risk capital. They risk years of their lives, relationships, and often their mental health on ideas the world may never embrace.

And — this is the harder truth, the one that should sit uneasily with anyone designing innovation policy — the mindset that makes a serious founder cannot be taught in a classroom. The skills can: how to validate an idea, build a product, manage cash flow, navigate growth. The mindset — the resilience to persist through repeated rejection, the wisdom to know when to pivot versus

persevere, the irrational optimism required to believe you can change the world when all the evidence suggests otherwise — is forged in the work itself, and absorbed by osmosis from people who have walked the road before.

This is why community matters more than curriculum. It is also why most government entrepreneurship programmes, no matter how well-funded, fail to produce the founders the country needs. They confuse the artefacts of successful ecosystems — the co-working spaces, the demo days, the startup competitions — with the essence of what makes those ecosystems work: genuine risk-taking, authentic knowledge transfer, and the patient capital required to build something transformative rather than merely incremental.

> **The best ecosystems are gardens, not factories. You can prepare the soil and ensure access to water and sunlight — you cannot make the seeds grow by pulling on the shoots.**

Every additional layer of oversight, every prescriptive programme, every attempt to pick winners squeezes out the randomness and rebellion that drives breakthrough thinking. The answer is not the absence of structure. It is deliberate design instead of direct control: thoughtfully created frameworks that enable rather than constrain — access to capital without dictating how it is spent, connections to mentors without mandating their advice be followed, spaces for collision without orchestrating who collides with whom.

All of this would matter in any decade. It matters more in this one, because the underlying conditions have shifted.

The world is fracturing along geopolitical fault lines. Supply chains are being redrawn. Trade alliances are being questioned in public for the first time in a generation. Large corporations that once felt comfortable in their market positions now face a stark choice: engage meaningfully with young companies or risk obsolescence. The difference between slowly dying and thriving may well be a partnership with a five-person team working on something a corporate R&D department dismissed as impossible.

And then there is artificial intelligence. Not another technology trend tracked on a consulting firm's hype cycle, but a fundamental shift in our species' capabilities. The internet and mobile computing changed how people access and share information. AI is changing how people think, create, and solve problems. It will touch every industry, reshape every job, and redefine what it means to be productive, creative, and useful. The companies being built in this space today are not just companies. They are architecting the next decade. The country that decides where they are built is the country that decides where the next decade's tax base, employment, and political agency live.

Meanwhile, governments throw money at innovation through grants that take more time to apply for than to build a working prototype, and then wonder why the resulting startup scene lacks dynamism. Corporations launch innovation labs with exposed brick walls and ping-pong tables, staffed by people who have never felt the terror of making payroll. Accelerators accelerate the production of pitch decks. Venture arms invest like pension funds. Entrepreneurs themselves often arrive with their own unrealistic expectations — that a good idea alone deserves funding, that growth should be linear and pain-free, that the market owes them attention because they took the risk. *Everyone knows their lines, and nobody is writing a compelling story.*

This paper is about getting out of that theatre.

> **We're doing this wrong.**
> **We don't have to.**

— DC CAHALANE · DUBLIN · JUNE 2026

CHAPTER 01 · PART I

01 / 16

The Age of Fragile Growth.

Why the old economic development model no longer absorbs enough risk.

KEY TAKEAWAYS

01 The large-employer model still works, but it concentrates economic destiny inside decisions a country does not control.

02 Trade volatility, geopolitics, and AI are intensifying the concentration risk simultaneously.

03 Resilience is not isolation. It is enough internal company-formation capacity to absorb shocks.

04 The political language has not yet caught up with the structural shift; the work of the next decade is to update it.

For most of the late twentieth and early twenty-first centuries, economic development had a recognisable shape. Countries tried to attract large employers. Regions competed for factories, headquarters, shared-services centres, R&D sites, and major public-sector commitments. Politicians announced jobs. Agencies reported investment. Universities supplied graduates. Local authorities built roads and housing around the hope that large employers would create stability.

The model was not foolish. For small economies and peripheral regions, large employers brought capabilities that would have taken decades to build alone. They trained workers, professionalised supply chains, created export disciplines, and made countries legible to global capital. The model worked often enough to become the default.

For three generations this single bargain anchored the political economy of every advanced country. Industrial estates, vocational schools, infrastructure plans, immigration rules, tax codes — each was bent toward the gravitational centre of the large employer. The bargain offered a clear and visible win condition: the minister knew what success looked like, because it landed on the front page of a newspaper with a number of jobs attached.

That clarity was its own form of value. It is also why letting go of the bargain is harder than the policy analytics suggest. The new posture has no equivalent single number. The next decade asks ministers to manage a portfolio of compounding domestic capacity rather than a sequence of announceable wins.

But economic models can work and still become insufficient. The weakness of the large-employer model is not that large employers are bad. It is that the model concentrates too much economic destiny inside decisions a country does not control. A regional economy may depend on a plant whose future is decided in another country. A graduate labour market may depend on recruitment cycles set by capital markets, automation strategies, and geopolitical tensions.

This was always true. What has changed is the intensity. The OECD warns that higher trade barriers and policy uncertainty are weighing on investment and trade [3]; the World Bank describes increased trade tensions as a substantial global headwind [4]; the IMF identifies downside risks from protectionism, fiscal vulnerabilities, labour-supply shocks, and financial-market corrections [5].

For political leaders, this means something practical. The old promise was: attract investment, create jobs, build prosperity. The new reality is: attract investment, but assume volatility; build jobs, but assume restructuring; grow exports, but assume disruption; welcome multinationals, but build domestic companies that survive if global decisions turn against you.

This is not pessimism. It is resilience. Resilience does not mean isolation. It means enough internal company-formation capacity to absorb shocks, create new opportunities, and retain more of the value generated by a country's people, research, and capital.

FIGURE 2 · CONCENTRATION RISK BY AXIS

DIAGNOSTIC

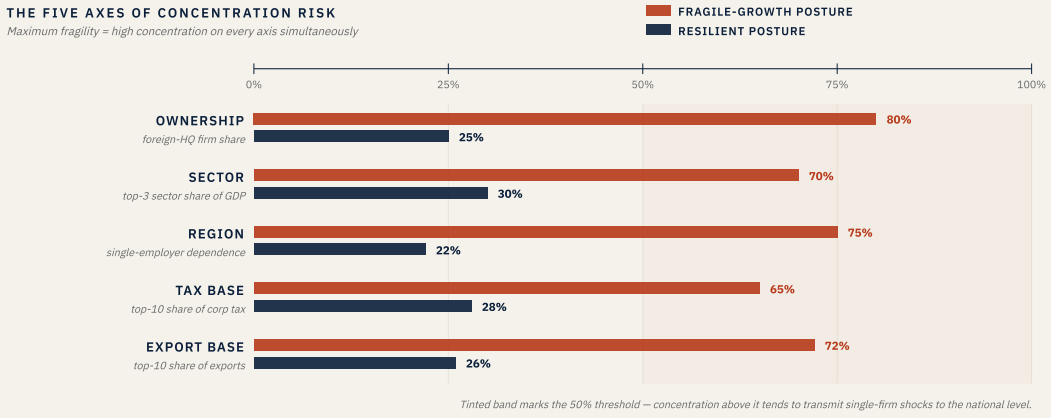

Concentration risk, mapped

Most ministers cannot answer one question without a brief: how concentrated is national employment inside foreign-headquartered firms, in any single sector, in any single region? The answer is usually more uncomfortable than the briefing paper suggests — because the five axes of concentration are rarely measured together, and the joint distribution is what matters.

CONCENTRATION AXIS	DIAGNOSTIC QUESTION	IMPLICATION IF HIGH
Ownership	What share of employer firms are foreign-headquartered?	Decision rights sit elsewhere; restructuring risk imported
Sector	What share of GDP comes from the top three sectors?	Single-sector shocks become national shocks
Region	What share of employment in any region depends on one employer?	One announcement closes a town
Tax base	What share of corporate tax comes from the top ten firms?	Fiscal exposure to firm-level decisions
Export base	What share of exports comes from the top ten firms?	Trade negotiations distorted by single-firm interests

None of these axes is, on its own, a problem. The problem is the joint distribution: countries that are concentrated on all five axes simultaneously are running maximum fragility. They are also unusually common. The OECD's Productivity Statistics show that across small advanced economies the top ten firms typically account for between a third and a half of corporate-tax revenue, and a similar share of exports [17].

The diagnostic is not abstract. Many of the most consequential economic-policy moments of the past five years — a microchip plant repositioned, a pharmaceutical mandate moved between sites, a financial-services unit relocated for regulatory reasons — are stories of concentration risk crystallising. None of those events were predicted by the standard headline indicators. All of them are visible on a five-axis radar drawn the year before.

> “
>
> **Stability and resilience are not the same thing, and they can pull in opposite directions.**

Three regime shifts, three lessons

The current fragility is not the first to be diagnosed. Three earlier regime shifts left lessons that the entrepreneurship economy can use. Each was originally read as a temporary crisis; each, in hindsight, was a permanent change in the operating environment for national economic policy.

SHIFT 01

The 1970s oil shocks

Energy-importing economies discovered overnight that sector concentration could become national crisis. The response — strategic stockpiles, energy diversification, sovereign wealth in some countries — became durable infrastructure. The entrepreneurship equivalent is firm-formation density: a strategic stock of domestic employers held against the next external shock.

SHIFT 02

The 1990s manufacturing offshoring

Countries that built skills and supply-chain capabilities outlasted those that competed on cost. The lesson for today is similar: complement, not substitution, for foreign capital. The countries that retained domestic capability around imported investment compounded; those that bid only on tax and labour cost found their factories portable.

SHIFT 03

The 2008 financial crisis

Economies dependent on a single financialised sector experienced the most violent restructuring. The slow rebuild was led by countries that had retained diversified domestic firm formation. The political lesson is that recovery capacity is in the firms that exist before the crisis, not in the programmes launched during it.

Each regime shift produced an institutional response. The next one — concentration risk in the age of AI, geopolitical fragmentation, and productivity-without-headcount — requires an institutional response too. The entrepreneurship economy is that response. The question is not whether countries will need it; the question is whether they will build it before the next shock or after.

What distinguishes the responses that worked from those that didn't, in each of the three earlier shifts, is the same property: the institutional reform was begun while the old model was still producing announceable wins. The countries that waited for the visible breakdown — the closed plant, the bank rescue, the queue at the labour exchange — spent the next decade catching up. The countries that began earlier compounded.

Diagnostic: are you running a fragile-growth strategy?

A short self-test for any government working its current portfolio. None of these questions is in itself definitive; all of them together are diagnostic.

01 Could the top ten employers in your country relocate any of their decision-making within five years without a public consultation? If yes, the decision rights sit outside the country.

02 Can your jobs strategy survive a single restructuring announcement of 800 roles in a regional town? If not, the strategy is balanced on a knife edge.

03 If foreign direct investment fell by 30 per cent for two years, what would the domestic firm-formation rate do? If the answer is "fall in parallel", the two are not independent — and one is a substitute for the other rather than a complement.

04 What share of your country's high-growth firms were founded inside the country in the past ten years? Twenty years? If the trend is downward, the renewal engine has slowed.

05 Can a minister name the senior official responsible for raising the company-formation rate? If not, the rate is nobody's KPI — and nothing without a KPI improves at the speed the moment requires.

06 How many of your country's twenty largest firms by employment are still headquartered domestically? If fewer than half, the long-run capital and talent recycling that mature ecosystems depend on is happening elsewhere.

The answers to these questions almost never produce comfortable reading on first attempt. **They are the starting point of an honest national entrepreneurship strategy, which is something different from the strategy most countries currently have.**

The mistake to avoid is treating the self-test as a scorecard. It is a diagnostic. A country scoring badly on every question is not failing; it is simply at the start of a useful conversation. The countries that have made the largest progress in the past decade are the ones whose senior officials walked into the test willing to read the answers honestly and act on them inside the political cycle that followed.

BRIDGE

Where this chapter lands

The diagnosis of fragile growth is the precondition for the policy stack proposed in Part III. If a country cannot see the concentration risk, it will not commit to the architecture that disperses it. Chapter 02 turns to the arithmetic that makes the architecture necessary.

CHAPTER 02 · PART I

The New Jobs Equation.

If companies are going to be half the size,
countries need twice as many companies.

KEY TAKEAWAYS

01 Software, automation, and AI have weakened the link between company growth and job growth.

02 If the average successful firm employs fewer people, the economy needs more successful firms.

03 Retraining and inward investment alone cannot close the gap; entrepreneurship must.

04 The math is not contested; the political implications are still being worked out.

The most important employment question of the next decade may fit on one line: what happens if successful companies need fewer people?

Joseph Schumpeter[21] described the entrepreneur as a combiner — someone who recombines existing resources, technologies, and labour into a new product, a new market, or a new form of organisation. A century later the description still fits. What has changed is the surrounding economics. The combiner now operates in an environment where capital markets reward efficiency, automation compresses headcount, and the gap between firm growth and job growth has widened. Schumpeter's gales of creative destruction still blow; what is new is that they no longer reliably scatter the seeds of mass employment in their wake.

Software reduced the people required for many business functions. Cloud reduced the need for internal IT teams. Marketplaces reduced distribution costs. Outsourcing modularised non-core work. Automation changed manufacturing and logistics. Remote work changed location. AI is now changing cognitive work, customer service, software development, legal process, finance, marketing production, administration, and research.

The result is not simply that jobs disappear. Firms become more productive without becoming proportionately larger. Stripe's payments business scaled to billions in revenue with a few thousand staff. Klarna publicly cut roles after AI deployment while continuing to grow [18]. Software companies routinely ship at scale with smaller engineering organisations than their predecessors. None of these are universal patterns, but a growing share of high-growth firms now exhibit them.

This is good for productivity. It is hard for employment. Asking companies to ignore productivity is not a serious national strategy. The serious strategy is to change the denominator. **If the average successful firm employs fewer people, the economy needs more successful firms.**

The wrong response is to treat the shock as a retraining problem alone. Retraining matters; lifelong learning matters; active labour-market policy matters. But a trained worker still needs an employer. If the number of employers is not growing fast enough, retraining becomes a queue-management system.

The new jobs equation

TREND	EMPLOYMENT CONSEQUENCE	POLICY IMPLICATION
Companies scale with fewer people	Job growth no longer follows revenue growth automatically	Create more companies, not just larger companies
AI automates tasks inside firms	Some functions shrink or reorganise	Treat founder creation as labour-market absorber
Large firms restructure globally	Local employment can fall despite firm profitability	Build domestic employer density
Supply chains fragment	Export and investment volatility rises	Increase local capacity and diversification
Young workers face fewer entry-level roles	Skills alone may not yield opportunity	Create more startups and early-company roles

FIGURE 3 · REVENUE AND EMPLOYMENT DECOUPLED

NEW JOBS EQUATION

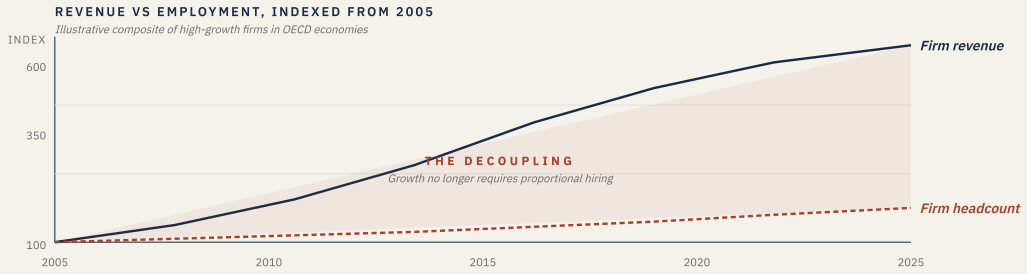

Sector by sector: where headcount compression is sharpest

The new jobs equation does not apply uniformly. A meaningful policy response distinguishes sectors that compress headcount fast from those that compress slowly or not at all.

FAST COMPRESSION

Cognitive-routine work

Customer support, paralegal, mid-skill admin, marketing production, software testing, basic finance. Headcount per output unit falling 20–40 per cent within a five-year horizon in firms that adopt available tools. Labour-market displacement is sharpest here.

REORGANISATION

Cognitive-judgement work

Software development, design, research, consulting, finance. Headcount stable or slowly declining; team shape changes; senior-to-junior ratio rises. The shock is to junior labour markets — exactly where founders are usually formed.

STABLE

Physical and regulated work

Healthcare delivery, construction, hospitality, regulated trades, in-person education. AI complements rather than substitutes; labour-market dynamics largely unchanged. Existing labour-market policy adequate.

The right response is to broaden the base of employers. More founders. More first-time employers. More spin-outs. More local suppliers. More exporters. More scalable SMEs. More immigrant founders, university founders, second-time founders, and founder-to-founder recycling. The question is not whether each becomes a unicorn. The question is whether enough become employers.

There is a related demand-side problem worth naming directly: the cost of being a founder has risen faster than founders' early-career incomes across most high-cost European capitals.

CASE STUDY

The cost-of-formation problem in high-cost European capitals

In Dublin, London, Paris, Amsterdam, and Stockholm, the post-tax income of an early-career graduate has fallen below the threshold at which founder-style risk is reasonable. A worked illustration: a graduate in Dublin earning roughly €52,000 takes home around €37,000 after tax. A one-bedroom rent of around €2,200 per month consumes approximately €26,000 of that. A small co-working desk and basic professional overhead consume another €7,000 to €8,000. Before paying themselves, the founder is around €50,000 below the breakeven of staying employed [22].

The point is not that founders need cheap rent. **Capital markets cannot solve a problem produced by housing markets.**

Retraining is necessary but not sufficient. A useful exercise for any government: estimate the size of the gap between the displaced labour pool and the employer-firm absorption rate.

Quantifying the gap retraining cannot close

- **Step 1.** Take the IMF's exposure-to-AI estimate for advanced economies — roughly 60 per cent of cognitive employment [2]. Assume only a fraction of this is displaced (say, 10–20 per cent over a decade).
- **Step 2.** Translate into worker numbers. In a country of five million workers, even the low end is 300,000 displaced workers over a decade.
- **Step 3.** Compare with current employer-firm formation rates. Most advanced economies form fewer than 10,000 new employer firms per year, of which fewer than 1,000 grow past ten employees within five years.
- **Step 4.** Compute the gap. Even on optimistic absorption assumptions, the formation rate would need to roughly double to absorb the labour-market displacement plus normal turnover.

The exact numbers depend on country size, sector mix, and assumptions. The shape of the answer does not. **The required formation rate is not 10 per cent higher than today's. It is a different order of magnitude.** Retraining without that doubling is queue management. Founder formation without retraining leaves displaced workers without paths. The two together are the entrepreneurship economy.

CHAPTER 03 · PART I

Entrepreneurship as Economic Sovereignty.

Why home-grown companies give countries more options.

KEY TAKEAWAYS

01 Home-grown companies are decision centres located inside the country.

02 Ownership matters: where wealth, IP, and reinvestment compound determines long-run capacity.

03 Sovereignty is not autarky; it is having enough domestic firms to give the country options.

04 The quietest form of sovereignty — where decision rights physically sit — is also the most consequential.

Sovereignty in economic policy is usually associated with currency, taxation, regulation, and trade. The quieter form of sovereignty is the location of decision-making capacity inside firms. A country can have political independence and very little economic agency if too many of its key decisions about employment, investment, and innovation are taken in boardrooms abroad.

Home-grown companies are decision centres inside the country. They retain headquarters functions. They keep significant intellectual property at home. They export from local addresses. They reinvest into local supplier ecosystems. They mentor and angel-invest into the next generation. They are not a substitute for foreign direct investment. They are a complement, and a hedge.

COUNTER-ARGUMENT

"Small economies should specialise, not diversify."

The economic-complexity literature argues that small countries grow fastest by deepening clusters around their existing comparative advantages, not by spreading effort across diverse sectors.

The argument here is not that countries should be sectorally indiscriminate. It is that *within* a chosen specialisation, more home-grown firms produce more options. Ireland's life-sciences cluster is more resilient with twenty domestic founders than with three multinational sites, even if the multinationals are larger.

Ownership matters because it shapes where wealth compounds. A successful subsidiary repatriates profit. A successful home-grown company tends to recycle wealth into local angel investment, philanthropy, mentorship, and follow-on companies. These are not romantic claims; they are observable in mature ecosystems and visible in the early years of every emerging one.

Mapping decision rights: where do your firms actually live?

A country's economic sovereignty is partly an empirical question. The map can be drawn — not by employment, but by the location of the decisions that matter.

DECISION CLASS	WHERE IT USUALLY SITS	SOVEREIGNTY IMPLICATION
Hiring & firing	Where the operating manager sits	Often domestic — but mandate-controlled
R&D direction	Where the global head of R&D sits	Almost always at the corporate HQ
Capital allocation	Board / CFO function	At HQ
Plant restructuring	Board, with country input	At HQ
Strategic exit	Board / shareholders	At HQ
Geographic footprint	Board / global strategy	At HQ

The pattern is consistent: the decisions with the largest national-economic consequences sit at headquarters. The question for any country is the share of its employer base whose headquarters are domestic versus foreign. The answer is rarely flattering and almost never publicly tracked.

Building decision-centre density

Decision-centre density is the quiet objective of an entrepreneurship economy. It cannot be achieved by foreign-investment attraction alone, even attraction of strategic-mandate FDI. It is built by domestic firm formation that survives to maturity with its decision rights intact.

- **Anchor late-stage retention.** Tax, regulatory, and capital-market design that keeps successful domestic firms domiciled rather than reincorporated abroad. The exit-recycling cycle (Chapter 09) does much of the work.
- **Make local capital deep enough to support late-stage rounds.** The Series C/D round is often where domestic firms become foreign firms — when the only available investors are based elsewhere and request reincorporation as a condition.
- **Build credible domestic acquirers.** A successful firm with no domestic acquirer is sold abroad. Domestic acquirer density is one of the least-tracked features of mature ecosystems.
- **Defend headquarters language in M&A.** Negotiated commitments around HQ location, R&D retention, and decision rights are standard in some jurisdictions and absent in others. The political cost of insisting is usually overestimated.

Sovereignty without protectionism

The argument for decision-centre density is sometimes confused with economic nationalism. The distinction matters.

NOT THIS **Protectionism**	Tariffs, export controls, subsidies aimed at insulating domestic firms from foreign competition. Politically tempting; economically expensive. Reduces national productivity over the long run.
BUT THIS **Decision-centre density**	Open trade, open investment, open talent — combined with deliberate effort to ensure that enough firms whose decisions affect the country are headquartered inside it. The state is not protecting firms from competition; it is ensuring there are firms whose mandate lives at home.

The countries that get this right run high-FDI policies alongside high-formation policies. The two are complements: strategic-mandate foreign investment is more attractive to a country with strong domestic firm density (Chapter 05), and domestic firm density is reinforced by talent and capital flows that high FDI brings.

> **Sovereignty in the entrepreneurship economy is the share of decisions that matter to the country which are taken inside it.**

CHAPTER 04 · PART I

The Employer Gap.

Why employment policy fails when it does not create employers.

KEY TAKEAWAYS

01 Most employment policy assumes the existence of employers and works on the supply side.

02 When employer creation slows, supply-side measures hit a ceiling.

03 Employment policy must own employer formation, not just employability.

04 In most countries, jobs and enterprise sit in different ministries with different KPIs. That separation is itself the problem.

Most employment policy assumes the existence of employers. It focuses on the supply side: skills, training, work-readiness, jobseeker support, wage subsidies, regional incentives, sector-specific programmes. These measures matter. But they all rely on a demand side that is taken as given.

When the demand side weakens — when fewer firms exist to absorb workers — the supply side hits a ceiling. A reskilled worker still needs an opening. A graduate with new credentials still needs a hire. If the number of employers is not growing fast enough, employment policy becomes a queue-management exercise dressed up as activation. This is not a failure of the people running the system. It is a failure of where the system applies its leverage.

Employment policy must own employer formation, not just employability. That means the same ministry, agency, or office that worries about jobseekers should also be accountable for the rate at which new employer firms come into being. In most countries those concerns sit in different departments under different ministers measuring different KPIs. **That separation is itself the problem.**

The future employer is unlikely to be the company that once employed the parents of today's graduates. It is more likely to be a young firm that did not exist five years ago, founded by someone who had to navigate twelve agencies and four tax regimes to start. The country that makes that journey shorter and clearer creates more employers. The country that does not, does not.

Why the jobs ministry should own employer formation

In most governments, jobs and enterprise sit in different ministries. Jobs reports on unemployment, retraining, vacancies, activation. Enterprise reports on FDI, sector strategy, business support. The two ministries rarely share KPIs. They often report on the same labour market without agreeing on what is happening to it.

The clean institutional fix is to make the same office accountable for both. There are three workable versions of this.

VERSION 01

Combined ministry

One cabinet portfolio covering jobs, enterprise, and entrepreneurship. Used in some smaller economies (Singapore, Estonia, several Nordic countries). Conceptually cleanest; politically expensive in countries where the two ministries have separate political histories.

VERSION 02

Joint accountability

Both ministers share a single set of KPIs — employer-firm formation, time-to-incorporation, survival to year three. The cabinet office assembles the dashboard. The political accountability is shared. Less dramatic than a merger; achievable inside one political cycle.

VERSION 03

Head-of-government office

A small unit in the Prime Minister's or Taoiseach's office holds the cross-ministerial KPI and the dashboard. Each ministry continues to operate; the unit is the integrator. Lightweight; depends on sustained political attention from the head of government.

Diagnosing the gap in your jurisdiction

A short audit a labour-market policy team can complete in two weeks, using existing data:

01 **Employer-firm formation rate.** Net new firms employing at least one person, per year, per 10,000 working-age adults. Compute the trend over ten years. A flat or declining trend is the diagnostic.

02 **First-employee firms.** The subset of new registrations that hire their first employee within 24 months. This is the meaningful population; pure registrations are noise.

03 **Survival to year three.** What share of new employer firms still employ at year three? The ratio matters more than the gross number.

04 **Hiring concentration.** What share of new hiring in the past year went to firms older than 25 years? Higher concentration here means fewer young firms are absorbing displaced workers.

05 **Public-spend allocation.** What share of jobs-policy spending is on employability (training, activation) versus employer formation? In most countries the answer is greater than 95 / 5. The reform target is 80 / 20 within a political cycle.

Questions to ask your jobs minister

A practical script for the head of government's office, the finance minister, or any colleague preparing for a jobs-portfolio meeting. Each question is designed to surface whether employer formation is being treated as a strategic objective or as someone else's problem.

- What is your KPI for net new employer firms this year? Is it published?
- What proportion of jobs-policy spending is on programmes that create employers, versus on programmes that train workers for them?
- Who, by name, in your ministry is accountable for the formation rate? Do they meet the enterprise minister's equivalent monthly?
- If foreign direct investment fell by 25 per cent for two years, what would happen to your employability targets?
- If the formation rate doubled over five years, what changes in your portfolio?

The point of these questions is rarely to embarrass — it is to surface whether the two halves of the labour-market problem are being held together by anyone at all. **In most countries, they are not. The cost of that gap is paid for in the next restructuring announcement.**

> A reskilled worker still needs an opening. Employment policy that does not create openings is solving the easier half of the problem.

CHAPTER 05 · PART I

05 / 16

The Multinational Illusion.

Why foreign direct investment matters,
but cannot be the whole strategy.

KEY TAKEAWAYS

01 Strategic-mandate FDI (R&D centres, decision HQs) is genuinely valuable and should be defended.

02 Satellite or branch FDI is more fragile and produces fewer local options.

03 The strategic question is not how much FDI; it is what kind, and what it sits beside.

04 Treating both kinds as equivalent in policy produces predictable surprises.

Foreign direct investment is not a single phenomenon. It comes in at least two distinct forms with very different long-run effects.

Strategic-mandate investment locates real decision-making, R&D, or product responsibility in a country. Intel's Leixlip operation runs a global wafer-fabrication mandate. IBM Research Zurich produces Nobel-prize-winning science. Pfizer Cork has been a centre of formulation science for decades. These are not branch offices. They are decision centres with deep local roots, supplier networks, and graduate-recruitment relationships that compound for decades.

Satellite or pass-through investment is different. Shared-services centres, contract-manufacturing sites, and routine sales offices are inherently more mobile. They are won on cost and won away on cost. They train people, generate tax, and create employment, but their decisions rest with a board elsewhere whose primary obligation is to its own shareholders.

Treating both as equivalent in policy has produced the awkward modern situation where a country can claim a record FDI year and still lose 800 jobs in a single restructuring announcement. The two are not contradictory. They are describing different categories of investment.

The strategic question is not how much FDI a country attracts. It is what kind, and what it sits beside. Strategic-mandate investment is best defended by domestic firm density. A country with two hundred home-grown software companies is a more attractive home for an R&D mandate than a country with twenty. The two strategies reinforce each other; **an honest entrepreneurship strategy strengthens FDI ambition rather than displacing it.**

Diagnosing your FDI portfolio

Most countries do not separately track strategic-mandate FDI and satellite FDI. The aggregate is reported; the composition is not. The diagnostic is straightforward to compute and unusually clarifying when it is.

QUESTION	WHAT THE ANSWER REVEALS
What share of FDI jobs are in roles where the decision-maker is local?	True mandate vs. satellite composition
What share of FDI R&D spend is on globally-mandated programmes?	Whether the country is in the value chain or downstream of it
What share of FDI revenue is exported from the country?	Mandate sites export; satellites are cost centres
How much of FDI hiring is at director-level or above?	Satellite operations rarely promote past mid-management
What share of FDI announcements in the last five years involved a strategic mandate (rather than capacity expansion of an existing satellite)?	The composition of the recent wave

A country running a successful FDI policy may have an unhealthy composition: high gross investment, high employment, but few sites with mandates. The reform is in the deal terms, not the volume.

Defending strategic-mandate FDI

Once a country has attracted strategic mandates, the policy challenge becomes retention. Mandates are at constant risk of being moved closer to headquarters or to lower-cost jurisdictions. The defence is rarely about subsidies; it is about embedded value.

- **Domestic firm density.** The single strongest defence. A site in a country with two hundred relevant home-grown firms is harder to relocate than one in a country with twenty. The site benefits from talent recycling, supplier networks, and a local market for its outputs.
- **Specialist supply.** Deep supplier networks lock in the mandate operationally. Building these requires sustained sectoral focus, not generic enterprise policy.
- **Long-run talent retention.** The country's labour market must keep producing senior people the corporate wants to promote. Otherwise the mandate gradually thins out — director by director — without any single announcement.
- **Quiet political work.** Mandate holders pay attention to whether the country defends their interests in EU/trade negotiations, in tax policy, and in regulation. The political signal is sometimes worth more than any financial incentive.

When satellite FDI is fine

Treating all FDI as either strategic or expendable is the wrong frame. Satellite operations have a role; the question is what they sit beside.

POSTURE 01 **Satellite FDI plus strong domestic formation**	The satellite operations train workers, generate tax, and provide volume employment. The domestic ecosystem absorbs shocks when satellites contract. This is the healthy posture.
POSTURE 02 **Satellite FDI without domestic formation**	The country depends on operations whose decisions are taken elsewhere. The next restructuring becomes a national event. This is the fragile posture — and it is more common than ministers admit.
POSTURE 03 **Strategic-mandate FDI plus weak domestic formation**	The mandates are valuable but exposed. Without domestic firm density, the country has no replacement capacity when any mandate moves. This is the medium-fragile posture.
POSTURE 04 **Strategic-mandate FDI plus strong domestic formation**	The two reinforce each other. Talent rotates between mandate sites and domestic firms; capital recycles; mandates are sticky because the underlying ecosystem is. This is the resilient posture.

The work of entrepreneurship policy is to move a country from postures 2 and 3 toward posture 4. **That work is not a substitute for FDI policy; it is the partner FDI policy has been lacking.**

CHAPTER 06 · PART I

06 / 16

The Local Multiplier.

Why a home-grown company compounds differently from an imported branch office.

KEY TAKEAWAYS

01 A successful home-grown company produces second-order effects branch offices rarely do.

02 Founder genealogies — alumni who go on to start the next company — are the engine of mature ecosystems.

03 Counting employees today is the wrong metric; counting future founders trained inside a firm is closer.

04 The data to track founder genealogies already exists; almost no country publishes it.

Endeavor's research on entrepreneurial multipliers, replicated across cities from Silicon Valley to Buenos Aires[11], shows that the most powerful long-run effect of a successful home-grown company is not its own employment. It is the founders, mentors, angels, and operators it produces.

PayPal alumni went on to found or co-found Tesla, LinkedIn, YouTube, Yelp, Palantir, and Affirm[11]. Stripe alumni have started dozens of companies; Klarna's, Skype's, Wise's alumni have populated the European fintech landscape. The named cases are dramatic; the underlying pattern is structural. A successful home-grown company creates the human capital and the financial capital that fund the next twenty companies.

FIGURE 4 · FOUNDER GENEALOGY, TWENTY YEARS ON

LOCAL MULTIPLIER

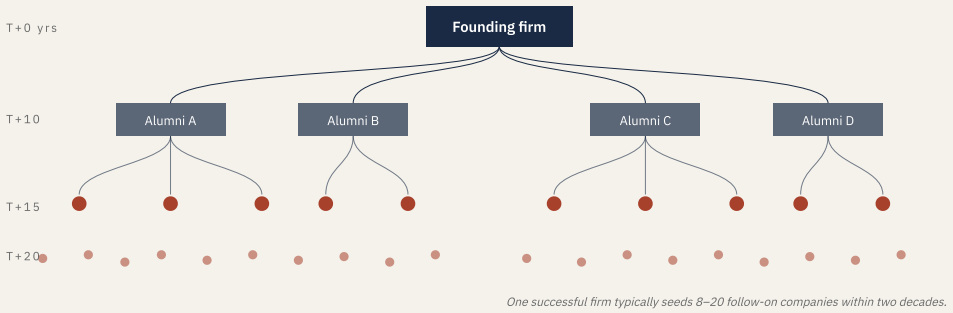

One successful firm typically seeds 8–20 follow-on companies within two decades.

This pattern operates only weakly inside satellite operations. Branch offices recruit, train, and promote. They rarely produce founders, because the founder pathway runs through firms that take strategic decisions locally, expose people to general management, and treat exits as opportunities for the next bet.

The right metric for a successful home-grown firm is not its peak headcount. It is the founder genealogy it produces a decade later. Most countries do not measure this. They could; the data exists in companies-house records and LinkedIn graphs. **The countries that begin to measure it will plan around it differently.**

CASE STUDY

Internally generated wealth across Europe

Some of Europe's most consequential companies were not the result of multinational investment but of internal founder creation. Ryanair began with a single route between Waterford and London-Gatwick and grew into Europe's largest airline by passenger numbers. Spotify reshaped the global music industry from Stockholm. SAP became the world's leading enterprise-software company from Walldorf. Skype gave rise to a generation of Estonian and Nordic founders who went on to build Wise, Bolt, and dozens of follow-on firms.

The transferable lesson is not that any one country can replicate any one of them — **internally generated wealth produces second-order effects that imported employment, however welcome, cannot.**

Mapping founder genealogies

The countries that have begun to track founder genealogies use the same playbook. The mechanics are unglamorous; the political return is high.

- **Step 1.** Identify the country's twenty most consequential home-grown firms of the past thirty years.
- **Step 2.** For each, list every employee in the first ten years (companies-house, professional-network, and alumni-record sources are typically sufficient).
- **Step 3.** Trace which of those alumni went on to found or co-found other firms.
- **Step 4.** For each follow-on firm, identify which became employers, which raised meaningful capital, and which produced their own founder alumni.
- **Step 5.** Visualise the tree. Publish it.

The result is invariably striking and politically useful. A single successful firm can be shown to have produced ten or twenty follow-on companies a decade later. Ministers can refer to specific firms when defending entrepreneurship investment. **The multiplier becomes visible.**

European parallels to the "PayPal mafia"

The most-cited founder genealogy is American — PayPal's alumni. Less-cited equivalents exist in Europe and elsewhere. Naming them publicly is part of the political work of an entrepreneurship economy.

STOCKHOLM **The Spotify / Klarna effect**	Several hundred alumni of these two firms have gone on to found, co-found, or become early operators at the next generation of Swedish technology companies. The Stockholm ecosystem is built on this base.
TALLINN **The Skype effect**	The single most-documented European case. Skype alumni founded or co-founded Wise, Bolt, Pipedrive, Veriff, and dozens more. A small country's entire technology economy was built on the alumni of one firm.
BERLIN **The Rocket Internet effect**	Imperfect cluster, but well-documented. Operators trained at Rocket have founded a wide cohort of follow-on firms across Europe and emerging markets.
DUBLIN / CORK **The Stripe / Intercom effect (in formation)**	Less mature, but the pattern is visible in early data: alumni founder rates substantially higher than baseline, with the first wave of follow-on companies emerging in the last five years.

What success at 20 years looks like

The mature outcome of a successful entrepreneurship economy is rarely a single dominant firm. It is a recognisable pattern of compounding founder genealogies, distributed across sectors, with regular new entries at the top.

- **Five to ten foundational firms.** Successful home-grown companies whose alumni populate the next wave. None individually larger than the country, but together they are an institutional layer.
- **Repeat founders as a visible category.** A non-trivial share of new firms started by founders with prior founding experience. Repeat-founder rate is one of the cleanest indicators of ecosystem maturity.
- **Domestic angel investment.** Wealth from earlier exits routinely flows back into new firms, mostly from operators, not financiers.
- **Visible mentorship.** Founders publicly help the next generation. Not as an industry; as a norm.
- **Acquirer density.** Mid-sized domestic firms acquiring smaller domestic firms — the ecosystem absorbs its own successes rather than losing them abroad.

No country starts here. The work is not building this state in one term; it is starting the compounding. **Twenty years is a long horizon. It is also exactly how long the work takes — and the countries that started fifteen years ago are already producing the genealogies that ministers point to today.**

> The right measure of a successful firm is not how many it employs today. It is how many founders it produces a decade later.

PART II · CHAPTERS 07 – 13

The architecture.

Seven chapters on the institutional layer that decides whether founders ever get to build — formation, customers, capital, universities, immigration, regions, and the AI shock that connects them.

CHAPTER 07 · PART II

07 / 16

The Missing Middle of Policy.

Why governments overfund programmes and underbuild formation architecture.

KEY TAKEAWAYS

01 Most entrepreneurship spending goes to programmes; most founder time is consumed by formation friction.

02 Formation architecture — digital ID, registration, taxation, banking, regulation — is the highest-leverage layer.

03 Estonia's e-residency is the canonical proof that this layer is buildable; the lesson is the philosophy, not the brand.

04 The blockers are institutional, not technical. Naming a single accountable owner is half the work.

Most countries fund entrepreneurship through programmes. Programmes are visible, announceable, and easy to evaluate by participation. They are also usually the wrong unit. A founder who spends three months getting registered, banked, taxed, and compliant has spent three months not building. Those three months are not solved by a programme. They are solved by infrastructure.

Formation architecture is the layer underneath the programmes. It is the digital identity, the company registration system, the tax registration, the banking access, the social-insurance registration, the cross-border payments rail, the regulatory clarity. When this layer is strong, programmes become marginal. When it is weak, programmes are compensating for friction the state itself created.

FIGURE 5 · THE FRICTION-TIME STACK

MISSING MIDDLE OF POLICY

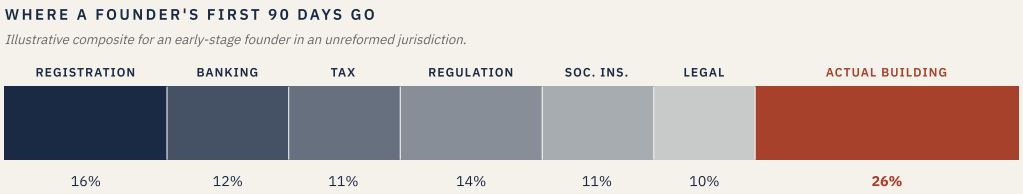

ROUGHLY 3 OF EVERY 4 FOUNDER-HOURS GO TO FRICTION THE STATE CREATED.

The point of formation architecture is to flip this ratio. A founder in Estonia spends a fraction of these hours; the building time expands accordingly.

Anatomy of formation friction

Most countries cannot tell a minister how long it takes a founder to legally start a company in their jurisdiction, because no single agency owns the end-to-end journey. Each agency optimises its own step; the founder's time is nobody's KPI. The table below maps the layer to the owner it usually needs.

FRICTION LAYER	WHAT IT CONTROLS	TYPICAL OWNER
Digital identity	Verifying a founder online	Interior / digital ministry
Company registration	Time-to-incorporation	Companies registry
Tax registration	VAT/PAYE onboarding	Revenue / tax authority
Banking access	Business account, payments rails	Central bank + commercial banks
Social insurance	Employing the first person	Social-insurance authority
Sectoral regulation	Permission to operate in regulated sectors	Sector regulators
Cross-border payments	Receiving foreign revenue	Central bank + finance ministry

A useful exercise for any minister: ask civil servants to walk through the journey personally — open a company, file the first VAT return, hire the first employee, receive the first international payment. Time it. Almost every country that does this is surprised by the answer.

CASE STUDY

Estonia — formation architecture as a national product

Estonia treats formation architecture as a national product. A founder anywhere in the world can become an e-resident [12], register an EU company online in roughly fifteen minutes, run that company digitally, and file tax electronically. As of end-2024, the e-residency programme had issued more than 120,000 IDs and supported more than 30,000 Estonian companies founded by non-residents.

The programme has paid for itself many times over in tax receipts and in soft-power influence. The genuinely transferable lesson is not the e-residency brand but the underlying philosophy: **treat the friction layer as a deliberate policy product, owned by a single agency, measured on time-to-incorporation and time-to-first-tax-filing.**

Three philosophies, three outcomes

Most countries' formation architecture sits in one of three postures. The posture is rarely chosen deliberately; it is inherited from the last decade of departmental settlements.

- **Product posture.** One agency owns the end-to-end founder journey. Time-to-incorporation is published. The friction layer is treated as a product with a roadmap. Estonia is the clearest example; Singapore, the UAE, and parts of the UK approach this for specific founder cohorts.
- **Portfolio posture.** Friction is managed across many agencies through coordination meetings. There is a strategy document but no single owner. This is the most common posture in mid-sized European economies. It is improvable but rarely improved.
- **Patchwork posture.** Each agency runs its own founder programme on top of unreformed friction. There is no consolidated view of the journey. This is more common than ministers realise; the test is whether anyone in the system can produce an end-to-end map on request.

Building the layer: a practical checklist

A government serious about formation architecture does not need to invent the technology. It needs to make five decisions in roughly this order. Each can be made inside a political cycle.

01 **Name a single accountable senior official.** Not a committee. The official's KPI is the published end-to-end founder time, measured quarterly. They have the authority to compel returns from departments.

02 **Publish the founder journey map.** A document, not a deck. Each step is named, owned, and timed. Friction points are visible. Founders can comment on it; the comments are public.

03 **Pick one digital identity rail.** Reuse whatever the country already has — a national eID, a digital wallet, a banking-grade KYC. Do not build a new one. Founder onboarding rides on the rail citizens already use.

04 **Set the time-to-incorporation target.** Estonia did fifteen minutes; the right target depends on the starting point. The target must be public, defended quarterly, and tied to the accountable official's role.

05 **Open a founder commentary channel.** A standing route by which founders can report friction they encountered last month. Read it; act on it; report what changed. The signal cost is low; the credibility return is high.

None of these steps requires legislation. All can be done administratively. The blocker is almost always institutional ownership, not legal authority.

Failure modes to watch

Most formation-architecture initiatives that fail do so for a small set of reasons. Each is predictable; each is avoidable if named in advance.

FAILURE 01 **Ownership is shared, so accountability is nobody's**	When the journey crosses five agencies and the responsible official is the chair of an inter-departmental committee, no agency feels the cost of friction. Publishing one number — total time, end-to-end — concentrates accountability where it belongs.
FAILURE 02 **The new system layers on top of the old one**	A founder portal is launched but the underlying agencies still demand paper, signatures, or in-person visits. The portal becomes a thirteenth step. The reform must replace, not append.
FAILURE 03 **The KPI is "people who used the portal"**	Activity metrics return through the back door. The right metric is the founder's total time and cost, not the agency's throughput.
FAILURE 04 **Banks are left out of the scope**	Most countries' formation friction is now concentrated in opening a business bank account. A formation-architecture programme that does not negotiate with the commercial-banking sector is solving the easy half of the problem.

The architectural layer is the cheapest, most under-celebrated reform in the entrepreneurship portfolio. It is rarely announced and rarely contested. **That is precisely why it is worth the political effort to own it.**

CHAPTER 08 · PART II

Government as First Customer.

Why procurement is an underused lever — with realistic limits, and what to build instead of another generic scheme.

KEY TAKEAWAYS

01 Public procurement can give a credible early company its first reference customer — sometimes its only path to existing.

02 The track record is mixed: SBIR, GovTech, and EU pre-commercial procurement produce winners, not high conversion rates.

03 Designing for narrow, high-fit pilots beats announcing large generic schemes.

04 The lever works in a defined set of sectors. Treating it as universal is part of why generic schemes underperform.

A company begins to feel real when someone pays it. For a serious early company in a regulated or institutional sector, the first paying customer is usually disproportionately hard to win, because trust is not yet earned and procurement systems are designed to manage risk rather than absorb it.

Government can be a powerful first customer in principle. In practice, the track record is mixed. The US Small Business Innovation Research programme has produced national champions but its conversion rate from Phase II to durable employer firm remains modest. UK GovTech procurement reform efforts have repeatedly stalled. The EU's pre-commercial procurement instrument has consistently underperformed its design intent.

What works, what doesn't

Procurement reform is one of the few entrepreneurship policies with thirty years of comparative evaluation behind it. The patterns of success and failure are unusually clear. Designs that work share five characteristics; designs that fail share their opposites.

WHAT WORKS	WHAT FAILS
Narrow problem framing	Broad open calls inviting "innovation"
A named buyer with a real budget	A challenge fund attached to no department
Short pilot timelines (90–180 days)	Multi-year frameworks before payment
A clear conversion path to a budget line	Pilots with no contractual route to renewal
A small unit allowed to depart from standard rules	Standard purchasing rules applied verbatim

The mechanism is straightforward. Early companies cannot survive eighteen-month procurement cycles. They can survive a four-month pilot that pays a real fee and leads to a real renewal. The state's job is to engineer the latter, not the former.

> The first paying customer is the moment a startup becomes a company. Government can be that customer — for credible companies in defined sectors — if procurement is designed for it.

Designing a startup procurement unit

The most replicable form of procurement reform is a small specialist unit with the authority to run pilots that the standard purchasing rules would not allow. Variants exist in the UK, France, Singapore, and the United States. The pattern is consistent enough to be specified.

- **Size.** Six to twelve people, including procurement specialists, technologists, a legal lead, and operational buyers seconded from line departments. Smaller fails to land; larger becomes a bureaucracy.
- **Mandate.** Authority to run pilots with firms under five years old, departing from standard purchasing thresholds, with audit-trail safeguards. The mandate is explicit and time-bounded.
- **Sponsor.** A senior minister or permanent secretary who personally chairs a quarterly review. The unit's authority is borrowed from this sponsor.
- **Pilot cadence.** Six to twelve pilots per year. Each pilot has a named buyer in a line department, a fixed timeline, and a published outcome.
- **Conversion target.** A published target for pilot-to-budget-line conversion (typically 25–40 per cent in the first three years). Visible, defended, used to refine the next year's pilots.

The unit does not run all government procurement; it runs the specific category where the friction is unsolvable inside the standard rules. The rest of the system can be left alone.

COMPARATIVE CASE

SBIR · GovTech · pre-commercial procurement — what each got right and wrong

US SBIR[15] (1982 –). Substantial scale, durable political support, well-defined Phase I/II/III architecture. Produces national champions. Conversion rate from Phase II to durable employer firm is modest; the strongest claim is that it would not have happened without the programme, not that the conversion rate is high.

UK GovTech Catalyst (2018 – 2023). Strong design, weak implementation. Pilots were funded; conversion to renewal contracts inside departments remained the bottleneck. Recent iterations have improved by attaching budget lines earlier.

EU pre-commercial procurement. Sound on paper, slow in practice. The cross-border legal framework adds friction that domestic schemes do not face. Useful where the buyer is a multi-country consortium; weaker as a national instrument.

The shared lesson is unglamorous: **the constraint is rarely on the supply side.** Founders can usually be found; the friction is in moving a credible early company from a pilot to a recurring budget line inside a line department. The reforms that work concentrate political effort there.

Sector playbook

Procurement is not a universal lever. It moves the needle in sectors where the public is a major buyer, where credibility is hard to manufacture privately, and where pilots can be scoped narrowly. In other sectors it is a distraction.

SECTOR 01 **Health technology**	National health systems are dominant buyers with high credibility transfer. Pilots in single hospitals, scaled to trusts and then nationally, are a proven sequence. The conversion bottleneck is clinical evaluation, not procurement law.
SECTOR 02 **Defence & security**	Long procurement cycles offset by very high willingness to pay for proven solutions. Dual-use commercial spin-off is common. Requires specialist procurement rails that most countries already operate.
SECTOR 03 **Critical infrastructure**	Energy, water, transport. Pilots run by infrastructure operators rather than central government. The state's role is regulatory permission for pilots, not direct purchase.
SECTOR 04 **Climate & sustainability**	Sub-national governments — cities, regions — are often the most credible early buyers, particularly for hardware and infrastructure pilots. National programmes that route through local buyers consistently outperform those that don't.
SECTOR 05 **Civic technology**	High-impact, low-political-risk pilots. Conversion bottleneck is internal IT integration, not procurement law. Best run alongside a digital-government reform programme.

CHAPTER 09 · PART II

Capital is Not Enough.

Why money without customers, talent, and speed does not create companies — and how to design state capital that catalyses rather than substitutes.

KEY TAKEAWAYS

01 Public capital should mobilise private capital, not replace it.

02 Capital without customers and talent does not create companies. It creates compliance.

03 The Yozma fund and the UK SEIS/EIS architecture are the clearest demonstrations of catalytic design.

04 Bad capital is a real cost. Founder-unfriendly early rounds destroy more value than capital scarcity.

There is a recurring temptation in entrepreneurship policy to declare a fund. Funds are easy to announce, easy to size, and easy to publicise. They are often the wrong unit. A fund without customers, talent, regulation, and exit pathways is a distribution mechanism for capital into an environment that cannot absorb it.

Public capital works best when it crowds in private capital rather than substituting for it. Matched funds, co-investment vehicles, tax-relief schemes that make angel investment attractive at the private margin, and pension-capital pathways that allow institutional money into venture have all been shown to work. Direct state-as-investor schemes have a worse track record, but matched and co-invested approaches have a strong one.

FIGURE 6 · THREE LAYERS OF CAPITAL ARCHITECTURE

CAPITAL DESIGN

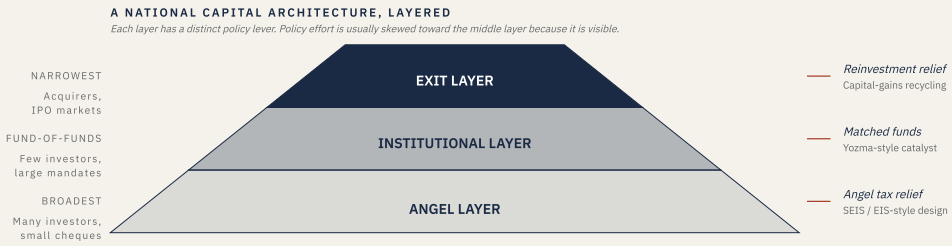

CASE STUDY · 1993

Yozma — Israel's matched-fund architecture

In 1993, Israel's Yozma programme [13] committed roughly $100 million of state capital to seed ten matched venture funds, each backed by a foreign general partner alongside Israeli investors. The state took a junior position with a buy-out option for the GPs. Within seven years the funds had been privatised, the state had recovered its capital, and the country had a functioning venture industry.

What Yozma demonstrates is not the size of the cheque. It is the design choice: **state capital was the catalyst, private capital was the engine, and the state stepped back deliberately once the engine started.** The transferable lesson is the structure, not the brand.

CASE STUDY · 2012 –

UK SEIS and EIS — capital architecture by tax relief

The UK's Seed Enterprise Investment Scheme and Enterprise Investment Scheme give individual investors substantial tax relief on early-stage investment, alongside capital-gains relief on successful exits. Cumulative SEIS investment from 2012 to 2023 exceeded £1.7 billion across more than 18,000 companies.

The scheme is not perfect — it concentrates capital in London, and some flows have ended up in tax-driven structures rather than serious companies — but the architecture has materially expanded the volume and quality of UK angel investment.

Bad capital is real. Founders take money from investors who cannot help them, on terms that make the company unfundable later, in rounds that solve the wrong problem. A serious entrepreneurship economy is honest about this and provides better signals for which capital is helpful and which is not.

The three layers of a national capital architecture

Most countries describe their capital ecosystem as a single thing. It is usefully decomposed into three layers, each with a distinct policy lever.

- **Angel layer.** Individuals investing their own money in companies they understand. Tax relief (SEIS/EIS-style) is the dominant lever. State direct investment is the wrong instrument here.
- **Institutional layer.** Venture funds, pension capital, sovereign wealth, family offices. Matched-fund schemes (Yozma-style) and pension-capital regulatory pathways are the dominant levers. Direct state investment is acceptable as a junior partner, never as the lead.
- **Exit layer.** Acquirers, IPO markets, secondary buyers. Largely outside the entrepreneurship portfolio but worth tracking: an ecosystem with no exit liquidity will eventually starve its own institutional layer.

Policy effort is usually skewed toward the institutional layer because it is visible. The angel and exit layers move more companies but require less announceable interventions. A serious capital strategy works on all three.

Designing a matched fund: operating guidance

A matched fund is the most replicable institutional-layer instrument. Most variants fail in implementation, not in design. The Yozma pattern is a good starting point; the operating choices that make or break it are well known.

01 **Match ratio.** 1:1 to 1:2 state-to-private. Higher ratios distort private-capital incentives and turn the scheme into a soft subsidy. Lower ratios fail to crowd capital in.

02 **State posture.** Junior position with a buy-out option for the private GP. Not first-loss capital; not equal-rights capital. The asymmetry is the whole point.

03 **GP selection.** Foreign GPs alongside domestic ones. The foreign GP brings discipline, network, and signalling. Pure domestic schemes consistently underperform.

04 **Time-boxing.** Five to seven years to demonstrate viability; explicit sunset language. State capital is not a permanent feature.

05 **Reinvestment clause.** Returns to the state can be reinvested in a successor scheme or returned to the exchequer. The former produces compounding; the latter signals discipline. Either is defensible; not deciding in advance is not.

The scheme's success is measured on private capital crowded in, not on portfolio-company outcomes. State funds that publish company returns rather than capital-multiplier returns are measuring the wrong thing.

Pension capital, ESOP reform, and the exit-recycling cycle

Three less-visible reforms have outsized capital effects, partly because they are politically uncontroversial.

PENSION CAPITAL

Regulatory permission, not state spending

Most European pension funds cannot allocate meaningfully to venture asset classes because of prudential rules written for a different era. Regulatory reform — explicit permission, with appropriate diversification caps — unlocks long-term institutional capital that does not exist in domestic venture funds today. The UK Mansion House compact and Australia's super-fund reforms are working examples.

ESOP TAXATION

The hidden tax on early hiring

When employee share options are taxed at grant rather than exercise, or when paper gains are taxed before liquidity, early hiring at startups becomes economically irrational. Reform is technical, cheap, and one of the highest-leverage interventions on early-talent supply.

EXIT RECYCLING

Where successful-founder wealth goes next

Successful founders recycle wealth into angel investment, mentoring, and follow-on companies — but only if the tax architecture rewards it. Capital-gains relief tied to reinvestment in early-stage companies is the cleanest instrument. The fiscal cost is small; the ecosystem return is observable in any mature cluster.

None of these reforms involves writing a state cheque. All of them are politically modest, technically straightforward, and disproportionately effective. They are the ground game of capital policy, and they are routinely under-prioritised in favour of more announceable funds.

CHAPTER 10 · PART II

Universities, Research, and the Founder Pipeline.

A third mission alongside teaching and research — and the institutional reforms that make it real.

KEY TAKEAWAYS

01 Public research is the largest underused source of company formation in most countries.

02 Most spin-out architectures are designed by technology-transfer offices optimising for institutional risk, not founder success.

03 Reframing entrepreneurship as a university's third mission allows the conversation to start on terms university leaders already use.

04 Reform is mostly administrative, not legislative. The blocker is institutional culture, not law.

Public research is the largest underused source of company formation in most countries. Universities and research institutes hold deep technical expertise, motivated graduate students, and intellectual property generated at public expense. Yet the rate at which this knowledge converts into operating companies is consistently low, and the standard explanation — "academics are not founders" — is too simple.

The harder truth is that most spin-out architectures are designed by technology-transfer offices optimising for institutional risk, not founder success. Equity terms are often punishing. Decision timelines are slow. Pre-incorporation IP rules are unclear. Researchers face unclear leave policies. Conflict-of-interest frameworks treat the act of starting a company as suspicious by default rather than valuable by design.

Why most spin-out architectures fail founders

The friction has a pattern. It is rarely the result of a single bad rule; it is the accumulation of risk-averse defaults across four areas.

FRICTION AREA	TYPICAL DEFAULT	WHAT IT COSTS THE COUNTRY
Institutional equity	20–40% retained by the university	The spin-out is unfundable at the next round
Decision timeline	6–18 months from disclosure to deal	Founders leave; companies form abroad
Pre-incorporation IP	Unclear what the researcher owns vs the institution	Investors cannot diligence; deals collapse
Founder leave	No written policy; case-by-case	Researchers must quit to find out

The pattern is institutional, not individual. Technology-transfer officers are usually thoughtful people working inside a system that rewards risk avoidance. The reform is to change the institutional default, not the individuals.

> The standard explanation that "academics are not founders" is convenient and wrong. Most academic founders did not stop being founders — their institutions stopped them.

Founder-friendly defaults: the new baseline

A small group of research-intensive universities — Oxford, Stanford, ETH Zurich, MIT, the Crick Institute, and a growing set of European peers — have converged on a similar set of reforms. The pattern is replicable; the political work is convincing institutions to adopt it.

01 **Standardised, founder-friendly spin-out terms.** Single-digit institutional equity for early-stage spin-outs (typically 5–10 per cent), pre-negotiated and published. The university wins by having more spin-outs, not by holding larger stakes in fewer.

02 **Decision SLA.** A published time-to-deal target (e.g. 60 days from disclosure to signed term sheet). Missed SLAs flag in the vice-chancellor's quarterly report.

03 **Pre-incorporation IP clarity.** A written policy that says, before any disclosure, what the researcher will be able to take with them. Published openly; not negotiated case by case.

04 **Founder leave.** A named, written sabbatical pathway — typically 12–24 months — with right of return. No founder should have to choose between research and a company because the university could not write a policy.

05 **PoC funding tied to formation.** Proof-of-concept grants that measure progress toward company formation, not toward another publication.

Each of these is administratively cheap. The aggregate effect is to move the institutional default from "block by default" to "support by default" — which is the entire intervention.

KPIs that move university behaviour

Universities respond to KPIs the way every other institution does. The current KPI set — publications, citations, grants, rankings — rewards exactly the behaviour the entrepreneurship economy needs less of: incremental risk-averse output. A reformed KPI set adds three measures without removing the existing ones.

KPI 01 **Spin-outs per €m of public research spend**	The right denominator. Countries currently report total spin-outs, which rewards size; this KPI rewards efficiency. Top-quartile research-intensive universities produce 0.5–1.5 spin-outs per €m of research spend; most produce a tenth of that.
KPI 02 **Alumni founders within 10 years of graduation**	Measures the university's effect on its own students, not just its faculty. Crucial for the founder genealogy of the next generation. Easy to compute from companies-house data and alumni records.
KPI 03 **Post-formation fundraising of spin-outs**	Whether spin-outs survive the first capital round. A spin-out that cannot raise is a registration, not a company. This KPI also exposes spin-outs created mainly to satisfy other KPIs.

Including these three in the public funding formula, even at low weight, changes vice-chancellor behaviour faster than any number of strategy documents.

The conversation with universities

University reform is one of the rare areas where the political work is harder than the technical work. The arguments that land — and those that don't — are well-rehearsed.

FRAME **Third mission, not commercial pressure**	Universities have always had two missions: teaching and research. A third mission — economic and social outcomes through company formation, civic engagement, and applied translation — is increasingly recognised, including by traditional research-intensives. Starting the conversation here changes its temperature.
FRAME **Founder alumni are a public good**	A successful spin-out produces alumni who donate, mentor, and angel-invest into the next generation. The university's long-run interest is to produce founder alumni, not to extract maximum equity from any single spin-out.
FRAME **The state's tools are funding and licence to operate**	Ministers cannot order universities to reform. They can adjust public funding formulas, publish league tables of the proposed KPIs, and ask vice-chancellors to defend their numbers. That is enough.

The countries that move fastest on university reform tend to share a feature: a small group of vice-chancellors who decide collectively to adopt the new defaults, and publish a joint statement. The state's job is to make that group's job easier.

CHAPTER 11 · PART II

The Immigration Question.

Why founder attraction is future-employer policy —
and how to design a founder visa that actually works.

KEY TAKEAWAYS

01 Founder immigration is one of the highest-leverage policies a country has.

02 The right test for a founder visa is the company being built, not personal income or net worth.

03 Immigrant founders disproportionately create employers, not just take jobs.

04 The scheme does not need to be perfect. It needs to exist, be visible, and be reasonably reliable.

Immigrant founders are a disproportionately important source of employer firms in every country with reliable data. They start more companies per capita, hire more aggressively, and are more likely to build internationally-oriented businesses than the native-born population. This is one of the most consistently replicated findings in labour economics.

Most countries' immigration systems are not built around this fact. They are built around employee work permits, skilled-worker quotas, and family reunification — useful categories for other purposes, but ill-fitted to founders, who do not have an employer because they are about to become one.

Five characteristics of a working founder visa

Countries that have run founder visas at scale — France, Estonia, Canada, Chile, the UK — converge on a similar pattern. Schemes that diverge from this pattern tend to fail in predictable ways.

01 **A fast track distinct from skilled-worker permits.** Founders cannot be processed through pipelines designed for employees of existing companies. A separate visa class with its own processing route is essential.

02 **A company-first test, not an income test.** The right question is what the founder is building — credibility of the venture, market potential, prior founding history — not their personal salary. Income tests filter for wealth, not capability.

03 **A credible runway period.** Two to three years before commercial milestones are required. Companies do not generate revenue in twelve months. Visas that require revenue at the twelve-month renewal force founders to optimise for survival, not building.

04 **Permanent residency tied to firm formation.** Not to salary, not to investment raised. The state's interest is that the company exists and employs people; the visa should reward that, not proxies for it.

05 **Active marketing in target source countries.** Schemes that exist on paper but are not actively promoted attract a fraction of the founders they could. Marketing is part of the scheme, not optional.

COMPARATIVE CASE

France · Estonia · Canada · Chile — what each got right

France · French Tech Visa (2017 –). Strong design, well-marketed, integrated with French Tech ecosystem support. Provides clear pathways for founders, employees of startups, and accompanying family. The scheme's strength is institutional buy-in across ministries.

Estonia · Startup Visa (2017 –). Tightly integrated with the country's broader digital-residency offer. Streamlined process, high transparency, and low friction. The model for small economies.

Canada · Start-Up Visa (2013 –). Permanent residency pathway from day one. Requires endorsement from a designated investor or incubator. Imperfect — the designated-organisation list is uneven — but the permanent-residency anchor is the single most-imitated feature.

Chile · Start-Up Chile (2010 –). Originally an equity-free grant programme that doubled as a founder-attraction vehicle. The first cohorts were transformative; the programme has narrowed since. The original design remains widely cited.

The shared lesson is unsentimental: **countries that treat founder immigration as a marketing problem outperform those that treat it as a legal problem.**

Designing a founder visa from scratch

A country starting from no founder visa today can have a working scheme within twelve to eighteen months. The sequence is well understood; the friction is political, not technical.

- **Step 1: Convene the inter-ministerial group.** Immigration, finance, enterprise, foreign affairs. Without alignment across these four, the scheme will stall in implementation.
- **Step 2: Publish the founder test.** What evidence the applicant must present. Make it specific. Vague tests produce inconsistent decisions and discourage applicants.
- **Step 3: Set the processing SLA.** Eight to twelve weeks from application to decision. Slower than this, and the scheme loses to faster jurisdictions for the same founders.
- **Step 4: Define renewal and permanent-residency rules upfront.** Founders need to know the full multi-year pathway before they commit. Year-by-year discretion is the most common reason talented founders pick a different country.
- **Step 5: Brief and brand.** Build a one-page site, a clear name, and a contact route. Founders find schemes through word of mouth — what is being said about the scheme matters as much as the scheme itself.

None of these steps require legislation in most jurisdictions. They require administrative coordination and a senior official willing to own the scheme publicly.

Political objections and how to answer them

Founder immigration is rarely contested on its merits but is occasionally caught in broader immigration debates. The objections that arise are predictable; the answers that work are too.

OBJECTION 01 **"We should support native founders first"**	Founder visas are not zero-sum with native founders. Immigrant founders disproportionately employ native workers, hire native co-founders, and recycle wealth domestically. The categories are complements, not substitutes, in every dataset.
OBJECTION 02 **"This is a back door for low-skill migration"**	Founder visa caps are typically two to four orders of magnitude smaller than broader immigration flows. The objection is real for some immigration debates; it does not apply to founder schemes designed with a company-first test.
OBJECTION 03 **"What if the company fails?"**	Most do. The right design rewards founder formation, not single-company success. A founder whose first company fails and who builds a successful second one inside the country is exactly the outcome the scheme should produce.
OBJECTION 04 **"What about housing capacity?"**	Legitimate at scale; not at the scale of any founder visa. Schemes typically attract hundreds to low thousands of founders annually. The housing conversation belongs in a different policy area.

CHAPTER 12 · PART II

Regional Renewal.

Why entrepreneurship is the most practical answer to hollowed-out places — with the base rates, the failures, and the patterns of success.

KEY TAKEAWAYS

01 Most regional ecosystem programmes fail. The honest base rate is low.

02 What the failures share is generic design copied from successful clusters that had different starting conditions.

03 Regional success requires anchor industries, repeat founders, local capital, and a decade-plus horizon.

04 It is often the least-bad available answer when the alternative is permanent decline.

Most regional ecosystem programmes fail. This is not a controversial observation in the policy-evaluation literature. Entrepreneurship-led regional renewal has a poor track record, and pretending otherwise damages the credibility of the better programmes that do exist.

What the failures share is generic design. They copy the visible features of successful clusters — incubators, accelerators, branding, conferences — without the underlying conditions that produced the original cluster: an anchor industry that generated repeat founders, local capital from successful exits, dense supplier networks, and a decade-plus horizon. The visible features are downstream of the underlying conditions, not a substitute for them.

COUNTER-ARGUMENT

"Regional entrepreneurship policy is mostly theatre."

The standard critique of regional ecosystem-building is that it creates buildings without companies. The empirical record gives the critique substantial support.

The honest answer is not to deny the critique but to ask what the alternative is. Industrial-attraction policy in hollowed-out regions has its own track record of failure. Regional entrepreneurship is not always the right answer; it is, often, the least-bad available answer, particularly when the alternative is permanent decline.

Regional success, where it has occurred, has shared a small set of features: an anchor industry that produced repeat founders; a credible source of local capital independent of the capital city; visible second-time founders mentoring the first wave; and patient public funding measured on a ten-to-fifteen-year horizon rather than a political cycle. None of these are easy. All of them are achievable. **None of them resemble the typical regional entrepreneurship strategy launched at a press conference.**

What anchor industries look like

An anchor industry is not a single large employer. It is a sector with enough density, talent rotation, and exit activity to produce repeat founders. The pattern is observable across mature regional clusters; the same pattern is detectable in regions before the cluster forms.

- **Talent rotation.** People move between firms regularly. The talent supply is shared across the region, not locked inside a single employer.
- **Exit liquidity.** Some firms have been acquired or floated. Their alumni have liquidity and reinvest locally.
- **Supplier density.** Specialist suppliers — lawyers, accountants, recruiters, sector-specific consultants — exist in the region, not only in the capital.
- **Recurring founder events.** Not the conference circuit; the working dinners, alumni groups, and informal gatherings that build trust between operators.
- **A local capital pool.** Independent of the capital city. Even modest — €20–50m in regional angel and seed activity per year — sustains a cluster.

Regions without these features should not run an entrepreneurship-led renewal strategy as the primary instrument. They should run one as a complement to other regional-development tools, with realistic expectations.

> The visible features of a successful cluster — the incubators, the branding, the events — are the consequence of underlying conditions, not a substitute for them.

Planning beyond the political cycle

The most important asymmetry in regional entrepreneurship policy is the mismatch between the timescale of cluster formation (10–20 years) and the political cycle (3–5 years). A serious strategy designs around the asymmetry rather than ignoring it.

MECHANISM 01

Cross-party endorsement

Regional strategies that survive a change of government tend to have explicit cross-party endorsement from inception. The political cost of opposing them is higher than the political return of dismantling them.

MECHANISM 02

Statutory regional bodies

Bodies with statutory authority and ring-fenced funding survive ministerial reshuffles. Bodies that rely on annual budget approval do not. Statutory standing is administratively cheap and politically valuable.

MECHANISM 03

Visible founder champions

A region with two or three successful, visible founders who advocate for the strategy gives ministers political cover. Identifying and supporting these champions early is part of the strategy.

MECHANISM 04

A published 15-year roadmap

Roadmaps with annual milestones are tactical; the strategic horizon is 15 years. Publishing both, and reporting against both, normalises the long horizon publicly.

Failure analysis: the four patterns

Failed regional programmes are not failures of effort. They are failures of design. The same four patterns recur across the policy-evaluation literature.

FAILURE PATTERN	WHAT IT LOOKS LIKE	WHAT IT ACTUALLY NEEDS
Building-first	Incubator opened; tenants recruited; no demand pipeline	Customer access before buildings
Branding-first	"Innovation district" rebrand; no underlying cluster	Anchor-industry density first
Event-first	Annual conference; no follow-through capital or buyer flow	Standing founder cohorts; ongoing buyer engagement
Programme-first	Accelerator cohorts; no exit liquidity to recycle	Patient local capital and acquirer relationships

A useful test for any regional strategy: ask whether it would still produce companies if the highest-visibility intervention were removed. If the answer is no, the strategy is the wrong shape. The underlying conditions — talent, capital, suppliers, customers — must do most of the work; the visible programmes are the surface.

The countries with the best regional outcomes have, in recent decades, been those that did less, more deliberately, over longer horizons. That is harder to announce. It is also what works.

CHAPTER 13 · PART II

The AI Employment Shock.

Why the labour-market answer to smaller companies is more companies — and the window in which policy can shape the outcome.

KEY TAKEAWAYS

01 AI changes the relationship between productivity and headcount across many sectors.

02 Frontier AI labs are real but small; they cannot absorb the scale of structural shift.

03 AI also lowers the cost of starting a company — the policy task is to convert that lower cost into more firms.

04 The window during which a country can shape the outcome is open and narrowing.

AI is changing the labour intensity of cognitive work in ways that the previous waves of computerisation did not. The IMF estimates that almost 40 percent of global employment is exposed to AI; in advanced economies the figure is closer to 60 percent. The structural consequence — even on optimistic assumptions about complementarity — is fewer people required for a given output in many cognitive functions.

COUNTER-ARGUMENT

"AI labs are the new mass employers — Anthropic, OpenAI, Mistral."

It is true that frontier AI labs employ thousands of people and are growing. As of 2025, Anthropic and OpenAI employ in the low thousands; even on aggressive growth assumptions, the entire frontier-lab sector cannot absorb a labour-market shift measured in tens of millions of roles.

The honest comparison is not against the labour-economics of legacy big-tech but against the labour-economics of the startup ecosystem they are spawning. The frontier labs themselves are not the answer; the companies built on top of them might be — and that is precisely an entrepreneurship question.

There is a partly hopeful corollary. AI also lowers the cost of starting and running a company. Tasks that once required a team of five — sales operations, design, customer support, basic legal — can now be done by one person plus tooling. A founder in 2026 can ship more product with less capital and fewer people than a founder in 2016.

This makes the policy task more, not less, achievable. If AI compresses the cost of formation, public policy that reduces the remaining frictions — registration, customer access, talent, capital — has higher leverage than ever. **The window during which a country can decide whether AI will produce more founders or fewer employers is open and narrowing.**

Sectoral exposure: where the shock lands first

"AI exposure" is not uniform. Some sectors compress headcount rapidly; others remain stable for longer. A policy that treats the shock as a generic event will misallocate effort. The clearer view distinguishes three exposure profiles.

EXPOSURE PROFILE	EXAMPLES	POLICY POSTURE
Compressing fast	Customer support, paralegal, marketing production, software testing, mid-skill admin	Active founder creation; retraining alone insufficient
Reorganising	Software development, design, finance, research, consulting	Reskill in place; expect smaller teams per output unit
Largely stable	Healthcare delivery, construction, hospitality, regulated trades	Existing labour-market policy adequate; AI complements

The compressing-fast sectors are where the labour-market displacement is sharpest and the founder-creation imperative is strongest. They are also where AI lowers the cost of founding a company the most — the founder of an AI-native customer-support company faces a fraction of the costs of a 2016 founder in the same vertical.

> **The sectors most exposed to AI displacement are also the sectors where AI most lowers the cost of starting a competitor. That coincidence is the policy window.**

Solo-founder economics

The fastest-growing founder category in 2026 is the solo founder with AI-augmented tooling. Most countries' entrepreneurship policies do not see them: they are not in incubators, they do not raise venture capital, and they often start as side projects from full-time employment.

IMPLICATION 01

Permission matters more than capital

The constraint is not money. It is permission — social, fiscal, regulatory — to leave full-time employment and try. Entrepreneurship-leave policies for public-sector workers, simple tax structures for sole-trader-to-company transitions, and clear founder visas are all higher-leverage than capital schemes for this cohort.

IMPLICATION 02

Formation friction is the bottleneck

A solo founder is exquisitely sensitive to formation friction. Three months of paperwork is three months of personal income foregone. Formation architecture (Chapter 07) is the right policy lever for this cohort.

IMPLICATION 03

Hiring rules matter when they hire #1

The transition from solo to first-employee firm is the moment many policy frictions activate (employment law, payroll, benefits). Reducing friction here turns more solo founders into employer firms — the central goal of the entrepreneurship economy.

IMPLICATION 04

Measurement gap

Most national entrepreneurship dashboards do not track solo founders becoming employer firms. The metric is invisible because the cohort is invisible. Adding it changes both visibility and prioritisation.

Public-sector AI use as procurement signal

There is a second-order opportunity in the AI shock that is rarely named in policy debates. Governments are themselves large buyers of cognitive work. As public-sector AI adoption accelerates, procurement becomes a powerful demand-side signal for an entire startup ecosystem.

- **Public-sector AI pilots create reference customers.** A national health service or revenue authority running pilots with AI startups is a credibility signal worth more than most grant programmes.
- **The state's own AI-driven productivity gains free fiscal headroom.** Modest reinvestment of efficiency gains into entrepreneurship policy is one of the cleanest political bargains available.
- **Public-sector AI deployment exposes regulatory gaps early.** Where the state's own use of AI is unclear under existing law, regulatory clarity is in the state's own interest — which makes it more likely to happen.

None of this requires accepting strong claims about AI's long-run effects. It requires noticing that the state has both a labour-market problem and a procurement opportunity, and that the two are connected. The countries that connect them most deliberately will produce the next decade's strongest entrepreneurship economies.

The AI shock is not a thing that happens to the labour market. It is a window in which countries can decide whether they create employers or wait for them to arrive.

CHAPTER 14 · PART III

14 / 16

The Policy Stack.

A practical architecture for the entrepreneurship economy, built around the six firsts.

KEY TAKEAWAYS

01 The policy stack is built around the moments where founders stop.

02 Six firsts: permission, customer, capital, talent, regulation, recovery.

03 Each first is owned by an identifiable office, not a programme.

04 The stack does not require winner-picking. It removes avoidable obstacles.

The entrepreneurship economy needs more than inspiration. It needs a stack. The stack is built around the six moments where founders stop, with one named owner per first.

COUNTER-ARGUMENT

"Government cannot pick winners — it should stay out."

The libertarian objection is that government is structurally incapable of identifying which founders or firms will succeed and should therefore stay out. There is real evidence behind the picking-winners critique.

The six firsts are designed precisely *not* to require winner-picking. Each first is a friction-reduction layer that benefits all serious early companies. **The state's job is to remove the avoidable obstacles, not to choose between founders.**

The six firsts

01 · FIRST PERMISSION

Many founders stop before they begin because entrepreneurship feels illegible, socially risky, or institutionally discouraged. Policy creates permission through founder pathways in universities, startup-leave policies, public-sector entrepreneurship fellowships, founder stories from ordinary backgrounds, and second-career founder programmes.

02 · FIRST CUSTOMER

A company becomes real when someone pays. Government can help with startup procurement pathways, challenge-based pilots, innovation purchasing units, public-sector reference-customer schemes, and corporate-startup customer networks.

03 · FIRST CAPITAL

Early capital should be simple, fast, and connected to follow-on. Angel tax incentives, matched seed funds, regional angel activation, founder-friendly grants, university PoC funding, and pension-capital pathways where appropriate. Public capital should mobilise private capital, not replace it.

04 · FIRST TALENT

Startups need people, and early hiring is risky. Startup apprenticeships, technical co-founder matching, employee-stock-option reform, founder visas, startup internships, and talent mobility from large firms into young firms.

05 · FIRST REGULATION

Regulation should protect the public without making formation unnecessarily slow. Regulatory sandboxes, one-stop compliance guidance, fast-track approvals for low-risk pilots, modern company law, and sector-specific startup guidance in regulated industries.

06 · FIRST RECOVERY

Failure is part of the system. Modern bankruptcy rules, second-chance founder pathways, mental-health support, honest failure data, and investor norms that distinguish fraud from ordinary commercial failure. **If failure permanently destroys founders, the country wastes the learning.**

FIGURE 7 · THE SIX FIRSTS AS ARCHITECTURE

POLICY STACK

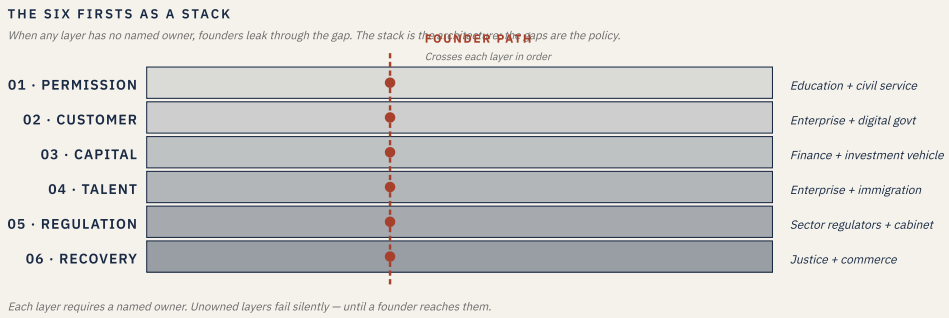

The stack at a glance

FIRST	WHAT GOES WRONG	GOVERNMENT ROLE	BETTER METRIC
Permission	The founder never starts	Legitimise and open pathways	Serious founder intents converted
Customer	The product never meets revenue	Use procurement and customer networks	First revenues and reference customers
Capital	The company dies before evidence	Mobilise fast early finance	Private capital crowded in
Talent	The founder cannot build a team	Reform equity, visas, mobility	First hires and co-founder formation
Regulation	Time kills momentum	Clarify and speed safe pathways	Time to compliant pilot
Recovery	Failure ends the founder	Enable restart and learning	Repeat-founder rate

Owning each first: who, what, with what budget

A stack with no owners is a slide. Each first needs a named accountable office, a small dedicated team, and a measurable target. The shape of each is more standardised than ministers usually realise.

FIRST	TYPICAL OWNER	TEAM SIZE	ANNUAL COST BAND
Permission	Education / civil-service ministry	3–6 people	Low — mostly policy work
Customer	Enterprise / digital-government unit	8–12 people	Medium — pilot budgets
Capital	Finance ministry + national investment vehicle	6–10 people	Programme-dependent
Talent	Enterprise + immigration ministries	4–8 people	Low — regulation + marketing
Regulation	Sector regulators + cabinet office	4–8 people	Low — coordination heavy
Recovery	Justice / commerce ministry	3–6 people	Low — legislative + cultural

The aggregate is a small, focused capability — fewer than 50 people, distributed across existing ministries, with named senior sponsors. It is structurally smaller than most countries' current entrepreneurship programmes and structurally more effective because it works on infrastructure rather than activity.

If you can only do three

Most governments cannot work all six firsts at once. They have political capital for two or three substantial reforms in a term. The order matters; not every sequence works equally well.

SEQUENCE 01

The minimum-viable trio

Permission · Customer · Capital. The three firsts that, together, move the largest cohort of would-be founders into operating companies. The other three matter, but a country that fixes these three first will produce more employers within a single term than one that spreads effort thinly across all six.

SEQUENCE 02

The technical-economy trio

Capital · Talent · Regulation. Right for countries with deep research bases and existing founder culture, where the bottleneck is institutional finance and talent mobility rather than permission. Most G7 countries are in this posture.

SEQUENCE 03

The renewal trio

Permission · Talent · Recovery. Right for countries where the bottleneck is founder formation rather than late-stage support. Smaller economies and post-shock regions usually need this sequence.

The point of the sequencing is not that the others wait forever — it is that the government can credibly defend three deep reforms in one political cycle, and add the rest in the next. **Three deep reforms outperform six shallow ones in every comparative study.**

CHAPTER 15 · PART III

What Governments Should Stop Doing.

The habits that quietly weaken entrepreneurship policy — and the substitute behaviours that change outcomes inside a single political cycle.

KEY TAKEAWAYS

01 Most governments do not need a new strategy first. They need to stop a few old behaviours.

02 The eight stops below are short, opinionated, and mostly free of fiscal cost.

03 Stopping is politically harder than starting. The defence of each old habit is usually emotional, not analytical.

04 Each stop is paired with a substitute behaviour. The reform is not absence — it is replacement.

A government working through its entrepreneurship portfolio for the first time typically reaches for new programmes. The faster intervention is the opposite. A candid audit of what to *stop*, defended publicly by the senior minister responsible, will move a country further inside twelve months than a new strategy launched at a press conference. The list below collects the eight most-replicated failure habits across the policy-evaluation literature. Each is described with what to do instead.

STOP 01

Treating entrepreneurship as a department

Entrepreneurship is affected by education, tax, procurement, immigration, regulation, universities, regional policy, finance, welfare, housing, transport, and industrial strategy. If it sits inside one small unit, that unit becomes a programme manager rather than a system architect. The unit will produce events, accelerators, and reports — none of which can move the levers the unit does not own.

What to do instead. Anchor the policy across cabinet, with a named senior sponsor in the head-of-government's office. The unit that runs operational entrepreneurship policy reports to that sponsor; its quarterly KPIs are cross-ministerial; its job is to compel returns from departments rather than to compete with them for visibility.

STOP 02

Measuring activity instead of company formation

Events held, workshops delivered, applications received, grants awarded, mentoring hours logged, incubator occupancy rates — none of these are the outcome. The outcome is serious companies created, employers formed, revenue generated, jobs created, exports produced, survival, productivity, and founder recycling. Activity metrics are the back door through which programme-thinking re-enters the system, because they are easy to measure and easy to announce.

What to do instead. Replace activity metrics with formation metrics in agency funding agreements within twelve months. Tie a meaningful share of agency funding to the cross-ministerial dashboard described in Chapter 16. Publish year-on-year change publicly. *The dashboard does the political work the strategy document cannot.*

STOP 03

Launching incubators without customers

A building full of startups with no customer pathway is entrepreneurship theatre. Founders need market access more than motivational posters or co-working desks. The most consistent finding in the incubator-evaluation literature is that the buildings make almost no difference to outcomes — the variable that does is whether the tenants can find paying customers. Most regional incubators were authorised before that variable was tested.

What to do instead. Before authorising any new incubator, ask one question: who will buy from the companies inside it within twelve months, and on what budget line? If the answer is hypothetical, the incubator is premature. Spend the budget on procurement reform (Chapter 08), corporate-customer brokerage, or sector-specific demand programmes. They are less photogenic and substantially more effective.

STOP 04

Making founders apply for everything

Application burden is a hidden tax on early companies. Every hour spent rewriting the same public-agency narrative is an hour not spent with customers, product, hiring, or fundraising. In countries with three or more major support agencies, a founder can spend several hundred hours per year applying for support — at a stage when founder time is the scarcest resource the economy has. The tax is real even when the support is generous.

What to do instead. Adopt a single-application principle: write once, share across agencies. This is administratively achievable inside one budget cycle and rarely requires legislation. Where legislation is required, it is straightforward, politically uncontroversial, and disproportionately valuable. *Treat founder time as a public good the state has so far been taxing.*

STOP 05

Designing policy around institutions rather than founders

Too much policy is built around what agencies, universities, funds, and departments can administer. Programme design follows administrative comfort: schemes look like the agency that runs them. The better question is what founders need next, in what order, with what trade-offs. The two views rarely converge by accident.

What to do instead. Document the founder journey end-to-end (Chapter 07). Walk it personally as a minister. Publish the map. Open it to founder commentary. *If a policy cannot be defended by reference to a specific point on the journey, it does not yet exist as policy — it is a budget line in search of a purpose.*

STOP 06

Confusing innovation with commercialisation

An idea, a patent, a research output, or a prototype is not yet a company. Innovation policy that stops at the prototype is solving the easier half of the journey. The expensive half — customers, pricing, sales, support, regulation, capital, talent, distribution — is where the failure rate is highest and where state friction is most damaging. Yet most innovation budgets sit on the cheap side of the curve.

What to do instead. Audit your innovation portfolio. What share of spend is on activities that produce patents and papers, versus activities that produce paying customers, hires, and exports? In most countries the ratio is dramatically skewed. Rebalancing within the existing portfolio — not new money — is the reform.

STOP 07

Punishing failure indiscriminately

Fraud should be punished. Recklessness should have consequences. But honest failure should be treated as learning. Punitive bankruptcy regimes, blanket director disqualifications, and tax treatments that follow a founder personally for years after a company closes all produce the same outcome: fewer second-time founders, and fewer first-time founders willing to take serious risks. The repeat-founder rate — the single strongest indicator of ecosystem maturity — collapses in countries that conflate failure with misconduct.

What to do instead. Modernise bankruptcy law with clear distinctions between commercial failure and misconduct. Publish honest failure data without identifying individuals. Build explicit second-chance pathways — re-incorporation, debt discharge, simple re-listing of directorships. Each is technically modest and culturally large.

STOP 08

Announcing funds without fixing demand

Capital without customers does not create a market. Public funding paired with no demand-side reform tends to produce two failure modes: capital deployed into companies that cannot reach revenue, or capital simply not deployed because the pipeline is shallower than the announcement assumed. The political win is in the launch; the policy cost is in the next three years of unimpressive results.

What to do instead. Pair every announced fund with at least one demand-side reform — procurement, export access, corporate-customer schemes, regulatory clarity. The most diagnostic question for any new entrepreneurship fund is: what demand-side reform is being announced alongside it? If none, the fund will underperform its own design.

CHAPTER 16 · PART III

16 / 16

The National Entrepreneurship Dashboard.

What countries should measure if they are serious — and how publishing the numbers changes the politics of the entrepreneurship economy.

KEY TAKEAWAYS

01 The metrics countries currently publish describe activity, not outcomes.

02 Seven measurement areas, named owners, annual publication.

03 The act of publishing the dashboard changes the political incentives.

04 Most countries already collect the data. The reform is publication, not collection.

A country serious about entrepreneurship publishes a small number of measures, on a fixed cadence, with a named senior owner accountable for each. The act of publishing changes the politics. Officials manage what is measured; ministers defend what is reported; agencies orient around what is visible.

FIGURE 8 · A NATIONAL ENTREPRENEURSHIP DASHBOARD (MOCK) DASHBOARD

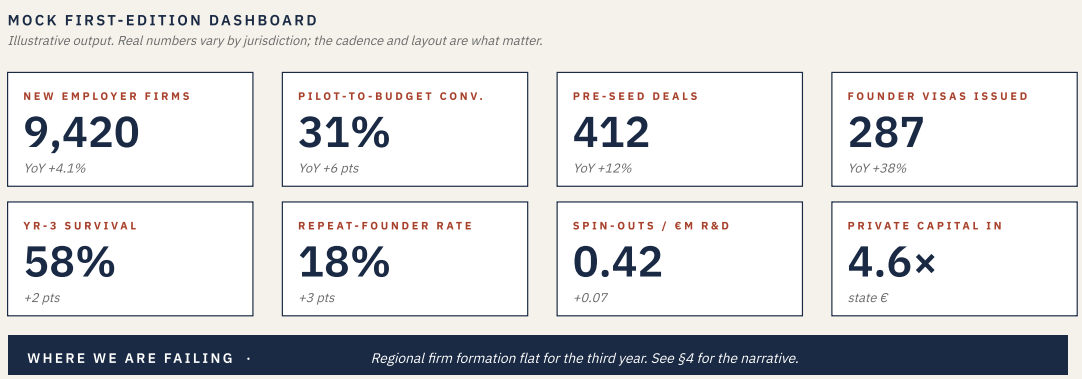

The seven measurement areas

AREA	WHAT TO MEASURE	WHY IT MATTERS
Company formation	New employer firms per year; survival to year 3; first-employee firms	Distinguishes serious companies from registrations
Founder diversity & geography	Founder demographics, regional distribution, immigrant founders	Tracks who is starting, where, with what support
Customer access	Public-sector procurement from firms <5 years old; pilot conversion rate	Whether procurement reform is actually buying
Capital formation	Pre-seed and seed deals; angel volume; private capital crowded in	Whether capital is reaching early companies
University conversion	Spin-outs per €m research spend; alumni founders; PoC funding	Whether public research becomes companies
Talent mobility	Founder visa volume; large-firm to startup transitions; co-founder formation	Whether the talent supply is responsive
Resilience	Repeat-founder rate; second-time founder share of new firms; exit recycling	Whether the system compounds

None of these measures requires invention. They require a single owner in a single agency with the authority to compel returns from departments. Most countries already collect the underlying data. **They simply do not publish it in a single place.**

Publishing the first dashboard

The first edition does not need to be perfect. It needs to exist. Five operational decisions determine whether the dashboard becomes a durable instrument or another report no one reads.

01 **Single agency, single owner.** One agency assembles the dashboard. One senior official is accountable. Returns from departments are compelled by the head of government's office, not negotiated annually.

02 **Fixed publication date.** Same week, every year. Tied to a budget statement or post-summer political cycle. Predictability creates discipline; flexible dates become indefinite delays.

03 **Baseline first, targets later.** The first edition publishes what is, not what should be. Targets are added in edition two, when the political conversation about the numbers has matured.

04 **Open data.** The dashboard is a website with downloadable data, not a PDF. Researchers, founders, and journalists can do their own analysis — which compounds the political effect.

05 **An honest "where we are failing" section.** The credibility of the dashboard rests on whether the government will publish numbers that embarrass it. If not, the dashboard becomes a marketing document and loses its political function.

Measurement pitfalls to avoid

Most dashboards that fail in their first three years fail in predictable ways. Naming the pitfalls in advance is most of the protection.

PITFALL 01

Counting registrations as companies

Companies-house data conflates serious early companies with structures created for tax, holding, or compliance reasons. The reform is to require an employee or revenue threshold before a registration counts toward the headline number.

PITFALL 02

Survival rates that flatter the system

Survival at year one is high in every country; survival at year three or five is the policy-relevant measure. Dashboards that report year-one survival are reporting noise.

PITFALL 03

Capital reported in headline totals only

Headline capital totals are distorted by a handful of large rounds. Pre-seed and seed deal counts — not total capital deployed — are the right early-stage measures. The detail is in the distribution, not the aggregate.

PITFALL 04

Failing to publish breakdowns by geography

National aggregates obscure the regional story that policymakers most need to see. Every measure should be broken down by region or NUTS-2 equivalent. The political conversation that follows is the point.

The political effect of publishing

The most under-appreciated feature of the dashboard is its second-order political effect. Numbers, once published, become a vocabulary that ministers, journalists, and opposition parties share. The conversation moves from "are we doing entrepreneurship policy?" to "what are the numbers, and what are we doing about them?" That is the substantive shift.

- **Ministers defend what is reported.** A minister whose portfolio includes one of the seven areas now has a public number to be accountable for. That number changes their incentives more than any internal target.
- **Civil servants manage to the measure.** When the dashboard exists, the agencies orient around it. Activity metrics fade from internal reporting because they no longer match the public ones.
- **Opposition has to engage substantively.** Critique becomes specific: "company formation is flat for the third year." The political debate improves.
- **Founders see themselves in the data.** When founder demographics, regional distribution, and immigrant-founder numbers are public, founders from underrepresented groups have a visible point of reference. The signal effect is larger than the data effect.

The dashboard is the cheapest, highest-leverage publication a government can produce in the entrepreneurship portfolio. It costs less than most events and lasts longer than any strategy. **The countries that publish it first will set the comparative reference for those that follow.**

IMPLEMENTATION ROADMAP

The first twelve months.

What a government can do quickly. Each three-month phase has a single deliverable and a named owner.

FIGURE · THE FIRST TWELVE MONTHS AT A GLANCE

YEAR ONE

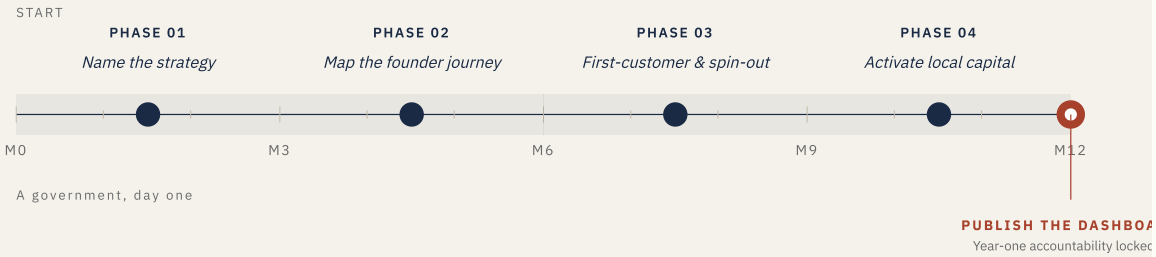

KEY TAKEAWAYS

01 A government does not need a five-year plan to start. It needs the first twelve months.

02 Each three-month phase has a single deliverable and a named owner.

03 Publishing the dashboard at month 12 locks in accountability for year two.

MONTHS 01 – 03

Name the strategy

Publish a national entrepreneurship statement. Name the senior official accountable. Map the entrepreneurship economy across the relevant departments — finance, enterprise, education, immigration, regional, university, procurement. Convene the first inter-departmental committee with founder, investor, and academic representation.

MONTHS 03 – 06

Map the founder journey

Document the founder journey end-to-end, from first idea to first exit. Identify every regulatory, administrative, financial, and informational point of friction. Publish the map. Open it to founder commentary.

MONTHS 06 – 09

Open first-customer pathways & reform spin-outs

Establish a procurement unit empowered to run pilots with firms under five years old. Reform university spin-out terms to founder-friendly defaults. Publish the new terms publicly so other universities can adopt them by reference.

MONTHS 09 – 12

Activate local capital

Reform or create the angel tax-relief scheme. Stand up the matched-fund instrument. Reform employee-stock-option taxation. Open the founder visa. Launch a startup-friendly regulatory sandbox in at least one priority regulated sector.

MONTH 12

Publish the dashboard

Publish the first national entrepreneurship dashboard. Include baseline data, not just success stories. Show the country how many employers it is creating, where they are, who is founding them, how many are surviving, how many are hiring, and where the system is failing. **Then repeat every year.**

SPECIFIC RECOMMENDATIONS BY OFFICE

Where this lands.

The argument of this paper does not land if every recommendation is addressed to "government." Below are the specific offices that need to do specific things.

OFFICE 01

Heads of Government

Name a senior cabinet member as the political owner of the entrepreneurship economy. Make the position visible and durable. Publish the entrepreneurship statement under your own signature. Defend it through one full political cycle so the strategy survives the next reshuffle.

OFFICE 02

Finance Ministers

Reform angel tax-relief and employee-stock-option taxation. Both are usually tax-cheap, politically uncontroversial, and disproportionately effective. Pair every announced startup fund with at least one demand-side reform. Resist the temptation to declare new funds without first fixing customer access.

OFFICE 03

Enterprise & Industry Ministers

Take ownership of the founder journey map. Empower a small procurement unit. Set the spin-out reform expectations for universities. Run the dashboard. Replace activity metrics with formation metrics in agency funding agreements within twelve months.

OFFICE 04

Treasury & Public-Expenditure Ministries

Allow the experimental procurement unit the legal latitude to depart from standard purchasing rules. Reform bankruptcy rules so that honest failure is recoverable. Instrument the dashboard inside the public-finance reporting cycle so its measures cannot be quietly dropped.

OFFICE 05 **Innovation & Enterprise Agencies**	Stop running entrepreneurship as a programme portfolio. Reorganise around the six firsts. Publish your conversion rates honestly. Treat repeat founders as alumni and stay in contact with them across firms.
OFFICE 06 **University Vice-Chancellors & Presidents**	Adopt the third-mission framing. Standardise founder-friendly spin-out terms. Make founder-leave a clear, written policy. Set institutional KPIs for founders produced, not patents filed.
OFFICE 07 **Regional Development Bodies**	Stop announcing buildings. Identify your anchor industries. Identify your repeat founders. Work on a fifteen-year horizon. Plan around what you have, not what you wish you had.

IN CLOSING

Conclusion.

The entrepreneurship economy is not a celebration of startups. It is a structural response to a structural change in how economies create work.

The world is becoming less predictable. Large companies are becoming more efficient. AI is changing the relationship between productivity and headcount. Global trade is becoming more political. Public finances are constrained. Regions are exposed. Young people need pathways into meaningful work. Universities need to convert knowledge into companies. Countries need more decision centres inside their own borders.

The next generation of economic policy will be judged by a single question: **did the country create enough employers?** The countries that understand this will not treat entrepreneurship as a cultural accessory or a departmental programme. They will treat it as national economic infrastructure.

The work is not glamorous. It is the patient compounding of small institutional choices — a registration form simplified, a procurement pathway opened, a visa scheme published, a failure regime modernised, a senior official named. None of these makes a front page. All of them, together, decide a decade.

Somewhere this morning, a minister is reading a restructuring announcement. Three kilometres away, a founder is writing the next pitch deck. They will never meet. They are the same story. **The only question that matters across the next ten years is whether the country builds the architecture between them, or leaves them to find each other in the dark.**

There is no neutral position. Doing nothing is a choice with a measurable cost: fewer employers, fewer first customers, fewer repeat founders, fewer chances. The cost is paid by the people the government promised to serve. The arithmetic is simple. The political will is the variable.

More serious attempts. More first customers. More repeat founders. More employers.

BACK MATTER

About the author.

THE AUTHOR

DC Cahalane

Founder · Investor · Government Advisor

Based
Cork, Ireland

Contact
me@dc.ie

DC Cahalane is a four-time founder, a venture-capital investor, and an adviser to public bodies on entrepreneurship and innovation policy. He is a Venture Partner at Sure Valley Ventures, where he invests in early-stage AI and software companies across Ireland and the UK, and Venture Capitalist-in-Residence at University College Cork. He has advised government departments and economic-development agencies on programme design, formation architecture, and ecosystem strategy.

This combination — operating, investing, and advising — is the basis of his published work on the structural design of entrepreneur ecosystems. He is the author of the forthcoming book *We're Doing This Wrong: How to Build Entrepreneur Ecosystems That Actually Work* (Cantillion Press, November 2026), and the companion white papers *Lost at Scale* (on the scale-up gap) and *The Equity Trap* (on capital design).

ENGAGEMENT

Continue the conversation.

This paper is the policy distillation of a longer argument. Three ways to engage further.

01

The book

We're Doing This Wrong: How to Build Entrepreneur Ecosystems That Actually Work will be published by Cantillion Press in November 2026. The book extends this paper into the underlying architecture: six principles of ecosystem design, country case studies, founder testimony, and an implementation manual.

02

Country editions

Localised editions of this paper — Ireland, the United Kingdom, the European Union, and a small number of other markets — are in preparation. Each will include country-specific data, named programmes, founder stories, and a country-tailored policy roadmap. To request a country edition or to participate in one, contact **me@dc.ie**.

03

Endorsement & circulation

This paper is intended to be circulated, quoted, and contested. Reproduction is permitted with attribution. Endorsement signatures from founders, investors, academics, and former officials are being collected for the next edition. To add yours, contact **me@dc.ie**.

Companion white papers

- **Lost at Scale** — on the scale-up gap and what happens to companies that get past first capital but stall at first €5 million in revenue.
- **The Equity Trap** — on capital design, dilution, and the founder economics of poorly structured early rounds.

BIBLIOGRAPHY

References & further reading.

Citations are footnoted in the body. The list below collects principal sources and points to additional reading. URLs are stable as of date of publication.

01 **World Economic Forum.** *Future of Jobs Report 2025.* Geneva, January 2025. weforum.org/reports/the-future-of-jobs-report-2025

02 **Cazzaniga, M. et al.** "Gen-AI: Artificial Intelligence and the Future of Work." IMF Staff Discussion Note SDN/2024/001, January 2024. imf.org/en/Publications/Staff-Discussion-Notes

03 **OECD.** *OECD Economic Outlook,* Volume 2025 Issue 1. Paris, June 2025. oecd.org/economic-outlook

04 **World Bank.** *Global Economic Prospects,* June 2025. Washington DC. worldbank.org/en/publication/global-economic-prospects

05 **International Monetary Fund.** *World Economic Outlook: Policy Pivot, Rising Threats.* October 2024 (updated April 2025). imf.org/en/Publications/WEO

06 **OECD.** *OECD SME and Entrepreneurship Outlook 2024.* Paris, OECD Publishing, 2024.

07 **OECD.** "Understanding Firm Growth: Helping SMEs Scale Up." OECD Studies on SMEs and Entrepreneurship, 2021.

08 **OECD.** "The Contribution of Scalers to Job Creation." OECD SME and Entrepreneurship Papers, 2024.

09 **World Bank.** "Small and Medium Enterprises (SMEs) Finance." Topic Brief, 2024.

10 **Haltiwanger, J., Jarmin, R., & Miranda, J.** "Who Creates Jobs? Small versus Large versus Young." *Review of Economics and Statistics* 95(2), 2013. Updated: Kauffman Foundation, "The Importance of Young Firms," 2024.

11 **Endeavor Insight.** "How Did Silicon Valley Become Silicon Valley? Three Surprising Lessons." 2014.

12 **Estonia e-Residency.** *e-Residency Annual Report 2024.* Tallinn, 2024. e-resident.gov.ee

13 **Erlich, Y.** "The Yozma Programme — Policy and Success Factors." Yozma Group archive, 1999. Senor, D. and Singer, S. *Start-Up Nation.* 2009.

14 **HM Revenue & Customs.** *Enterprise Investment Scheme, Seed Enterprise Investment Scheme and Social Investment Tax Relief: Statistics.* London, 2024.

15 **US Small Business Administration.** *SBIR/STTR Annual Report.* Most recent edition. sbir.gov

16 **Hausmann, R. & Hidalgo, C.** *The Atlas of Economic Complexity.* MIT Press, 2014.

17 **OECD.** *OECD Productivity Statistics.* Ongoing series.

18 **Anthropic, OpenAI, Mistral, Stripe, Klarna.** Company employee disclosures, financial filings, and public statements, 2023–2025.

19 **Mazzucato, M.** *The Entrepreneurial State.* 2013. Hwang, V. and Horowitt, G. *The Rainforest.* 2012. Feld, B. *Startup Communities.* 2nd ed., 2020.

20 **Eurostat.** "Self-employed persons by sex, age and citizenship of country of residence." 2025 release, dataset lfsa_egais. ec.europa.eu/eurostat. Surfaced and analysed in McWilliams, D. "Ireland Is Killing Its Entrepreneurs." *The David McWilliams Podcast,* 6 June 2026.

21 **Schumpeter, J. A.** *The Theory of Economic Development.* Harvard University Press, 1934. *Capitalism, Socialism and Democracy.* Harper & Brothers, 1942.

22 **OECD city-level cost-of-living and rent indices,** 2024–2025; Numbeo Cost of Living Index, accessed 2025; national rental statistics for Dublin, London, Paris, Amsterdam, and Stockholm. The Dublin worked example draws on McWilliams, D. "Ireland Is Killing Its Entrepreneurs." *The David McWilliams Podcast,* June 2026.

DEFINED TERMS

Glossary.

Defined here for non-specialist readers. Each definition is the working definition used in this paper, not necessarily the canonical academic one.

Spin-out.

A company formed from research, intellectual property, or knowledge generated inside a university or public institution.

Scaler.

An SME (typically 10–249 employees) growing employment or turnover at above-average rates over a sustained period; the OECD's standard definition is used.

Foreign direct investment (FDI).

Investment by a firm headquartered in one country into operations in another. Distinguished here between strategic-mandate FDI (decision-centre, R&D, full responsibility) and satellite FDI (cost-driven branch operations).

Formation architecture.

The legal, administrative, digital, financial, and regulatory infrastructure that determines how easy or hard it is to start a company. Distinct from programmes that support founders after formation.

First customer.

The earliest paying buyer for a company's product, in the policy sense used here including public-sector first customers via procurement.

DEFINED TERMS · CONTINUED

Glossary *(cont.)*

Six firsts.

The six policy domains where founders most commonly stop: first permission, first customer, first capital, first talent, first regulation, first recovery.

Founder visa.

An immigration pathway designed specifically for founders intending to start companies, distinct from skilled-worker permits.

Repeat founder.

An entrepreneur founding their second or later company, with the failures and successes of prior ventures behind them.

Exit.

The sale of a company through acquisition or initial public offering, or in some uses the wind-down of a venture; the founder-economic event.

Dashboard.

In this paper, the proposed national entrepreneurship dashboard: a small set of measures, published annually, with a single named senior owner per area.

A CLOSING LINE

A country that does not create employers will eventually run out of employment policy.

More serious attempts. More first customers. More repeat founders. More employers.

FORTHCOMING · NOVEMBER 2026

We're Doing This Wrong

How to Build Entrepreneur Ecosystems That Actually Work

IMPRINT

CANTILLION PRESS · CORK, IRELAND

CONTACT: ME@DC.IE