

CHARLES UNIVERSITY
Faculty of Social Sciences
Centre for Economic Research and Graduate Education

Master's Thesis

CHARLES UNIVERSITY
Faculty of Social Sciences
Centre for Economic Research and Graduate Education

**Why ICT Matters Less in Europe:
Financial Structure and Organizational
Investment**

Master's Thesis

Author: Petr Rusnok
Study programme: Economic Research
Supervisor: Prof. Ing. Michal Kejak, M.A., CSc.
Year of the defence: 2026

Declaration

1. I hereby declare that I have compiled this thesis using the listed literature and resources only.
2. I hereby declare that my thesis has not been used to gain any other academic title.
3. I fully agree to my work being used for study and scientific purposes.

In Prague on May 5, 2026

Petr Rusnok

Bibliographic Record

RUSNOK, Petr. *Why ICT Matters Less in Europe: Financial Structure and Organizational Investment*. Praha, 2026. Master's thesis. Charles University, Center for Economic Research and Graduate Education – Economics Institute, A joint workplace of Charles University and the Economics Institute of the Czech Academy of Sciences. Supervisor Prof. Ing. Michal Kejak, M.A., CSc.

Acknowledgment

I owe my deepest thanks to my supervisor, Michal Kejak, for steady guidance throughout this project. I am also grateful to Stephanie Ettmeier and Alexander Popov, each of whom commented on the thesis at various stages, and to the CERGE-EI Macro Reading Group for a careful discussion of an earlier draft. The empirical analysis rests on open data from EU KLEMS & INTANProd at the LUISS Lab on European Economics, the Bank for International Settlements credit statistics, and the World Bank's Global Financial Development Database. I thank those institutions for making the data publicly available. The replication package, including all data extraction and regression scripts, is available at <https://github.com/Lemkei/ict-europe-replication-final>.

Abstract

After 1995, frontier economies with developed equity markets pulled ahead in productivity while economies dependent on bank lending fell behind. I argue financial structure shapes which firms take on a class of investment I call *risky innovation*, where the firm's expected cash flow rises but so does risk. A risk averse owner financed by debt absorbs the worst states herself. Under equity the investor shares every state, and the partial insurance lets her take on riskier projects with higher productivity gains. I prove that the optimal contract for such an owner is a call option, the mirror image of debt. From these theoretical foundations I test three predictions. First, I revisit Levine's canonical result that financial structure does not matter for aggregate growth ($\beta = +0.019$, $p = 0.97$). I extend it into a two period model and show that the canonical result misses transitional dynamics to a technology shock ($\beta = +5.56$, $p < 0.0001$ post 1995). For H2 and H3 I use EU KLEMS & INTANProd data on thirteen Western economies, twelve of them in a Rajan Zingales style identification with the United States serving as a benchmark country. *H2* relates the labour productivity contribution of intangible assets and ICT to financial structure across European sectors and finds the pattern the micro foundations predict. *H3* tests the interaction of the intangible complement with IT adoption and shows capital reallocation toward IT in equity heavy countries ($\beta = +19.42$), with bank credit and GDP per capita placebos loading wrong signed.

Keywords

risky innovation, ICT, financial structure, organizational capital, intangible investment, EU KLEMS, equity finance, productivity, European economy

JEL Classification

E22, G32, O33, O47

Title

Why ICT Matters Less in Europe: Financial Structure and Organizational Investment

Abstrakt

Vyspělé ekonomiky, kde akciové trhy hrají větší roli než banky, začaly růst rychleji, zatímco ekonomiky závislé na bankách zpomalily. Důvodem je, že banky nejsou optimální pro rizikovou inovaci, kterou definuji jako firemní investici, která sice zvýší očekávané zisky, ale i riziko. Vlastník firmy s averzí k riziku pro takové projekty dluh nevezme, protože ty nejhorší stavy disproporčně ovlivní jeho užitek. Equity investor bere jen konstantní podíl ze zisku, takže částečně pojistí vlastníka firmy před rizikem, a to umožňuje financování rizikovějších projektů s většími zisky. Dokazuji, že optimální kontrakt pro takového vlastníka je kupní opce, zrcadlový obraz dluhu a tedy v podstatě jeho opak. Z těchto teoretických základů testuji tři predikce. Nejprve reviduji Levineovo obecně uznávané zjištění, že finanční struktura nemá vliv na ekonomický růst ($\beta = +0,019$, $p = 0,97$). Rozšiřuji jeho dataset a replikuji jeho klíčové zjištění, ale zároveň ukazují, že opomíjí tranzitní dynamiku při technologickém šoku ($\beta = +5,56$, $p < 0,0001$ po roce 1995). Pro H2 a H3 používám data EU KLEMS & INTANProd pro třináct západních ekonomik, dvanáct z nich v identifikační strategii ve stylu Rajana-Zingales s USA jako referenční zemí. H2 dává do souvislosti příspěvek nehmotných aktiv a ICT k produktivitě práce s finanční strukturou napříč evropskými odvětvími a nachází zjištění, které můj teoretický rámec předpovídá. H3 testuje interakci nehmotného komplementu s adopcí IT a ukazuje realokaci kapitálu směrem k IT v akciově orientovaných ekonomikách ($\beta = +19,42$), přičemž bankovní úvěr a HDP na obyvatele jako placebo nenacházejí signifikantní výsledek.

Klíčová slova

riziková inovace, ICT, finanční struktura, organizační kapitál, nehmotné investice, EU KLEMS, akciové financování, produktivita, evropská ekonomika

JEL kódy

E22, G32, O33, O47

Název práce

Proč ICT v Evropě není produktivní: finanční struktura a investice do organizačního kapitálu

Contents

1	Introduction	1
2	Innovation theory and risky innovation	4
3	The equity-debt wedge for risky innovation	13
4	What labour productivity measures, and what it does not	23
4.1	Labour productivity versus income per capita	24
4.2	The Jorgenson decomposition	24
4.3	What survives, what does not	26
4.4	Implications for the empirical strategy	27
5	Empirical results	28
5.1	Specification and sample	30
5.2	Extending Levine’s dataset	33
5.3	Country-level patterns	36
5.4	Equity, the intangible-ICT bundle, and labour-productivity growth . .	38
5.5	Capital reallocation: the input side	41
5.6	Robustness	43
6	Limitations	47
7	Conclusion	50

1 Introduction

This thesis connects risk, innovation, and financial structure. Endogenous growth theory puts innovative firms at the center stage of economic growth. If anything happens to them, the growth trajectories of whole economies should change. If financial structure matters for these firms, as recent work suggests (Brown et al., 2017; De Haas & Popov, 2023; Draghi, 2024; Garcia-Macia, 2017), it should matter for growth as well. Yet it famously does not (Levine, 2002). The reconciliation is timing. Levine’s panel ends in 1995, before the ICT revolution and the rise of intangibles began showing up in aggregate growth. Equity heavy economies exploited the technological shock ICT and intangibles created better than bank heavy ones, disrupting existing growth trajectories.




FIGURE 1. **Labour productivity relative to the United States, 1950–2024.** GDP per hour, US = 100. EU-12 is the unweighted mean of the twelve Western European economies in the sample. Source: BCL Long-Term Productivity Database v2.7 (Bergeaud et al., 2016).

The scale of divergence the ICT and intangible shock created is evident in Figure 1. The graph is largely a story of the US pulling ahead through its adoption of ICT (Bergeaud, 2024). If equity share were the only factor, equity heavy countries like the UK should

be doing better. To isolate the relationship from other factors I split the data into two periods, account for country and period effects, and find a strong association between financial structure and productivity in Section 5.2. Levine's result then isn't wrong, but it misses important transitional dynamics.

That leaves us with a deeper problem. If innovative firms drive growth both before and after 1995, why did their relationship with financial structure change? I argue the answer is risk. However, to make this argument I first need to deal with the following issues.

Strictly speaking, risky innovation is an oxymoron. Standard mathematical definitions of an increase in risk (Rothschild & Stiglitz, 1970) preserve the mean. A mean preserving spread cannot raise expected payoff. Innovation, by definition, raises it. An innovative project pays above the safe rate in expectation, otherwise it is mere investment. The resolution is to restrict attention to a subset of concave utility functions for which the project strictly lowers expected utility despite raising mean payoff. Within that class, risky innovation is a coherent object.

Innovation is not inherently risky. The distinction between project-level and firm-level risk matters. A large firm running a diversified portfolio of research bets makes its cashflow distribution safer, an observation Schumpeter made about innovation in mature firms (Schumpeter, 1942). A robust definition compares the status quo cashflow distribution to the post innovation distribution. I take risky innovation to be a convex, monotonically increasing transformation whose costs are large enough to put the firm at risk in the worst states. This is a subset of risky innovation, but it is the subset that matters here.

Information technology adoption is risky innovation even at large firms. To turn hardware into productivity, firms must commit to large complementary intangible investments such as organizational restructuring, retraining, and software customization (Bloom et al., 2012; Bresnahan et al., 2002; Brynjolfsson & Hitt, 2003; Brynjolfsson et al.,

2021). The whole innovation project then becomes very costly. Crucially, it cannot be isolated to a few products the firm sells. The processes that change are core to the whole firm, and if it goes wrong, the entire product portfolio is affected. Intangibles themselves carry risk properties Section 2 unpacks. Pre-ICT frontier innovation ran largely through safer tangible investment. The switch to intangibles cuts deeper to the core of the firm than buying a machine ever would.

The last theoretical piece is showing that credit fails for risky innovation. The empirical pattern is clear, since venture capital funds are equity investors for a reason. In standard corporate finance, however, the result does not fall out cleanly. When the entrepreneur can adjust project riskiness and an increase in risk raises expected payoff, debt prices itself break even and the entrepreneur is indifferent (Biais & Casamatta, 1999). I resolve this by making the entrepreneur risk averse. Equity then plays a dual role, providing funds and insuring against bad states, an observation Schumpeter also made about innovation finance (Schumpeter, 1934). Credit fails because it leaves the entrepreneur worse off in bad states and better off in good states, a mean preserving spread of consumption for an agent who would prefer it flat.

The chain is now visible. The post-95 technology shock came in the form of risky innovation. Risky innovation cannot be funded by debt, because debt makes the risk-averse entrepreneur worse off in the bad states. Equity can. Equity heavy economies therefore captured the shock. Bank heavy economies did not.

To put the theory into test, I distinguish economic competencies from innovative properties. Both classes of intangibles play a role in risky innovation, but innovative properties at large firms can take the form of safe, diversified research portfolios. The framework predicts that economic competencies depend on financial structure while innovative properties do not. Conditional on being funded, the innovative properties equity heavy economies select should disproportionately raise labour productivity. I test this with a Rajan-Zingales style identification using US benchmark growth in

labour productivity contribution as the sectoral exposure, and run the same exercise for ICT and economic competencies.

For ICT adoption specifically, I show that equity heavy economies reallocated capital toward IT-intensive sectors faster than bank heavy ones. The same interaction built on bank credit comes out wrong-signed, ruling out a generic credit-deepening story.

2 Innovation theory and risky innovation

Innovation in existing theory

Innovation enters most economic models as the residual the model cannot explain. Firms do not report it on their balance sheets, and it is not always clear that firms themselves know how much of it they do. Giving innovation a statistical structure, including a notion of risk, is therefore a tricky exercise. I start with Schumpeter, who supplies the canonical definition of innovation and the canonical statement on how innovation relates to risk, and revisit the question first in his early work and then in the later one.

For Schumpeter (1934), innovation is the carrying out of a new combination. It is not limited to a new product. The innovator can introduce a new method of production, open a new market, conquer a new source of supply, or reorganise an industry. In the early framework it is the entrepreneur who, often from outside the existing economic order, brings the new combination into production by founding a new firm. The entrepreneur is not a researcher: she does not, as a rule, invent the new technology. Her contribution is bringing it to fruition. Schumpeter draws the distinction sharply: “As long as they are not carried into practice, inventions are economically irrelevant” (Schumpeter, 1934, ch. II).

For Schumpeter, the entrepreneur is not the bearer of risk; the investor is. “The en-

entrepreneur is never the risk bearer. In our example this is quite clear. The one who gives credit comes to grief if the undertaking fails” (Schumpeter, 1934, ch. II). The investor lends to many entrepreneurs and absorbs failures across the portfolio. The entrepreneur cannot diversify in the same way: her firm is the project. Optimal innovation therefore requires someone willing to bear the residual risk, and that is the function of the capitalist. Section 3 formalises the optimal financial contract for the risk-averse entrepreneur, but the structural point is already in Schumpeter: the financing investor is constitutive of innovation, not an afterthought.

In his later works Schumpeter (1942, ch. VIII) looks at how innovation works in larger firms, and, crucially for us, he finds that innovation loses its risky character. Mature capitalism internalises the entrepreneurial function inside the large established firm. “The first thing a modern concern does as soon as it feels that it can afford it is to establish a research department every member of which knows that his bread and butter depends on his success in devising improvements.” Innovation becomes the routine work of teams of trained specialists. The firm funds the research department from internal cash flow. External finance disappears from the picture, and the risk-bearing investor of the 1934 framework disappears with it.

Innovation is not inherently risky. Risk as a concept does not make sense outside a specific payoff function. The same project can be risky for one agent and safe for another, exactly as we saw in the contrast Schumpeter drew between large and small firms: in 1942 the routine R&D of an established firm is safe in a way that the same project would not be in the hands of a 1934 entrepreneur. Large firms can take on risky innovation too. The question is whether the project shifts the firm’s cash flow distribution in a way that constitutes risk in the formal sense.

The framework of risk

Defining “riskier” formally requires comparing the integrals of the cumulative-distribution functions. Let X denote the firm’s stochastic future cash flow, a random variable with cumulative distribution F_X and finite mean μ . Rothschild and Stiglitz (1970) characterise Y as riskier than X if

$$\int_{-\infty}^t F_Y(s) ds \geq \int_{-\infty}^t F_X(s) ds \quad \text{for all } t, \quad (1)$$

with equality at $t = \infty$, which forces equal means.

The integral characterisation and the mean-preserving-spread version are equivalent. For our purposes we need to relax the definition, and the integral form is the more convenient starting point.

Innovation does not preserve the mean. The relaxation is impossible without having a specific concave payoff function in mind. Here, that function is the debt contract. One distribution is riskier than another when a credit investor would have to raise the interest rate to compensate for the additional expected shortfall. Formally, with $\mathcal{D} = [\underline{D}, \bar{D}]$ denoting the range of face values that a bank could competitively offer and a sane owner would accept,

$$\int_{-\infty}^{\bar{D}} F_{X_1}(s) ds > \int_{-\infty}^{\bar{D}} F_{X_0}(s) ds \quad \text{for every } \bar{D} \in \mathcal{D}. \quad (2)$$

The integrals are the expected shortfall below \bar{D} under each distribution. The face value \bar{D} is loosely defined here; Section 3 pins it down formally. The definition of risk above is sufficient for what follows.

Safe and risky research project

As we saw in Schumpeter's later work, the same project can be both safe and risky depending on the firm that takes it on. It is not enough to look at the statistical properties of a project in isolation; we need to look at what it does to the firm's cash flow distribution.

First consider a research project that ends in a patent the firm licenses out rather than uses itself. The project costs c . In both success and failure the firm's cash flow falls by c relative to baseline. If the project succeeds, the firm earns an independent payoff Y on top, and the post-project cash flow becomes $X - c + Y$. For a credit contract with face value D this is not risky as long as $X - c > D$ in every realisation of X : the bank is paid in full whether the project succeeds or fails. For the firm, the variance increases by the project's variance, but as long as c is small relative to baseline the additional risk is modest.

Next consider a project that improves the firm's existing product directly: in failure the cash flow is $X - c$, in success it is $sX - c$ for some scalar $s > 1$. The bank's claim is unchanged from the previous case as long as $X - c > D$. The firm's variance is higher than baseline because the success state amplifies the entire X distribution. In both this case and the licensing case, the more independent products and independent research projects the firm has, the more the law of large numbers concentrates the total around its mean. This replicates Schumpeter's 1942 observation that mature large firms internalise innovation routinely without taking on substantial risk.

A project becomes risky for a large firm only when it affects all products at once. ICT adoption is the canonical example: the firm restructures its workflows, retrain staff, redesigns the organisation, and every product line is touched. Diversification across products is no longer available. Two further conditions tighten the picture. The project's effect on cash flow has to be both increasing and convex in the underlying state, so

that it amplifies good outcomes more than bad ones rather than scaling everything by a constant. The cost has to be substantial enough that the whole business model is at stake. These three features (economy-wide reach within the firm, increasing convex amplification, and substantial cost) are what make risky innovation possible for large firms.

The implication for R&D spending follows. Existing firms can fund many kinds of R&D projects whose effect on cash flow is safe in the sense above, so bank-heavy economies can spend heavily on R&D and the cross-country relationship between R&D intensity and financial structure is not strong. The composition of R&D differs across countries even when the total does not. If the riskier R&D projects are systematically more productive, equity-heavy economies should grow faster in R&D-intensive sectors, exactly the pattern Brown et al. (2017) documents. Aggregate R&D intensity smooths over the composition difference, which is why R&D serves as a clean placebo for the financing channel in our empirical strategy: equity finance raises composition-shifted productive R&D, but does not raise total R&D.

The relation to the endogenous-growth literature is complicated. In this thesis we model the decision of a single firm choosing how risky its innovation project should be. Most risk in endogenous growth models arises from creative destruction at the level of incumbent firms rather than from the project decisions themselves. Aghion and Howitt (1992) treats successful R&D events as Poisson arrivals that lift aggregate productivity by a multiplicative factor $\gamma > 1$. Each new innovation comes from an outside research firm, and the previous incumbent's rents are erased on arrival. The risk in the model is the timing of displacement, borne passively by the incumbent and controlled (in expectation) by aggregate research effort. Our framework gives the incumbent an active choice over the risk profile of its own project, which has no clean analog in this kind of model. There may exist Schumpeterian growth models that describe at the economy-wide level something similar to what we describe at the firm level, but that is beyond the scope of this thesis.

Are intangibles riskier?

Consider a firm that has the option to invest a fixed amount in either a tangible or an intangible asset, and that asset is the firm's only one. Does the asset choice affect the credit contract the bank would offer?

Likely yes, although the case requires some care. Kermani and Ma (2023) hand-collect liquidation analyses from Chapter 11 filings of US non-financial firms across 48 industries. Property, plant, and equipment recover about 35 percent of replacement cost on average; non-goodwill book intangibles recover a comparable 37 percent; goodwill, which captures organisational capital and acquisition premiums, recovers essentially zero. The gap between tangibles and intangibles concentrates in the embedded portion. For an asset with lower value in some states to be a sensible investment, it must have higher value in other states, otherwise no firm would invest in it. Intangibles earn the compensation through going-concern productivity gains, which typically accrue above the bank's face value \bar{D} , in the region the firm captures rather than the bank. The bank consequently raises the interest rate to break even. In a complete-information model with a single investor the expected cost of capital is the same regardless of the asset; Garcia-Macia (2017) formalises this in a structural collateral model with recovery rates of 0.65 for tangibles and 0.26 for intangibles. Our framework instead treats liquidation value as risk in the sense of the previous subsection: when the firm is funded by both equity and credit and the credit investor is senior, lower liquidation value falls disproportionately on the credit investor, even when the total cost of capital is held constant.

The picture worsens with depreciation. Bontadini et al. (2023) estimate the annual depreciation rate for organisational capital at roughly 40 percent, faster than software (around 30 percent), R&D (around 20 percent), machinery (13 to 19 percent), or structures (1 to 3 percent). ICT hardware, although technically tangible, depreciates at rates close to intangibles: roughly 30 percent per year for IT equipment, around 12 percent

for CT equipment. For a fast-depreciating asset to be a sensible investment, the firm must capture compensating cash flow during the asset's productive life. As long as that compensating cash flow does not accrue solely below \bar{D} , the bank does not see it. The asset depreciates, the bank's collateral falls with it, and the compensation lies in a region the bank cannot claim.

Lower liquidation value also reduces collateralisability. Lian and Ma (2021) document that asset-based lending accounts for only about 20 percent of US non-financial corporate debt by value while cash-flow-based lending accounts for the remaining 80 percent; among large US firms the median share of cash-flow-based debt reaches 85 percent. Continental Europe leans the other way: with bankruptcy regimes that favour liquidation over reorganisation, claims over specific assets are more important, and lending tilts toward asset-based contracts. The shift in legal infrastructure restricts the menu of credit contracts available to firms with mostly intangible assets. The restriction does not cleanly imply higher project risk, but it does mean that European credit investors face a contract space less suited to firms with low liquidation value.

Both channels work in the same direction. Under the assumption that some portion of the project's compensating gains accrues above \bar{D} , investment in assets with low liquidation value and fast depreciation is riskier than investment in tangibles with high liquidation value and slow depreciation.

Rise of intangibles and ICT revolution

This theoretical asymmetry became consequential in the 1990s and 2000s, when two parallel transformations of the frontier economy tested it: the diffusion of information and communications technology and the rise of intangible capital. The two events were not independent. The productive deployment of ICT requires firm-specific reorganisation, training, software customisation, and the accumulation of organisational know-how. The complement to ICT hardware is intangible by construction. As a result,

the post-ICT frontier requires firms to fund innovation through assets riskier than the pre-ICT tangible base of factories, machinery, and inventory. This was an exogenous technological shock to existing economies. Countries with deeper equity markets could finance the riskier complement; bank-heavy economies could not.

The complementarity between ICT and organisational capital is documented at the firm level. Bresnahan et al. (2002) establish that ICT investment, organisational change, and skilled labour move together within firms. Bloom et al. (2012) show that US multinationals operating in Europe extract roughly four times the productivity per ICT unit through superior management practices. Brynjolfsson et al. (2021) document a productivity J-curve: intangible complements accumulate before measured output catches up, leaving early-period statistics to understate the eventual gain.

The era also produced the modern superstar firm. The economic logic is to invest heavily in the intangible complement, build a defensible moat that competitors cannot easily replicate, and then extract a steady flow of rents. De Ridder (2024) estimates structurally that the rise of intangible-intensive firms raised aggregate markups by 14.6 percentage points in the United States and 5.6 percentage points in France.

The composition of innovative investment matters more than its absolute level. European ICT investment intensities through the 1990s and 2000s were broadly comparable to the United States; what differed were the productivity gains. Inklaar et al. (2008) document the Danish paradox: Danish market services invested 6.2 percent of value added in ICT during 1995–2004, above the US share of 5.9 percent, yet Danish multi-factor productivity in market services fell by 0.7 percent per year while the US rose by 1.3 percent. The difference is not in ICT spending but in the organisational and managerial restructuring that accompanies it. Bank-heavy economies undertook the safer adoption: they bought the hardware but did not commit to the lumpy, organisationally disruptive complement. Their productivity gains concentrated in the middle of the firm distribution rather than the superstar tail.

The asymmetry compounds. Once superstar firms are established with their intangible moats in place, subsequent technological shocks reinforce the existing distribution rather than levelling it. Cloud computing in the 2010s was largely deployed by the established big-tech incumbents that had built ICT-era organisational capital. Artificial intelligence in the 2020s is following the same pattern. Each new general-purpose technology arrives at firms whose ability to undertake risky organisational restructuring has already been shaped by the financing environment that funded the previous generation.

Risky innovation: formal definition

Section 3 works with a parametric form that captures the three conditions identified above: economy-wide reach within the firm, increasing convex amplification, and substantial cost. The terminal cash flow under risky innovation is

$$X_1(\gamma, \tilde{\varepsilon}) = x_0 \cdot \tilde{\varepsilon} \cdot (1 + \Phi(z))^\gamma, \quad (3)$$

where $\tilde{\varepsilon} = \exp(\sigma z)$ with $z \sim \mathcal{N}(0, 1)$ is a multiplicative log-normal industry shock, Φ is the standard normal CDF, σ is industry log-volatility, x_0 is the firm's pre-investment scale, and $\gamma \geq 0$ is the firm's amplification choice. The multiplier $(1 + \Phi(z))^\gamma$ acts on every realisation of $\tilde{\varepsilon}$ (economy-wide reach), is convex and strictly increasing in $\Phi(z)$ for $\gamma > 0$ (amplification of the upper tail above the lower), and is paired in Section 3 with the cost function $I(\gamma) = \alpha\gamma^2 A$ (substantial cost rising in γ). The form is the canonical reduced expression of risky innovation in the sense of the relaxed risk inequality. Section 3 pairs this technology with a CRRA-2 owner who chooses γ optimally under different contract forms.

3 The equity-debt wedge for risky innovation

Credit can fund a risky innovation at the same expected cost as equity. As risk rises the lender raises the interest rate to compensate for losses on bad outcomes, leaving the owner's expected payment unchanged. The puzzle is not why banks cannot fund risky innovation, but why owners do not borrow. The answer is a second concavity. The debt payoff is concave in firm cash flow, and so is the owner's utility. Combining the two makes debt close to the worst possible contract, because both the debt payoff and risk-averse utility put the most weight on the bottom of the distribution.

Under debt, bad states leave the owner with nothing once the lender is paid. Concave utility makes her want exactly the opposite: insurance against bad outcomes, a constant consumption in every state. Equity provides partial insurance because the investor shares in the loss. Debt is worse than self-financing. With most of her wealth and her future employment locked in one firm, she prefers not to innovate even when the expected gain is several times larger.

Setup

Consider a business owner who holds her entire wealth A inside an existing firm. She has a one-shot opportunity to undertake a risky innovation in the sense of equation (3) and chooses how risky to make it through $\gamma \geq 0$. The owner draws a fixed wage $w = rA$ in both periods, paid before any investor claim, so she is not solely dependent on the firm's outcome. She is free to design her financial contract: any function $\delta : \mathbb{R}_+ \rightarrow [0, 1]$ that pays the investor a share $\delta(X_1)$ of the realised project payoff, subject to the investor breaking even in expectation.

She solves

$$\max_{\gamma, \delta} u(c_0) + \beta \mathbb{E}[u(c_1)] \quad (4)$$

subject to

$$\begin{aligned} \mathbb{E}[\delta(X_1)X_1] &= (1+r)I(\gamma) && \text{[investor break-even]} && (5) \\ 0 \leq \delta(X) \leq 1 &\quad \forall X && \text{[limited liability]} \end{aligned}$$

with $c_0 = w$, $c_1 = w + (1 - \delta(X_1))X_1$, $w = rA$, $I(\gamma) = \alpha \gamma^2 A$, and CRRA-2 utility $u(c) = -1/c$ with $\beta(1+r) = 1$. The CRRA-2 specification places the owner in the empirically defensible range of risk aversion for entrepreneurs (Halek & Eisenhauer, 2001). The terminal cash flow $X_1(\gamma, \tilde{\varepsilon})$ is the parametric form of equation (3), with $\tilde{\varepsilon} = \exp(\sigma z)$ a multiplicative log-normal industry shock and γ the owner's amplification choice. Pre-investment scale is normalised so the passive baseline $\gamma = 1$ is fairly priced: $\mathbb{E}[X_1(1, \tilde{\varepsilon})] = (1+r)A$.

The remainder of Section 3 refers back to these primitives without re-listing them. The symbols introduced here ($\gamma, \sigma, \alpha, A, r, \beta, z, \tilde{\varepsilon}, \Phi, X_1, I, w, c_0, c_1, u, \delta, \lambda$) carry their meanings to the end of the section.

The optimal contract is convex

Hold γ fixed. The owner chooses $\delta : \mathbb{R}_+ \rightarrow [0, 1]$ to maximise her expected utility subject to the investor's break-even condition. Since $c_0 = w$ is constant in δ , the choice problem reduces to maximising $\beta \mathbb{E}[u(c_1)]$ alone. The Lagrangian is

$$\mathcal{L}(\delta, \lambda) = \beta \mathbb{E}[u(w + (1 - \delta(X_1))X_1)] + \lambda (\mathbb{E}[\delta(X_1)X_1] - (1+r)I(\gamma)), \quad (6)$$

with multiplier $\lambda > 0$ pricing a unit of break-even slack. Limited liability bounds $\delta(X) \in [0, 1]$ at every realisation X .

PROPOSITION 1 (Optimal contract is convex). *The investor's payoff under the optimal contract is*

$$\delta^*(X)X = \max(X - K, 0), \quad (7)$$

a convex function of the realised project payoff X . Equivalently, the optimal contract is

$$\delta^*(X) = \max\left(0, 1 - \frac{K}{X}\right), \quad (8)$$

with threshold $K = \bar{c} - w$, where \bar{c} is the unique constant solving the break-even condition $\mathbb{E}[\max(X_1 - K, 0)] = (1 + r)I(\gamma)$.

PROOF. The proof proceeds in five steps.

Step 1: Pointwise FOC. Differentiate the integrand of \mathcal{L} in equation (6) with respect to $\delta(X)$ at a single realisation X . The chain rule gives $\partial c_1 / \partial \delta(X) = -X$, so

$$\frac{\partial \mathcal{L}}{\partial \delta(X)} = -\beta u'(c_1(X))X + \lambda X = 0.$$

Dividing by $X > 0$:

$$u'(c_1(X)) = \frac{\lambda}{\beta}. \quad (9)$$

Step 2: c_1 is flat on the interior. The right-hand side of equation (9) does not depend on X . With CRRA-2 utility, $u'(c) = 1/c^2$ is strictly decreasing, so the FOC has a unique solution:

$$c_1(X) = (u')^{-1}\left(\frac{\lambda}{\beta}\right) \equiv \bar{c}, \quad \text{at every interior } X. \quad (10)$$

The owner's $t = 1$ consumption is the same constant \bar{c} in every state where the contract is interior.

Step 3: Solve for $\delta(X)$ on the interior. From $c_1(X) = w + (1 - \delta(X))X = \bar{c}$:

$$\delta(X) = 1 - \frac{\bar{c} - w}{X}.$$

Define $K \equiv \bar{c} - w$. Then $\delta(X) = 1 - K/X$ on the interior.

Step 4: Limited liability and the corner. The interior expression $1 - K/X$ lies in $[0, 1]$ if

and only if $X \geq K$. For $X < K$ the interior expression gives $\delta < 0$, which violates limited liability. The optimum binds at $\delta(X) = 0$, leaving the owner with $c_1 = w + X < \bar{c}$. Combining the two regions yields equation (8), and the investor's payoff $\delta^*(X)X = \max(X - K, 0)$ in equation (7).

Step 5: Pin \bar{c} from break-even. On the interior region $X \geq K$ the investor receives $X - K$. On the truncation region $X < K$ the investor receives 0. Hence

$$\mathbb{E}[\delta^*(X_1)X_1] = \mathbb{E}[\max(X_1 - K, 0)] = (1 + r)I(\gamma).$$

The left-hand side is strictly decreasing and continuous in K , so the equation has a unique solution that pins K and equivalently $\bar{c} = K + w$.

Convexity. The investor's payoff $\max(X - K, 0)$ is the kinked call option with strike K , a convex function of X . The debt contract's investor payoff $\min(D, X)$ is concave. The optimal contract is the structural inverse of debt. Figure 2 compares the three. \square




FIGURE 2. **Investor's payoff under three contracts (schematic).** Debt's $\min(D, X)$ is concave: the investor takes the lower tail and gets capped at D . Equity's δX is linear. The optimal contract's $\max(X - K, 0)$ is convex: the investor gets nothing on the lower tail and takes the upper tail.

The owner would prefer δ to be negative in the bottom of the distribution. With negative δ the investor would pay her in bad states, topping up her consumption to \bar{c} everywhere. Limited liability rules this out, leaving the corner where $\delta = 0$ and her consumption falls to $w + X$. What the owner wants is not just funding but insurance. A risk-averse agent undertakes risky innovation only when her future consumption is insulated from bad realisations. Credit cannot deliver this insulation, and owners prefer the status quo to projects whose expected payoff is otherwise positive.

The amplification wedge

The optimal contract from Proposition 1 is not generically available. Banks issue debt that pays the lender $\min(D, X)$ for some face value D . Equity markets issue claims that pay the investor a constant share $\bar{\delta}$ of cash flow. Both restrict the owner to a specific functional form for δ . The question is how her optimal γ , the riskiness of the project she undertakes, changes under each.

Without specific parameter values, the optimal γ has no closed form. The goal here is to identify which parameters matter and how the optimum differs between debt and equity. From the previous subsection we know the unrestricted-contract optimum is constant consumption \bar{c} . Deviation from \bar{c} hurts the owner. A higher \bar{c} helps. For any given γ , equity and debt deliver the same mean consumption $\mathbb{E}[c_1]$, since both contracts pin the investor at break-even. What differs is the distribution of c_1 around that mean.

To isolate the contract-shape effect from scale, define the wealth-normalised consumption $\tilde{c} \equiv c_1/A$ and the wealth-normalised project payoff $g(z; \gamma, \sigma, r) \equiv X_1(z; \gamma)/A$. Under each contract \tilde{c} takes a specific functional form, with the contract constants $\bar{\delta}(\gamma)$ and $d(\gamma)$ pinned by the investor's break-even condition:

$$\tilde{c}_{\text{eq}}(z; \gamma) = r + (1 - \bar{\delta}(\gamma)) \cdot g(z; \gamma, \sigma, r), \quad (11)$$

$$\tilde{c}_{\text{debt}}(z; \gamma) = r + \max(g(z; \gamma, \sigma, r) - d(\gamma), 0). \quad (12)$$

The owner's objective under each contract is $U_C(\gamma; A) = u(c_0) + \beta \mathbb{E}[u(c_1)]$ with CRRA utility $u(c) = -1/c$. The amplification choice is bounded by a technological-and-market ceiling γ_{\max} : the production process has a maximum scaling that no single firm can exceed, and the addressable market sets a further cap when fragmentation prevents the firm from capturing the upper tail. Both objectives factor as $U_C(\gamma; A) = -(1/A) \cdot F_C(\gamma; \sigma, \alpha, r)$, so the location of the FOC in γ is independent of A . Appendix 3 gives the derivation. The owner's optimum under each contract reduces to:

$$\gamma_{\text{eq}}^* = \arg \max_{\gamma \leq \gamma_{\max}} \mathbb{E}_z[u(\tilde{c}_{\text{eq}}(z; \gamma, \sigma, \alpha, r))], \quad (13)$$

$$\gamma_{\text{debt}}^* = \arg \max_{\gamma \leq \gamma_{\max}} \mathbb{E}_z[u(\tilde{c}_{\text{debt}}(z; \gamma, \sigma, \alpha, r))]. \quad (14)$$

The only difference between equations (13) and (14) is the functional form of \tilde{c} . Under equity, \tilde{c}_{eq} is smooth in z : raising γ stretches every quantile of \tilde{c} upward, including the worst states where the owner's marginal utility is highest. Under debt, \tilde{c}_{debt} stays pegged at the wage floor r in every state where the project pays less than the bank's claim ($g(z) \leq d$). Raising γ does nothing for these states. The gain only reaches the upper region $\{g(z) > d\}$, where the owner's marginal utility is already small. Equity therefore pushes γ to the ceiling γ_{\max} to capture the upside across the full support. Debt picks γ well below the passive baseline $\gamma = 1$. The owner takes the project but at a hedged-down amplification: each unit of γ raises the bank's claim faster than it raises her consumption in the states that matter, so her optimum lands well short of where equity lands.

Numerical illustration. The wedge is large for empirically plausible parameters. Calibrating to $\sigma = 0.5$, $\alpha = 0.1$, $A = 100$, $r = 0.05$, $\gamma_{\max} = 8$:

Contract	γ^*	U^*	Notes
Equity	8.00	-0.2029	corner: γ_{\max} binds
Debt	0.80	-0.2096	interior, below the passive baseline $\gamma = 1$

The wedge in amplification is $\gamma_{\text{eq}}^* - \gamma_{\text{debt}}^* = 7.2$. Equity rides the technological ceiling. Debt hedges below the passive baseline: the owner picks an amplification level lower than no investment at all, shrinking her firm's risk profile rather than expanding it.

Robustness. The wedge is structural rather than a level effect. Wealth A scales every term in both objectives linearly, so the FOC location in γ is invariant: a hundred-fold increase in A leaves both γ^* exactly unchanged (Appendix 3). Aggregate stock-market volatility in the empirical literature varies between roughly 0.2 and 0.7 (Bloom, 2009), with tech and biotech at the upper end of cross-sectional dispersion. Higher σ widens the wedge: at $\sigma = 0.2$ debt picks $\gamma_{\text{debt}}^* = 1.48$, at $\sigma = 0.7$ debt picks 0.50, while equity stays at the corner throughout. Investment cost intensity has a similar effect: at $\alpha = 0.5$ debt's optimum collapses to $\gamma_{\text{debt}}^* = 0.19$, while equity remains at the ceiling.

Market potential and γ_{\max} . The sharpest sensitivity is to the technological ceiling itself:

γ_{\max}	γ_{eq}^*	γ_{debt}^*	wedge
2	2.00	0.80	1.20
8	8.00	0.80	7.20
20	20.00	0.80	19.20
50	50.00	0.80	49.20

γ_{\max} binds equity's amplification at every value tested. Debt's amplification is invariant in γ_{\max} , since γ_{debt}^* stays interior at every empirically relevant ceiling. The wedge scales one-to-one with the ceiling.

The model is illustrative; in a structural calibration γ_{\max} would be set sector-by-sector and country-by-country, jointly capturing the underlying technological potential and

the size of the addressable market. The split matters because γ_{\max} separates the financing channel from the market-fragmentation channel. Consider retailers. The underlying retail technology is similar in the United States and Europe, but Walmart can expand across all fifty states while European retailers typically operate in one or two countries: differences in culture, language, regulation, and consumer preferences shrink the effective market and lower the relevant γ_{\max} for European firms. In sectors that serve the global market γ_{\max} is similar across the Atlantic, which is why a small country like Sweden can still operate as the factory for European unicorns.

Wealth invariance: derivation of the FOCs

Readers may skip this derivation. The headline numerical result is reported in the amplification wedge subsection above; the empirical tests follow in Section 5. This subsection shows in detail that the optima γ_{eq}^* and γ_{debt}^* in equations (13)–(14) are determined by integral equations involving $(\gamma, \sigma, \alpha, r)$ alone. Wealth A factors out of both.

Equity

The objective is $U_{\text{eq}}(\gamma; A) = u(c_0) + \beta \mathbb{E}[u(c_1)]$ with $c_0 = w = rA$, $c_1 = w + (1 - \bar{\delta}(\gamma))X_1$, $\beta = 1/(1 + r)$, and $u(c) = -1/c$.

Step 1: Project payoff factorises as wealth times shape. Substitute $x_0 = (1+r)A / \exp(\sigma^2/2)$ into $X_1 = x_0 \exp(\sigma z)(1 + \Phi(z))^\gamma$:

$$X_1(z, \gamma) = A \cdot g(z; \gamma, \sigma, r), \quad g(z; \gamma, \sigma, r) \equiv (1 + r) e^{-\sigma^2/2} e^{\sigma z} (1 + \Phi(z))^\gamma. \quad (15)$$

The function g does not contain A .

Step 2: Equity break-even share inherits the scaling.

$$\bar{\delta}(\gamma) = \frac{(1 + r) \alpha \gamma^2 A}{\mathbb{E}[X_1]} = \frac{(1 + r) \alpha \gamma^2 A}{A \cdot \mathbb{E}[g]} = \frac{\alpha \gamma^2 (1 + r)}{\mathbb{E}[g(\gamma, \sigma, r)]}.$$

A cancels in numerator and denominator.

Step 3: Owner's $t = 1$ consumption factorises.

$$c_1 = w + (1 - \bar{\delta})X_1 = rA + (1 - \bar{\delta})Ag = A \cdot [r + (1 - \bar{\delta})g(z)] \equiv A \cdot h_{\text{eq}}(z; \gamma, \sigma, \alpha, r).$$

Step 4: Utility factorises as $-1/A$ times an A -free term.

$$\begin{aligned} u(c_1) &= -\frac{1}{A h_{\text{eq}}} = -\frac{1}{A} \cdot \frac{1}{h_{\text{eq}}}, & u(c_0) &= -\frac{1}{rA} = -\frac{1}{A} \cdot \frac{1}{r}, \\ U_{\text{eq}}(\gamma; A) &= -\frac{1}{A} \left[\frac{1}{r} + \beta \mathbb{E} \left[\frac{1}{h_{\text{eq}}(z; \gamma, \sigma, \alpha, r)} \right] \right] \equiv -\frac{1}{A} \cdot F_{\text{eq}}(\gamma; \sigma, \alpha, r). \end{aligned}$$

Step 5: First-order condition. The factor $-1/A$ cancels:

$$\frac{\partial U_{\text{eq}}}{\partial \gamma} = -\frac{1}{A} \cdot \frac{\partial F_{\text{eq}}}{\partial \gamma} = 0 \iff \mathbb{E} \left[\frac{\partial h_{\text{eq}} / \partial \gamma}{h_{\text{eq}}^2} \right] = 0. \quad (16)$$

The integral equation involves only $(\gamma, \sigma, \alpha, r)$. Together with the box constraint $\gamma \leq \gamma_{\text{max}}$, it determines γ_{eq}^* .

Debt

Same setup with debt $\delta_{\text{debt}}(X) = \min(D, X)/X$, break-even $\mathbb{E}[\min(D, X_1)] = (1 + r)\alpha\gamma^2 A$.

Step 1. $X_1 = A \cdot g(z; \gamma, \sigma, r)$, identical to the equity case.

Step 2: Break-even pins D linear in A . Try the ansatz $D = A \cdot d$. By linearity of the minimum,

$$\mathbb{E}[\min(D, X_1)] = \mathbb{E}[\min(Ad, Ag(z))] = A \cdot \mathbb{E}[\min(d, g(z; \gamma, \sigma, r))].$$

Setting equal to $(1+r)\alpha\gamma^2 A$ and dividing by A :

$$\mathbb{E}[\min(d, g(z; \gamma, \sigma, r))] = (1+r)\alpha\gamma^2. \quad (17)$$

A cancels. The implicit function $d(\gamma; \sigma, \alpha, r)$ is well-defined and pinned by an A -free integral.

Step 3: Owner's $t = 1$ consumption factorises.

$$c_1^{\text{debt}} = w + \max(X_1 - D, 0) = A \cdot [r + \max(g(z) - d, 0)] \equiv A \cdot h_{\text{debt}}(z; \gamma, \sigma, \alpha, r),$$

with

$$h_{\text{debt}}(z) = \begin{cases} r & \text{on } \{g(z) \leq d\}, \\ r + g(z) - d & \text{on } \{g(z) > d\}. \end{cases}$$

Step 4: Utility factorises identically.

$$U_{\text{debt}}(\gamma; A) = -\frac{1}{A} \cdot F_{\text{debt}}(\gamma; \sigma, \alpha, r), \quad F_{\text{debt}} = \frac{1}{r} + \beta \mathbb{E}\left[\frac{1}{h_{\text{debt}}}\right].$$

Step 5: First-order condition. On the truncation region h_{debt} is constant in γ , so the contribution to the integral vanishes. On the interior $\partial h_{\text{debt}}/\partial\gamma = \partial g/\partial\gamma - \partial d/\partial\gamma$.

Hence

$$\mathbb{E}\left[\frac{(\partial g/\partial\gamma - \partial d/\partial\gamma) \cdot \mathbf{1}_{g(z)>d}}{h_{\text{debt}}^2}\right] = 0. \quad (18)$$

The implicit derivative $\partial d/\partial\gamma$ comes from differentiating equation (17):

$$\frac{\partial d}{\partial\gamma} \cdot \mathbb{P}(g > d) + \mathbb{E}\left[\frac{\partial g}{\partial\gamma} \cdot \mathbf{1}_{g \leq d}\right] = 2(1+r)\alpha\gamma.$$

Both $\partial g/\partial\gamma$ and $\partial d/\partial\gamma$ are A -free. The integral equation determining γ_{debt}^* involves only $(\gamma, \sigma, \alpha, r)$, alongside the box constraint $\gamma \leq \gamma_{\text{max}}$.

Wealth invariance and the kink

Two structural results.

(i) *Wealth invariance.* Both γ_{eq}^* and γ_{debt}^* are determined by integral equations that contain no A . The CRRA-2 specification produces exact homogeneity: a scaling of A rescales the level of utility but not the trade-offs governing γ^* .

(ii) *Structural difference.* The functional difference between \tilde{c}_{eq} in equation (11) and \tilde{c}_{debt} in equation (12) is the kink at $g(z) = d(\gamma)$ in the debt case. Equity directs γ -amplification across the entire support of z ; debt directs it only above the kink. Marginal utility is highest at the bottom of the consumption distribution, which equity reaches and debt does not. The amplification wedge in Section 3 is the consequence.

Numerical solution of equations (16) and (18) proceeds by Gauss-Hermite quadrature.

4 What labour productivity measures, and what it does not

The model in Section 3 predicts that equity-rich economies fund higher- γ projects within each asset class. To test this, the thesis uses EU KLEMS & INTANProd, which decomposes labour productivity growth across non-ICT tangible capital, ICT tangible capital, and three intangible asset classes. The decomposition assumes perfect competition and constant returns, so each asset earns its competitive marginal product. The extra productivity risky innovation pulls out is not credited back to the asset and shows up in residual TFP. Section 5 reads the wedge off total labour-productivity growth.

4.1 Labour productivity versus income per capita

The thesis uses labour productivity per hour as the outcome variable throughout. Income per capita decomposes as

$$\frac{Y}{N} = \frac{Y}{L} \times \frac{L}{N}, \quad (19)$$

the product of labour productivity and the employment-to-population ratio. The first factor is the productivity object the model in Section 3 predicts. The second moves with population aging, immigration, and labour-force participation, none of which the financial-structure mechanism speaks to. Pooling them in income per capita masks the productivity mechanism, the consequence Section 5.2 demonstrates directly.

Output per worker is itself contaminated. Labour productivity decomposes one layer further as $Y/L = (Y/(L \cdot H)) \cdot H$, where H is hours per worker. In countries with rigid labour markets, firms adjust to demand shocks by cutting hours per employee rather than laying workers off, so output per worker understates the productivity adjustment in continental Europe relative to the United States, where headcount is the active margin. The post-1980 European labour-time reductions, including the 35-hour week and earlier retirement, push in the same direction. Output per hour strips both layers in one step.

4.2 The Jorgenson decomposition

EU KLEMS & INTANProd builds on the Jorgenson growth accounting framework (Bontadini et al., 2023; Jorgenson & Griliches, 1967). Under constant returns to scale and perfect competition, each input earns its marginal product, and the cost share of each input equals its output elasticity. Value added growth at the country sector level

decomposes additively into the contributions of each input:

$$\Delta \ln V_{c,i} = \sum_k \bar{s}_{c,i}^k \Delta \ln S_{c,i}^k + \Delta \ln A_{c,i}, \quad (20)$$

where $\bar{s}_{c,i}^k$ is the cost share of capital type k , $\Delta \ln S_{c,i}^k$ is the log growth of capital services from k , and $\Delta \ln A_{c,i}$ is the residual (TFP).

Subtracting hours growth from both sides and folding labour quality into the residual gives the labour productivity version,

$$\Delta \ln LP_{c,i} = \sum_k \bar{s}_{c,i}^k \Delta \ln(S_{c,i}^k/H_{c,i}) + \Delta \ln A_{c,i}, \quad (21)$$

which EU KLEMS publishes year by year as the LP1ConX family of contribution variables. The classes are tangible non-ICT capital, tangible ICT (the IT and CT hardware), software and databases, innovation property (R&D, design, mineral exploration), and economic competencies (organisational capital, brand, firm specific training). The intangible extension follows Corrado et al. (2005) as implemented in EU KLEMS & INTANProd (Bontadini et al., 2023). ICT assets are deflated using harmonized hedonic price indices following Schreyer (2002), which removes the typical cross country comparability problem in ICT investment series. The within country LP growth dependent variable in Section 5 is the cumulative sum of LP1_G. The asset class moderators are the cumulative sums of LP1ConX_i^{US}, with the United States serving as the benchmark economy and excluded from the analysis.

The cost share equals elasticity identity breaks under imperfect competition. With markup $\mu = P/MC \geq 1$, the cost share equals the output elasticity scaled by the inverse markup,

$$\alpha_{k,j} = \mu_j \cdot s_{k,j}. \quad (22)$$

Measured cost shares understate true output elasticities by exactly the markup factor. Equity rich economies should carry systematically higher markups, since intangibles

raise fixed costs and lower marginal costs (De Ridder, 2024). De Ridder (2024) estimates structurally that the rise of intangible inputs raised aggregate markups by 14.6 percentage points in the United States and 5.6 percentage points in France, accounting for roughly half of the empirical rise in each economy. The fixed cost share of total costs rose from 13 to 23 percent in the United States and from 9.5 to 14 percent in France between the late 1990s and the mid 2010s. The gap in measured cost shares between equity rich and bank heavy economies therefore understates the true gap in output elasticities by exactly the markup ratio. Section 5.6 flags where this matters.

4.3 What survives, what does not

Equation (21) works at the level of a representative country sector cell. EU KLEMS does not see firm level heterogeneity within a sector, and the constant returns and perfect competition assumption applies to the representative firm. Two firms in the same cell, both investing the same fraction of value added in innovative properties, are assigned the same contribution, even if one runs a portfolio of incremental projects and the other bets the firm. The model in Section 3 predicts the aggregation hides exactly the firm level composition the wedge operates on. Two countries that allocate the same fraction of value added to innovative properties can do so on entirely different distributions of project risk.

The averaging is not a small concern. Within four-digit US manufacturing industries, Syverson (2011) reports a 90th to 10th percentile plant TFP ratio of 1.92, so the high productivity plant produces nearly twice the output of the low productivity plant from the same inputs. Bartelsman et al. (2013) document the same pattern across eight countries and trace cross country differences in aggregate productivity to the covariance between firm productivity and size rather than to mean technology. A representative firm aggregator collapses this distribution onto its mean and discards exactly the margin on which the equity wedge operates, namely the right tail willing

to bet the firm.

The wedge between equity rich and bank heavy economies still shows up, but in different places for different asset classes. For innovative properties, US and European economies invest comparable amounts as a share of value added, so average LP contributions are similar across countries. The within firm composition difference migrates to total industry LP via the residual TFP rather than showing up in LP1ConInnovprop directly. For economic competencies the picture is different. Organisational capital is riskier by construction. It cuts to the core of the firm's cash flow rather than sitting in one product line, so safe diversified investment is not available. Equity rich and bank heavy economies invest different amounts, and the gap shows up in LP1ConEconComp itself.

What survives the country level aggregation is the pattern of the diffusion shock across sectors. The United States frontier from 1996 to 2019 absorbed an ICT driven productivity acceleration whose magnitude varied across sectors. Variation in $LP1ConX_i^{US}$ across sectors captures exposure to the shock. Interacting it with country financial structure identifies the differential ability of equity rich and bank heavy economies to absorb the shock, and the regression result bounds the equity wedge from below.

4.4 Implications for the empirical strategy

Three commitments for the empirical strategy follow. The macro Levine extension in Section 5.2 reads off labour productivity per hour from PWT 10.0 (Feenstra et al., 2015) and the BCL frontier panel (Bergeaud et al., 2016) rather than income per capita, since the demographic and hours layers pool the productivity object with confounding factors. The cross sector regression in Section 5.4 uses EU KLEMS $LP1ConX_i^{US}$ flow as the sector exposure benchmark, not a stock intensity, since stock measures the size of the US asset base while flow measures how much productivity it actually delivered. The capital reallocation regression in Section 5.5 reads off published nominal stocks

K_{IT} and K_{GFCF} , which does not depend on the cost share equals elasticity identity and is robust to the markup wedge of equation (22).

The treatment variable in all three is the country level equity share, the share of stock market capitalisation in the sum of stock market capitalisation and bank credit, averaged over 1990 to 1995. Section 5.1 pins down the construction and the sample.

5 Empirical results

The model in Section 3 makes three predictions. First (H1), aggregate growth differs across the financial structure cleavage post-1995 for frontier economies. Second (H2), the difference fires through complementary intangible and ICT investment. Third (H3), equity rich economies extract more labour productivity per unit of intangible spending, because the project mix they fund is more productive.

Figure 3 shows the asset class breakdown of mean labour productivity contributions for the United States against the European 12 average over 1996 to 2019. The gap is concentrated in tangible ICT, software and databases, economic competencies, and residual TFP. Innovation property contributions are nearly identical across the two (US 0.10 pp/yr, EU 0.08), and tangible non ICT capital contributes more in Europe than in the United States. Two of the three intangible classes therefore look the same on average across the Atlantic, even though the model predicts equity rich economies fund more productive intangible investment. I argue this is an artefact of the representative firm, constant returns, and perfect competition framework EU KLEMS rests on. The wedge migrates to residual TFP and to the tangible capital of other firms in the industry rather than showing up in innovation property contributions directly. The remaining sections recover the wedge through cross sector identification with the United States as the benchmark country.

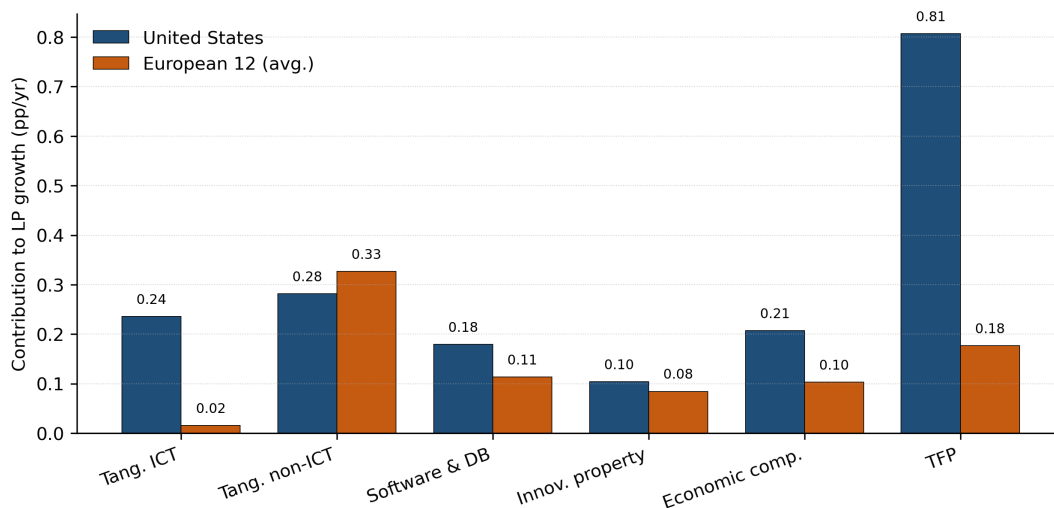


FIGURE 3. Mean contribution to labour productivity growth by asset class, 1996–2019, total economy. The European 12 averages over Austria, Belgium, Denmark, Finland, France, Germany, Italy, the Netherlands, Portugal, Spain, Sweden, and the United Kingdom. Source: EU KLEMS & INTANProd (2025 release), variable LP1Con-.

The ICT contribution gap is too large to be taken at face value. ICT investment intensities are broadly similar across the Atlantic (proxy of $\delta \cdot K_{IT}/VA$ at $\delta = 0.32$ averages 0.49 percent in the United States against 0.54 percent across the European 12), so the gap cannot reflect a hardware investment shortfall in Europe. Hedonic price deflation combined with the 30 percent annual depreciation of ICT hardware makes the contribution mechanically sensitive to the investment trajectory near the endpoint. After 2008 European ICT investment slowed and the contribution turned negative in several countries, dragging the long run average toward zero. A 1996 to 2005 cutoff lifts the EU 12 ICT contribution from 0.02 to 0.04 pp/yr while the United States rises to 0.37, so the gap widens rather than closes. The thesis explains part of this gap through the wedge in complementary intangible investment. Equity rich economies fund the organisational restructuring that turns ICT hardware into productivity, and the model does predict a positive cross country gap. The remaining gap is too severe to take at face value. The alternative reading is that Europe extracted essentially no productivity gain from ICT over twenty four years, which is inconsistent with European productivity sitting close to the United States in aggregate. I suspect the residual reflects a measurement artefact

rather than a real productivity gap of this magnitude.

5.1 Specification and sample

Financial structure is the relative size of equity markets versus bank intermediation. I measure it as the six year average over 1990 to 1995 of stock market capitalisation as a share of total external finance:

$$ES_c = \frac{1}{6} \sum_{t=1990}^{1995} \frac{mcap_{c,t}/GDP_{c,t}}{mcap_{c,t}/GDP_{c,t} + bankcredit_{c,t}/GDP_{c,t}}. \quad (23)$$

Stock market capitalisation comes from the Global Financial Development Database (World Bank, 2022, September 2022 release, variable DM.01). Bank credit is the Bank for International Settlements total credit series for lending to the private non-financial sector from domestic banks, as a share of GDP, with the break adjusted series (Bank for International Settlements, 2026). Portugal's stock market capitalisation in GFDD begins only in 1997. I supplement with historical data from Kuvshinov and Zimmermann (2022) for 1990 to 1996.

The bank credit measure covers the entire private non-financial sector rather than corporations alone. Bank lending to households, primarily mortgages, competes with business lending for the same balance sheet capacity. Spain shows the pattern. With bank credit at 63.9 percent of GDP and an equity share of 0.28, Spanish banks directed credit toward the housing expansion of the 1990s and 2000s, while intangible investment remained the lowest in the sample. Denmark runs the other way, with the highest bank credit ratio in the sample (114.5 percent) and the lowest equity share among Nordic countries (0.20). Total private credit therefore captures the constraint on equity dependent business investment better than corporate credit alone (Langfield & Pagano, 2016).

The 1990 to 1995 window predates the productivity divergence van Ark et al. (2008)

dates to 1995, which reduces reverse causality without eliminating it, since stock market valuations are forward looking. Reverse causality is further attenuated by the persistence of financial structure. Political economy coalitions established by interwar inflation shocks shape modern pension policy and stock market capitalisation (Perotti & Schwienbacher, 2009; Scharfstein, 2018). Rajan and Zingales (2003) show that the political economy framing fits the 1913 to 1999 reversal in cross country market cap to GDP, which legal origin alone cannot explain (La Porta et al., 1998). The German case anchors the mechanism. Interwar hyperinflation wiped out middle class holdings of bonds, pensions, and rental income, and the postwar settlement built a public pension system on the surviving wage base rather than restoring private capital markets (Perotti & Schwienbacher, 2009). Perotti and Schwienbacher estimate that a single interwar inflation shock reduces a country's 2004 private retirement assets by 58 percent of GDP. Scharfstein (2018) extends the chain to the outcomes the thesis measures, with a one standard deviation rise in the public replacement rate cutting stock market capitalisation by 21.8 percentage points and raising the bank loan share of corporate debt by 6.2 points. The 1990 to 1995 equity share is therefore the outcome of decades old political economy choices, not a short run productivity response.

One objection to a market capitalisation numerator is that it omits venture capital and private equity, which finance many of the risky young firms the model targets. Black and Gilson (1998) answer this directly. Venture funds need an IPO exit to recycle capital, so deep public markets and active venture markets coexist. The 1994 cross section makes the point. The United States raised \$34 billion across 591 venture funds, Germany 8.3 billion DM across 85, and venture investment ran at 0.18 percent of GDP in the United Kingdom against 0.06 percent in France. The equity share therefore proxies the wider equity ecosystem rather than missing it.

The sample for the panel work comprises 13 Western frontier economies, listed in Table 1. Risky innovation is a frontier phenomenon, so transition economies are excluded. Their investment came largely from abroad and catch up growth is not what

the model speaks to. Ireland and Greece are excluded for lack of stock market data before 1997, Japan for its distinct lost decade trajectory, and Luxembourg because its financial sector serves the euro area. The Levine extension in Section 5.2 keeps the broader sample close to Levine’s original, with the dependent variable the only substantive change. The industry panel uses only the 18 NACE letter sectors with a distinct economic competencies measurement.

TABLE 1. Equity share and components, 1990–1995 average

Country	Market cap / GDP (%)	Bank credit / GDP (%)	Equity share
United States	71.2	47.3	0.597
United Kingdom	92.1	71.3	0.562
Netherlands	51.7	69.2	0.425
Belgium	32.9	66.5	0.330
Sweden	43.6	103.6	0.294
Spain	24.6	63.9	0.277
France	29.0	76.2	0.276
Finland	20.5	54.3	0.265
Portugal	16.1	47.9	0.249
Germany	20.3	77.8	0.206
Italy	13.9	53.9	0.204
Denmark	28.3	114.5	0.198
Austria	13.4	70.9	0.159

Market capitalisation from GFDD (DM.01) except Portugal (Kuvshinov & Zimmermann, 2022). Bank credit from BIS Total Credit Statistics (private NFS, domestic banks, % GDP, break-adjusted). Equity share = $mcap / (mcap + bankcredit)$, averaged over 1990–1995.

I estimate the panel regressions by OLS with country and industry fixed effects, clustering standard errors at the country level. Clustering matches the level at which financial structure varies and allows within country error correlations across industries. The cluster count is small (12 in the LP growth panel, 11 in the capital reallocation panel after Portugal drops for missing K_{GFCF} coverage). Following Cameron et al. (2008), I report p values using the $t(G - 1)$ distribution with a small sample correction, since asymptotic Wald inference over rejects with few clusters.

5.2 Extending Levine's dataset

Levine (2002) ran the canonical cross country regression of GDP per capita growth on financial structure across 48 countries over 1980 to 1995 and reported a null. Reproducing his Table IV column 1 specification on the same sample gives $\beta = +0.019$ ($p = 0.97$), matching his result to two decimal places. The null reproduces cleanly on the controls that are publicly available. I argue it is regime specific. Levine's window closes the year before the ICT driven productivity divergence between equity rich and bank heavy frontier economies began. The equity wedge of Section 3 predicts that financial structure should become newly load bearing for productivity in this post 1995 regime. A small literature has revisited Levine's null on later samples. Rousseau and Wachtel (2011) document that the finance growth link weakened sharply after 1990 across the standard cross country panels. Arcand et al. (2015) push the same data further and find a non monotonic relation, with private credit losing predictive power above roughly ninety percent of GDP. Demirgüç-Kunt et al. (2013) show the bank versus market mix shifts toward markets as countries grow richer. These papers all use GDP per capita and treat structure as a level. The contribution here is to recover a regime specific structure effect on labour productivity.

Table 2 runs a Chow specification on a Penn World Tables 10.0 sample of 61 countries (97 country period observations across pre and post 1995) with two dependent variables side by side. Column 1 is GDP per capita growth, the variable Levine used. Column 2 is labour productivity per hour growth, the productivity object the model predicts. Stock market capitalisation and bank credit enter as separate regressors, averaged within each period, with their Post1995 interactions. A log initial dependent variable convergence control absorbs country effects, mirroring Levine's cross section design. The contrast is sharp. On GDP per capita, stock market depth is zero in both periods. Only bank credit shows a regime change ($\text{credit} \times \text{Post1995} = -0.025$, $p = 0.006$), the financial development weakening of Rousseau and Wachtel (2011). On labour

productivity per hour, stock market depth flips sign across 1995, from $\beta = -0.017$ pre to $+0.016$ post ($\text{mcap} \times \text{Post1995}$ $p = 0.041$). The structure regime change is real, but invisible on income per capita because the demographic and hours per worker factors mask it.

TABLE 2. **Pre/post-1995 Chow test on a single Penn World Tables sample, two dependent variables side by side.** Column 1: real GDP per capita growth (Levine 2002’s dependent variable). Column 2: labour productivity per hour growth (the productivity outcome the model predicts). Same 97 country-period observations across 61 countries on both columns. Stock-market capitalisation and private credit enter as separate regressors at the percentage-of-GDP level, with their Post1995 interactions. Convergence control: log of initial dependent-variable level. Country-clustered standard errors in parentheses. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	DV: real GDP/capita (Levine 2002 outcome)	DV: LP per hour (productivity outcome)
Stock-market cap / GDP	+0.0008 (0.0075)	-0.0172** (0.0075)
× Post1995	-0.0014 (0.0079)	+0.0156** (0.0076)
Private credit / GDP	+0.0360*** (0.0082)	+0.0417*** (0.0103)
× Post1995	-0.0249*** (0.0091)	-0.0343*** (0.0106)
Post1995 (intercept shift)	+1.8237** (0.7853)	+1.6114** (0.8072)
$\log(\text{DV}_{\text{init}})$	-1.4376*** (0.2424)	-1.2713*** (0.2220)
Observations	97	97
Countries	61	61
R^2	0.271	0.378

To let the data pick the break year, I run a panel DiD on annual BCL data from 1965 to 2020 for the twelve frontier economies in the BCL coverage. BCL was the original dataset for this analysis. The Chow regression above moved to PWT for a cleaner Levine comparison, but break detection is not sensitive to the dataset choice. For each candidate $T^* \in \{1980, \dots, 2010\}$ I estimate

$$\Delta \ln \text{LP}_{i,t} = \alpha_i + \delta_t + \beta(T^*) \cdot (\text{ES}_i \times \mathbf{1}\{t > T^*\}) + \varepsilon_{i,t}, \quad (24)$$

applying the Andrews (1993) sup Wald procedure for unknown break date. Out of 31

candidates the data picks $T^* = 1995$ at $\text{sup-}F = 38.88$, over four times the Andrews 5 percent critical value of 8.85, with 95 percent confidence interval on \hat{T}^* of [1990, 1995]. The corresponding coefficient is $\beta = +5.56$ ($p < 0.0001$). Figure 4 plots the F statistic across candidates.

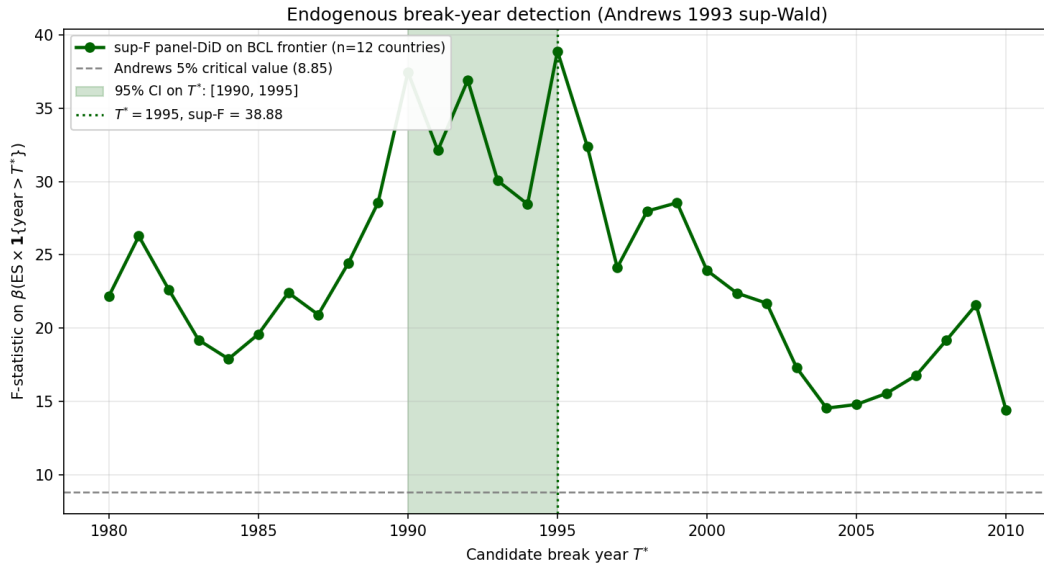


FIGURE 4. **Endogenous break-year detection by the Andrews (1993) sup-Wald procedure.** F-statistic on $\beta(T^*)$ from the panel-DiD specification of equation (24), plotted against candidate break year T^* on the BCL frontier panel (12 countries, 1965–2020 annual). The dashed horizontal line is the Andrews 5 percent critical value of 8.85. The shaded band is the 95 percent confidence interval on \hat{T}^* obtained by inversion of the F-statistic threshold. The maximum is reached at $T^* = 1995$ with $F = 38.88$.

A 0.1 unit increase in 1990 to 1995 ES is associated with 0.56 percentage points higher annual labour productivity growth in the post 1995 era. Across the BCL frontier sample with ES ranging from approximately 0.13 (Austria) to 0.60 (United States), this implies a $0.47 \times 5.56 \approx 2.6$ percentage point differential in annual productivity growth between high ES and low ES frontier economies. Compounded over 25 years, the cumulative gap is approximately 90 percent. This is the right order of magnitude for the transatlantic productivity divergence documented in van Ark et al. (2008) and Bergeaud (2024).

Robustness battery. A five year bucket event study shows pre period coefficients insignificant and post period coefficients decisively positive and stable around +5 from 1995 onwards, the textbook DiD picture. Leave one out across the twelve frontier countries

gives $\beta \in [+5.10, +6.49]$ with $p < 10^{-7}$ on every drop. Frontier threshold sensitivity from 65 to 90 percent of US 1990 GDP per capita gives $\beta \in [+5.56, +7.37]$, growing as the threshold tightens. The subsection documents two simultaneous regime changes around 1995 on disjoint outcome regressor pairs. Financial development lost its predictive power for income growth, the Rousseau Wachtel finding. Financial structure became newly load bearing for productivity growth, the new finding.

5.3 Country-level patterns

Table 3 regresses the 1996 to 2019 mean labour productivity contribution of three asset classes on the equity share. Economic competencies loads strongly: $\beta = +0.43$ ($R^2 = 0.80$, $p < 0.001$, $N = 13$). A ten percentage point increase in the equity share predicts 0.043 additional percentage points of annual labour productivity contribution from economic competencies. Tangible ICT loads moderately ($\beta = +0.41$, $R^2 = 0.51$, $p = 0.009$, $N = 12$). Innovation property is null ($\beta = -0.02$, $p = 0.83$, $N = 13$). Figure 5 plots the economic competencies and innovation property scatters. The wedge between equity rich and bank heavy economies surfaces at the country level on economic competencies and ICT, but not on innovation property.

TABLE 3. Labour productivity contribution regressed on the equity share, 1996–2019 country averages

Dependent variable (pp/yr)	$\hat{\beta}$	R^2	p	N
Economic competencies (LP1ConEconComp)	+0.43	0.800	< 0.001	13
Tangible ICT (LP1ConTangICT)	+0.41	0.512	0.009	12
Innovation property (LP1ConInnovprop)	-0.02	0.005	0.83	13

Bivariate OLS regression of country-mean labour-productivity contribution from each asset class (1996–2019 average) on the 1990–1995 equity share. Belgium drops for ICT due to missing K_{IT} coverage.

These results rest on the EU KLEMS perfect competition and constant returns framework, which prices each capital unit at the same competitive marginal product. The framework can see whether the capital stock of each asset class grew faster in equity rich economies, but not whether the innovation those investments funded was

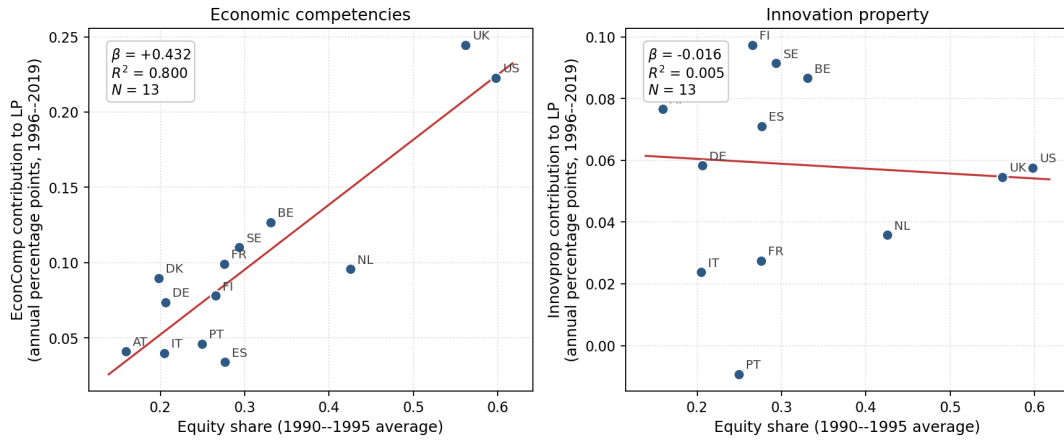


FIGURE 5. **Equity share and labour-productivity contributions of intangible asset classes.** Left panel: each country’s 1996–2019 mean annual contribution to labour productivity from economic competencies (LP1CONECONCOMP), regressed on the 1990–1995 equity share. Right panel: same construction for innovation property (LP1CONINNOVPROP). Source: EU KLEMS & INTANProd (LUISS 2025), GFDD, BIS.

riskier or more productive. Economic competencies loads strongly because equity rich economies invest more in organisational capital. The weak ICT and null innovation property results reflect the framework’s measurement limit. The within asset class project mix wedge migrates to residual TFP, where the country level scatter cannot pick it up.

A second caveat concerns measurement. Intangibles in EU KLEMS are largely measured from the wages of specific occupations, primarily managers for organisational capital. The classification of who counts as a manager varies across countries, which can bias the level of measured intangible capital (Nonnis et al., 2024). The country level analysis above uses averages of growth based contributions rather than levels, which absorbs much of the country specific level bias. The concern remains that the growth rates themselves can carry residual country specific noise. Section 5.4 mitigates this further, since country fixed effects strip out any country level measurement artefact, and identifies the within asset class wedge via a Rajan-Zingales style identification with the United States as the benchmark country.

5.4 Equity, the intangible-ICT bundle, and labour-productivity growth

For each asset class X I regress cumulative country sector labour productivity growth over 1996 to 2019 on the equity share interacted with the United States sector labour productivity contribution from X over the same window:

$$\Delta \ln LP_{c,i}^{[1996,2019]} = \beta_X \cdot ES_c \times T_X^{US}(i) + \alpha_c + \alpha_i + \varepsilon_{c,i}. \quad (25)$$

$\Delta \ln LP_{c,i}$ is the cumulative sum of LP1_G over the window. $T_X^{US}(i)$ is the cumulative US sector contribution from asset class X over the same years. Country fixed effects α_c and sector fixed effects α_i absorb both moderator levels. The sample covers 12 European economies and 18 NACE letter sectors. The United States is the benchmark and stays off the left hand side.

The two way fixed effects do most of the identification work. Country effects strip out ES_c as a level, ruling out a rich countries grew faster story. Sector effects strip out $T_X^{US}(i)$ as a level, ruling out an intangible intensive sectors grew faster story. What remains is the cross country by cross sector interaction. In equity rich countries, sectors that the US frontier showed could be transformed by asset class X should grow disproportionately faster in labour productivity. The benchmark T_X^{US} is a realised contribution rather than a stock intensity, since a stock ratio measures the size of the US asset base while a flow measures what productivity it actually delivered. Section 5.6 reports the contrast across both benchmarks.

The downside of stripping out both levels is that most of the variation is absorbed, leaving the interaction to identify off whatever cross country by cross sector variation survives. With that caveat, all five intangible and ICT asset classes load on the equity share interaction. Table 4 reports every interaction at $p \leq 0.012$. Innovation property loads heaviest ($\beta = +18.80$), software and databases next ($\beta = +13.27$), ICT hardware

third ($\beta = +9.10$), all intangibles combined fourth ($\beta = +7.57$), and economic competencies fifth ($\beta = +5.13$). EU KLEMS does not decompose the labour productivity contribution below the economic competencies aggregate, which bundles organisational capital with brand and firm specific training. Direct OrgCap only LP loadings are therefore not observable in this dataset, and the EconComp coefficient is the cleanest available proxy. Two facts make it informative. First, organisational capital is the largest of the three components by investment volume across the EU KLEMS panel, so the proxy weights toward the channel the framework targets. Second, the loading survives at $p = 0.002$ after the stock intensity specification returned a clean null in Section 5.3, which the realised flow benchmark would not produce if the EconComp signal were driven by brand or training. The estimate is best read as a lower bound on the OrgCap specific effect.

Innovation property moves from insignificance in the country level scatter (Section 5.3) to the strongest coefficient here. The shift reconciles with Brown et al. (2017), who finds higher labour productivity growth in equity heavy countries for sectors with high US R&D intensity. Brown et al. (2017) runs an industry by country panel of value added growth on stock market value traded interacted with US R&D intensity, 38 countries 1980 to 2005, and reports an interaction of 0.0263 (SE 0.0116). Two design differences narrow the present test. First, the dependent variable is realised LP growth rather than value added growth, so the channel is read off productivity rather than expansion. Second, the benchmark spans five EU KLEMS asset classes rather than R&D alone, so innovation property is one of five readings rather than the whole exercise. The wedge model in Section 3 supplies the theoretical wrapper Brown's reduced form lacks. The shift also points to a structural limitation of EU KLEMS & INTANProd. The dataset cannot capture the productivity of each asset class properly under the representative firm, constant returns, and perfect competition assumptions. Superstar firms drive sector productivity disproportionately from their intangible and ICT investment, but their above competitive returns leak into residual TFP and into the tangible capital of

other firms in the industry, since the Jorgenson framework works with the industry’s capital mix rather than the superstar’s own. Autor et al. (2020) document that between firm reallocation among US manufacturing survivors accounts for essentially all of the 18.5 percentage point fall in the compensation share of value added between 1982 and 2012, with concentration rising fastest in patent and TFP intensive industries. Bajgar et al. (2025) extend the pattern to 13 OECD economies. A 10 percentage point rise in industry intangible investment intensity predicts a 1.8 percentage point rise in CR8 concentration four years later, with effects loading on innovative property rather than economic competencies. The cross sector identification recovers the wedge by interacting country financial structure with the United States frontier benchmark, where the superstar contribution surfaces.

TABLE 4. RZ-style panel: cumulative country-sector LP growth on $ES \times T_X^{US}$, 1996–2019

Asset class X	$\hat{\beta}_X$	SE	p	R^2
ICT hardware (LP1CONTANGICT)	+9.10	2.85	0.009	0.599
Software and databases (LP1CON_SOFT_DB)	+13.27	3.29	0.002	0.584
Innovation property (LP1CONINNOVPROP)	+18.80	6.22	0.012	0.618
Economic competencies (LP1CONECONCOMP)	+5.13	1.27	0.002	0.560
All intangibles (LP1CONINTANG)	+7.57	1.83	0.002	0.601

Two-way fixed-effects regression of cumulative country-sector LP growth over 1996–2019 on the equity share interacted with the cumulative US sector labour-productivity contribution from asset class X over the same window. Country and sector fixed effects included. Standard errors cluster by country ($G = 12$). Sample: 12 European economies, 18 NACE-letter sectors. $N \in [204, 216]$ depending on EU KLEMS coverage of the asset-class variable.

The point estimates are economically meaningful. With country and sector fixed effects absorbing both levels, the coefficient identifies how much more strongly equity rich economies outperform bank heavy ones in high US frontier exposure sectors. The United Kingdom and Austria sit at the ends of the European equity share distribution ($ES = 0.57$ and $ES = 0.16$). At a top quartile innovation property sector with $T_{Innov}^{US} \approx 6$ percentage points of cumulative US LP contribution, the within sector differential between the two countries is $18.80 \times 0.41 \times 6 \approx 46$ percentage points of cumulative LP growth over 1996 to 2019, or roughly two percentage points per year. Outside a general equilibrium structural model, such comparisons are illustrative rather than authorita-

tive. The magnitude is in the same range as the transatlantic and intra European LP gaps documented in Bergeaud (2024).

5.5 Capital reallocation: the input side

A productivity result on its own invites the suspicion that the moderator is picking up something other than capital reallocation. If equity rich economies grow faster in high frontier exposure sectors only because their workers are better trained or their firms are better managed in ways orthogonal to investment, the LP regression cannot tell. The model in Section 3 predicts a sharper input side prediction. Equity rich economies should also tilt their capital composition toward IT in sectors where the complementary organisational restructuring is heaviest. This subsection runs that test directly on capital stocks. Section 5.4 identified the channel on productivity; this subsection identifies it on capital allocation. The specification is

$$\Delta \log\left(\frac{K_{IT}}{K_{GFCF}}\right)_{c,i,1996 \rightarrow 2010} = \alpha_c + \gamma_i + \beta (ES_c \times T_{org}^{US}(i)) + \delta \log\left(\frac{K_{IT}}{K_{GFCF}}\right)_{c,i,1996} + \varepsilon_{c,i}, \quad (26)$$

where $T_{org}^{US}(i) = K_{US,i,2005}^{OrgCap} / VAadj_{US,i,2005}$ and the initial state control absorbs β convergence at the cell. The sample is 195 country sector cells across 11 non US countries (Portugal drops for missing K_{GFCF} coverage) and 18 sectors. The treatment side benchmark is the US organisational capital stock to value added intensity at 2005, a Rajan-Zingales structural complementarity proxy rather than the LP contribution flow benchmark of Section 5.4.

Equity rich economies reallocated capital toward IT in organisational capital intensive sectors. The interaction is positive and highly significant. Column 1 of Table 5 reports $\beta = +19.42$ ($p = 0.002$). The estimate is identified off variation that survives both country and sector fixed effects. In organisational capital intensive sectors, IT capital intensity grew faster in equity oriented countries than in bank oriented countries, even

TABLE 5. Primary RZ-style panel: $\Delta \log(K_{IT}/K_{GFCF})_{1996 \rightarrow 2010}$ on $ES \times T_{org}^{US}$

	(1) Main	(2) + IP placebo	(3) Full horse-race
$ES \times T_{org}^{US}$	+19.42*** (4.67)	+19.62*** (5.97)	+16.66** (5.27)
$ES \times T_{ip}^{US}$		+0.10 (0.78)	+0.07 (0.78)
$BankCredit \times T_{org}^{US}$			-5.41 (3.11)
$\log(GDPpc) \times T_{org}^{US}$			-11.12 (6.77)
$\log(K_{IT}/K_{GFCF})_{1996}$	-0.31*** (0.08)	-0.31*** (0.09)	-0.33*** (0.09)
N	195	195	195
Country fixed effects	yes	yes	yes
Sector fixed effects	yes	yes	yes
R^2	0.43	0.43	0.45

Cluster-robust standard errors at the country level ($G=11$) in parentheses, with the small-sample correction described in Section 5.1. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. $T_{org,i}^{US} = K_{OrgCap}^{US,2005,i} / VAadj^{US,2005,i}$. $T_{ip,i}^{US}$ is the analogous innovation-property ratio. Bank credit is private non-financial sector lending, 1990–1995 average from BIS, expressed as a share of GDP. GDP per capita is from World Bank WDI, 1996, PPP constant 2017 international dollars.

after absorbing every country level confounder (legal origin, GDP per capita, baseline ICT intensity) and every sector level confounder. The initial state coefficient $\delta = -0.31$ is negative and significant, consistent with β convergence at the country sector cell.

Column 2 adds the moderator side placebo $ES_c \times T_{ip}^{US}(i)$, which enters at +0.10 ($p = 0.90$). For IT capital reallocation, the productive complement is the organisational restructuring around the new hardware, comprising management practices and process redesign captured by organisational capital (Bloom et al., 2012). Bloom et al. (2012) pin this down at the establishment level. In a UK panel of 7,121 establishments matched to the CEP management survey, US multinationals raise labour productivity by 4.9 percent for a doubling of IT capital, against 1.2 percent at domestic or non US firms. The premium vanishes once a direct people management score is interacted with IT intensity, pinpointing organisational practice rather than innovation property as the binding ICT complement. Innovation property funds risky innovation in general,

but it is not the specific complement to ICT adoption. Its null here confirms that the wedge is not a generic intangible finance story but specifically about risky restructuring around newly bought IT hardware. Column 3 runs the full horse race. Bank credit interacted with T_{org}^{US} enters at -5.41 ($p = 0.11$), consistent with bank heavy economies pulling capital allocation away from organisational capital intensive sectors. GDP per capita interacted with T_{org}^{US} enters at -11.12 ($p = 0.13$), ruling out a rich countries do everything more alternative. The equity coefficient survives essentially intact at $+16.66$ ($p = 0.010$).

The bank credit interaction is the cleanest placebo for a generic credit deepening story. If credit volume drove ICT reallocation, $BankCredit \times T_{org}^{US}$ should load positive. It loads wrong signed at -5.41 . Bank heavy economies pulled capital allocation away from organisational capital intensive sectors, the opposite of what a financial deepening alternative predicts. The wedge is about contract form, not about how much external finance an economy carries.

A one standard deviation increase in the $ES \times T_{org}^{US}$ interaction predicts $19.42 \times 0.110 \times 0.056 = 0.120$ additional log points of IT capital share growth, about 12.7 percent over the 1996 to 2010 window. For a marginal sector with average $T_{org}^{US} = 0.118$, IT capital share grew about 0.83 log points more in the United Kingdom ($ES = 0.56$) than in Italy ($ES = 0.20$). As in Section 5.4, this comparison is illustrative rather than authoritative outside a general equilibrium model. The wedge operates at a magnitude comparable to half the observed IT adoption gap between bank heavy and equity rich frontier economies.

5.6 Robustness

The headline coefficient rests on 195 country sector cells, two strong anchor points, and a single LP growth window. Each of these is a place a sceptic could break the result. A leverage driven coefficient would collapse on the leave one out drop. A window specific

coefficient would not survive endpoint variation. A benchmark specific coefficient would not survive switching to a flow based moderator. The wedge would not generalise across asset classes if it were not the lumpy slice of intangible investment doing the work. The remainder of this section runs each of these tests in turn.

The headline result in Section 5.4 and the input-side corollary in Section 5.5 survive a battery of robustness checks. I summarise five: window sensitivity for both panels; moderator timing for the capital-reallocation specification; leave-one-out sensitivity to country and sector composition; a 2×2 specification grid that crosses the two treatment-side benchmarks with the two outcome variables; and a reconciliation of the innovation-property placebo behaviour across stock-intensity and flow-contribution benchmarks.

Window sensitivity. For the capital-reallocation specification, extending the LHS endpoint from 2005 to 2007, 2010, or 2015 strengthens the result (peak $\beta \approx +20$ at $1996 \rightarrow 2010$); extending to 2019 attenuates to $+6.75$ ($p = 0.53$), consistent with the post-2015 cloud-leasing decoupling of on-premises K_{IT} from underlying technology adoption. For the LP-growth specification in equation (25), a six-window sweep across 1996–2005, 1996–2010, 1996–2019, 2001–2010, 2006–2015, and 2011–2019 traces a coherent pattern. All five asset-class coefficients hold their sign and significance through the early windows, with magnitudes within a factor of two of the headline. The 2001–2010 and 2006–2015 windows attenuate, with most coefficients losing significance. The attenuation traces mechanically to the cross-section of ES_c rather than to the mechanism: between 1991–1995 and 2006–2010 the United Kingdom-to-Austria range fell from 0.42 to 0.23 as continental European stock markets matured, and a moderator with shrinking variance has less identifying power. The 2011–2019 window shows the most economically interesting pattern: economic competencies flips sign to $\beta = -5.23$ ($p = 0.42$), while ICT strengthens to $\beta = +17.12$ ($p = 0.013$) and innovation property to $\beta = +33.54$ ($p = 0.075$). The first reads as the organisational-capital diffusion shock having passed: post-2010 United States frontier economic-competencies gains shifted

to retail and services, where equity has nothing to add. The second reads as the equity premium for IT and innovation widening after Basel III tightened intangible-collateral lending (Basel Committee on Banking Supervision, 2017). Both patterns are what the model predicts. Tightening the LP growth window to the diffusion only 1996 to 2005 endpoint sharpens precision rather than the point estimates. All five asset classes load positive at $p < 0.005$, with the intangibles aggregate at $\beta = +8.33$ ($p < 10^{-5}$) and innovation property at $\beta = +19.06$ ($p = 0.001$), essentially identical to the 1996 to 2019 headline of +18.80. The headline does not depend on post 2010 observations.

Moderator timing. For equation (26), the result holds whether $T_{\text{org}}^{\text{US}}(i)$ is measured at 1995 ($\beta = +10.66$, $p = 0.06$), 2000 (+10.05, $p = 0.01$), 2005 (+8.97, $p = 0.013$ in the diffusion-only window), or averaged over 1995–2005 (+10.36, $p = 0.018$); switching from nominal K_{OrgCap} to the real quantity index K_{OrgCap}^q at 2005 gives +7.72 ($p = 0.012$). Replacing $T_{\text{org}}^{\text{US}}(i)$ with the broad $T_{\text{intang}}^{\text{US}}(i) = K_{\text{US},i,2005}^{\text{Intang}}/\text{VAadj}_{\text{US},i,2005}$ delivers $\beta \approx +0.07$ ($p = 0.93$); the organisational-capital-specific story does not generalise to broad intangibles, which is the prediction the model in Section 3 makes when only the lumpy slice carries the wedge.

Leave-one-out sensitivity. Dropping each of the 11 sample countries in turn yields a coefficient in [+17.4, +25.3] with $p < 0.025$ in every case except the United Kingdom; dropping the United Kingdom attenuates the coefficient to +8.79 ($p = 0.42$). The United Kingdom is the equity-rich anchor of the sample (its equity share of 0.562 is more than two standard deviations above the next-highest, the Netherlands at 0.425), so dropping it compresses the effective range of the equity-share variation and reduces statistical power without reversing the channel. Dropping each of the 18 sectors in turn yields $\beta \in [+17.7, +25.3]$ with $p \leq 0.022$ except for sector A (agriculture), which has the lowest United States organisational-capital intensity ($T_{\text{org},A}^{\text{US}} = 0.002$); dropping it removes the low- T anchor that pins the regression slope. Both fragility points reflect the mechanics of the test rather than alternative interpretations.

Specification grid. The treatment-side benchmark and the dependent variable can each take one of two natural forms. The treatment is either United States sector organisational-capital stock-to-value-added intensity at 2005 (the structural-complementarity reading) or United States sector economic-competencies real-stock log-growth from 1995 to 2010 (the dynamic-complementarity reading aligned to the diffusion window). The dependent variable is either $\Delta \log(K_{IT}/K_{GFCF})_{c,i,1996 \rightarrow 2010}$ (capital reallocation) or the average ICT contribution to labour-productivity growth over 1996–2010 (productivity outcome). Crossing the two yields four specifications. Three deliver coefficients positive at the 5-percent level: $\beta = +19.42$ ($p = 0.002$) for the organisational-capital-intensity treatment with the capital-reallocation outcome (the headline), $\beta = +3.75$ ($p = 0.019$) for the economic-competencies-growth treatment with the same outcome, and $\beta = +0.32$ ($p = 0.019$) for the economic-competencies-growth treatment with the LP-contribution outcome. The fourth cell, organisational-capital-intensity treatment with the LP-contribution outcome, is right-signed but not statistically significant ($\beta = +1.23$, $p = 0.13$). The wedge therefore is identified across three of four reasonable combinations of the two treatment-side benchmarks and the two outcome variables.

Stock-versus-flow innovation-property asymmetry. The innovation-property placebo is null when the treatment is a stock-intensity benchmark (Section 5.5, Column 2: $\beta = +0.10$, $p = 0.90$) and approximately null when it is a real-stock-growth benchmark, but it is not null when either treatment is paired with the LP-contribution outcome (Section 5.4, Table 4: $\beta = +18.80$, $p = 0.012$). The asymmetry is mechanical and consistent with the model. Equity-rich economies extract higher productivity per unit of innovation-property capital because they fund a different composition of projects within the same aggregate R&D investment intensity. The perfect-competition and constant-returns assumption inside the EU KLEMS growth-accounting framework averages this composition effect away by construction (Section 4.3). The stock-intensity benchmark, which reads off published nominal stocks, does not see the composition effect; the LP-contribution benchmark, which weights by realised factor returns, does.

Both behaviours are predictions of the model: the wedge binds on the lumpy slice of intangible investment and shows up in capital reallocation under the stock benchmark, while it shows up in productivity returns under the flow benchmark precisely because equity-rich economies pick higher-return projects within each asset class.

The four results trace the wedge from macro to micro. Financial structure becomes load-bearing on productivity at 1995 (Section 5.2). The pattern shows up at the country level on organisational capital (Section 5.3), at the cross-sector level on the full intangible-ICT bundle (Section 5.4), and on the input side as capital reallocation toward IT (Section 5.5). The magnitudes match what the model in Section 3 predicts.

The four results line up with the model. They do not exhaust the ways the design could be challenged. Sample size is small at every layer. The equity share is measured as an average over a five year window, and stock market valuations are forward looking. Two anchor observations carry disproportionate weight in the cross sector panel. The capital reallocation window stops at 2010 for measurement reasons that are themselves an empirical choice. Section 6 takes each of these in turn.

6 Limitations

The thesis claims that financial structure binds on risky innovation in frontier economies after 1995. Six choices the design makes leave that claim open to challenge. I take each in turn.

Sample. The country level scatter uses 13 frontier economies. The cross sector LP growth panel uses 12 (the United States is the benchmark, off the left hand side). The capital reallocation panel uses 11. Portugal drops because K_{GFCF} is missing. The Levine extension uses 61 PWT countries on the Chow specification (97 country period observations across pre and post 1995) and the same 12 BCL frontier economies on the panel DiD. The thesis does not claim the wedge fires the same way in catch up

economies, where Levine’s null held in his original sample and where bank credit funds tangible accumulation that fits collateral well. External validity is bounded to frontier economies after 1995.

Predetermination of the equity share. Stock market valuations are forward looking, so the equity share could in principle respond to expected future productivity. Three things mitigate the concern. First, ES is measured at 1990 to 1995, before the post 1995 outcomes. Second, financial structure persists over generations through the political economy and legal origin channels developed in Section 5.1 (La Porta et al., 1998; Perotti & Schwienbacher, 2009; Rajan & Zingales, 2003; Scharfstein, 2018), so 1990 to 1995 valuations are the outcome of decades old political economy coalitions rather than short run productivity expectations. Third, the Section 5.2 endogenous break test picks 1995 from 31 candidate years rather than imposing it. If 1990 to 1995 valuations were endogenous to expected post 1995 productivity, the break date should sit later in the window, not at the boundary.

Two leverage points. The capital reallocation panel has two anchor observations. The United Kingdom is the equity rich anchor, with $ES = 0.56$, two standard deviations above the next highest country. Sector A (agriculture) is the low T anchor, with US OrgCap intensity of 0.002. Dropping the United Kingdom attenuates the coefficient from +19.62 to +8.79. Dropping sector A attenuates similarly. Neither drop reverses the sign. Both fragility points reflect the test’s identification mechanics. When an anchor is removed the regression loses range, the standard error widens, and the coefficient is harder to distinguish from zero. The channel is unchanged.

Manager classification across countries. The moderator $T_{org}^{US}(i)$ uses the United States as the benchmark, which sidesteps the cross country measurement problem in OrgCap, but the dependent variable in the H3 panel is K_{IT}/K_{GFCF} rather than European OrgCap, so the headline result is insulated. The wider concern, documented by Nonnis et al. (2024), is that OrgCap own account investment is built from manager compensation

under ISCO classifications that are not harmonised across countries. Levels of K_{OrgCap} are therefore unreliable for cross country comparison. Growth rates are cleaner but not clean. Any specification that uses European OrgCap stocks on the right hand side, including the Section 5.4 EconComp channel, inherits this measurement noise and biases coefficients toward zero.

Representative firm aggregation. The Jorgenson framework prices each capital unit at the same competitive marginal product within a country sector cell, so the within sector heterogeneity in firms' γ choice is collapsed into a single average. The wedge the model formalises operates on this margin and cannot be observed directly in EU KLEMS. The cross sector regression identifies a lower bound on the equity wedge through the diffusion shock channel, with the within sector composition channel showing up only as residual TFP and as the tangible capital of other firms in the industry. Closing this gap requires firm level data EU KLEMS does not provide.

Window dependence. The capital reallocation result peaks at the 1996 to 2010 window with $\beta = +19.6$ and attenuates at 1996 to 2019 to $\beta = +6.75$. The likely explanation is the post 2015 shift from owned IT hardware to leased cloud services. Once firms substitute leased compute for on premises capital, K_{IT} becomes a capital formation accounting object rather than a measure of technology adoption. The LP growth panel does not suffer the same drift because cumulative LP growth is built from value added growth, not from K_{IT} . The thesis uses 1996 to 2010 as the headline window for the capital reallocation specification for that reason.

ICT contribution measurement. The European 12 mean LP contribution from tangible ICT is 0.02 pp/yr against the United States 0.24, a gap larger than the model alone predicts and larger than ICT investment intensity differences justify. Hedonic price deflation combined with fast depreciation and the post 2008 European investment slowdown are the most plausible mechanical sources, but the residual is severe enough that I treat it as a measurement caveat rather than a real productivity gap. Section 5

reports the headline result on the LP growth panel and the K_{IT} panel without relying on this comparison.

Aggregate magnitude is not pinned down. The empirical coefficients are industry level. Translating them into a country level productivity gap magnitude requires a general equilibrium model with factor substitution, firm dynamics, and the demographic factor of Section 4.1. The implied 12.7 percent differential in IT capital share growth from a one standard deviation increase in the interaction is large enough to matter when compounded with the EconComp channel, but the compound calculation needs structure the thesis does not build. Closing the magnitude gap is the natural direction for future work.

Model abstractions. The model is two period, with no recurring cash flow, so the timing of risky innovation collapses to a single decision. Equity and credit are assumed to read the firm's profit distribution equally well, so the screening role of finance is absent. The model speaks most cleanly to larger frontier firms with public track records, less cleanly to entrants where screening competes with insurance. Effort moral hazard (Innes, 1990; Stiglitz & Weiss, 1981) is left for future work. Risk is the binding friction for risky innovation in this thesis, not effort. Garcia-Macia (2017) formalises the collateral channel as a complementary within asset class mechanism that fits inside the framework.

7 Conclusion

Financial structure shapes which risky innovations firms undertake. A risky innovation is a project with positive mean payoff that raises the firm's cash flow in some states and lowers it in others. Debt and equity can both fund it at the same expected cost. The lender raises the interest rate to compensate for the bad states. What differs is who carries the consumption risk. Under debt the owner takes the bottom of the distribution.

Under equity the investor shares every state. A risk-averse owner prefers equity for the same reason anyone prefers insurance. In bank-heavy economies, where she can only borrow, the calculation cuts the other way. She does not undertake the project at all. That is the wedge.

The thesis sits between two literatures. Schumpeterian endogenous growth models put risk in the timing of incumbent displacement, where the incumbent has no choice over the risk profile of its own production (Aghion & Howitt, 1992; Klette & Kortum, 2004). The financial structure literature, following Levine (2002), treats structure as a level effect on aggregate growth. Adding a firm level γ choice, and the financing contract that shapes it, connects the two. Equity rich economies fund higher γ projects within each asset class, and the cross country gap in growth that follows is the macro shadow of a micro project mix decision.

Three findings test the model. *H1, the macro break.* Levine's 1980 to 1995 null on financial structure and growth reproduces at $\beta = +0.019$ ($p = 0.97$). On labour productivity per hour after 1995, the same variation predicts growth at $\beta = +5.56$, with the break year picked endogenously at 1995 ($\text{sup-}F = 38.88$). *H2, the cross sector channel.* Across 12 European economies and 18 NACE letter sectors, the equity share interacted with each US frontier asset class LP contribution loads at $p \leq 0.012$, with innovation property heaviest at $\beta = +18.80$. *H3, the input side check.* The same variation interacted with US OrgCap stock intensity loads at $\beta = +19.42$ on capital reallocation toward IT, while bank credit and GDP per capita placebos load wrong signed. The three findings trace the wedge from macro break to sector to capital allocation.

The lens applies whenever firms choose how aggressively to undertake an organisational transformation. The energy transition is the cleanest near term test. De Haas and Popov (2023) show that firms in market based economies transition faster to low carbon production through enhanced green patenting, and that equity financing reduces firm level carbon intensity. The empirical design generalises directly. Build an exogenous

US sector benchmark for the relevant technology, interact it with the local equity share, and predict the local cell's adoption rate. Cloud computing in the 2010s and artificial intelligence adoption in the 2020s are the natural follow ups. The policy implication is that bank heavy economies face a structural disadvantage in any technology shock requiring complementary intangibles, and the disadvantage compounds across waves.

The thesis unites several strands of the financial structure literature into a single mechanism. The owner of an innovating firm chooses how risky her project should be, and her financing contract changes which choice she makes. The single mechanism is the framework's strength. It does not assume other channels are absent. The cross sector empirical signature shows up in Brown et al. (2017) on R&D intensity and De Haas and Popov (2023) on green innovation. Other channels sit beside the mechanism. Bank regulation tightens the credit menu in ways the model takes as given. Existing bank exposure to non innovative incumbents creates a displacement friction the model assumes away. Garcia-Macia (2017) derives an intangibles collateral channel from the joint presence of debt and equity claims on the same firm, complementary to the across firm γ choice studied here. Draghi (2024) aggregates several mechanisms (bank screening, pension depth, regulatory fragmentation) the model does not separate.

The natural next steps follow from the limitations in Section 6. A general equilibrium model with γ_{\max} calibrated sector by sector and country by country would translate the industry coefficients into national accounts magnitudes and separate the financing channel from market fragmentation. Firm level data would identify the within sector γ heterogeneity directly. The micro foundations need more work. Adding effort moral hazard would test whether equity remains viable once the standard agency cost of equity finance is in play. The risky innovation lens names a channel. The data show the channel firing on the ICT case. The mechanism does not depend on the specific empirical example.

References

- Aghion, P., & Howitt, P. (1992). A model of growth through creative destruction. *Econometrica*, 60(2), 323–351. <https://doi.org/10.2307/2951599>
- Andrews, D. W. K. (1993). Tests for parameter instability and structural change with unknown change point. *Econometrica*, 61(4), 821–856. <https://doi.org/10.2307/2951764>
- Arcand, J. L., Berkes, E., & Panizza, U. (2015). Too much finance? *Journal of Economic Growth*, 20(2), 105–148. <https://doi.org/10.1007/s10887-015-9115-2>
- Autor, D., Dorn, D., Katz, L. F., Patterson, C., & Van Reenen, J. (2020). The Fall of the Labor Share and the Rise of Superstar Firms. *The Quarterly Journal of Economics*, 135(2), 645–709. <https://doi.org/10.1093/qje/qjaa004>
- Bajgar, M., Criscuolo, C., & Timmis, J. (2025). Intangibles and Industry Concentration: A Cross-Country Analysis. *Oxford Bulletin of Economics and Statistics*. <https://doi.org/10.1111/obes.12659>
- Bank for International Settlements. (2026). Total credit statistics [Private non-financial sector credit from domestic banks, % of GDP, adjusted series. Accessed: 2026-03-01].
- Bartelsman, E. J., Haltiwanger, J. C., & Scarpetta, S. (2013). Cross-country differences in productivity: The role of allocation and selection. *American Economic Review*, 103(1), 305–334. <https://doi.org/10.1257/aer.103.1.305>
- Basel Committee on Banking Supervision. (2017). *Basel iii: Finalising post-crisis reforms* (tech. rep.). Bank for International Settlements. Basel. <https://www.bis.org/bcbs/publ/d424.htm>
- Bergeaud, A. (2024). The past, present and future of european productivity. In *Ecb forum on central banking 2024: Monetary policy in an era of transformation*. Frankfurt am Main, European Central Bank.
- Bergeaud, A., Cette, G., & Lecat, R. (2016). Productivity trends in advanced countries between 1890 and 2012. *Review of Income and Wealth*, 62(3), 420–444. <https://doi.org/10.1111/roiw.12185>
- Biais, B., & Casamatta, C. (1999). Optimal leverage and aggregate investment. *Journal of Finance*, 54(4), 1291–1323. <https://doi.org/10.1111/0022-1082.00147>

- Black, B. S., & Gilson, R. J. (1998). Venture capital and the structure of capital markets: Banks versus stock markets. *Journal of Financial Economics*, 47(3), 243–277. [https://doi.org/10.1016/S0304-405X\(97\)00045-7](https://doi.org/10.1016/S0304-405X(97)00045-7)
- Bloom, N. (2009). The impact of uncertainty shocks. *Econometrica*, 77(3), 623–685. <https://doi.org/10.3982/ECTA6248>
- Bloom, N., Sadun, R., & Van Reenen, J. (2012). Americans Do IT Better: US Multinationals and the Productivity Miracle. *American Economic Review*, 102(1), 167–201. <https://doi.org/10.1257/aer.102.1.167>
- Bontadini, F., Corrado, C., Haskel, J., & Iommi, M. (2023). *Euklems and IntanProd: Industry productivity accounts with intangibles* (MPRA Paper No. 126218). Munich Personal RePEc Archive.
- Bresnahan, T. F., Brynjolfsson, E., & Hitt, L. M. (2002). Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence. *Quarterly Journal of Economics*, 117(1), 339–376. <https://doi.org/10.1162/003355302753399526>
- Brown, J. R., Martinsson, G., & Petersen, B. C. (2017). Stock markets, credit markets, and technology-led growth. *Journal of Financial Intermediation*, 32, 45–59. <https://doi.org/10.1016/j.jfi.2016.07.002>
- Brynjolfsson, E., & Hitt, L. M. (2003). Computing Productivity: Firm-Level Evidence. *The Review of Economics and Statistics*, 85(4), 793–808. <https://doi.org/10.1162/003465303772815736>
- Brynjolfsson, E., Rock, D., & Syverson, C. (2021). The Productivity J-Curve: How Intangibles Complement General Purpose Technologies. *American Economic Journal: Macroeconomics*, 13(1), 333–372. <https://doi.org/10.1257/mac.20180386>
- Cameron, A. C., Gelbach, J. B., & Miller, D. L. (2008). Bootstrap-based improvements for inference with clustered errors. *Review of Economics and Statistics*, 90(3), 414–427. <https://doi.org/10.1162/rest.90.3.414>
- Corrado, C., Hulten, C., & Sichel, D. (2005). Measuring Capital and Technology: An Expanded Framework. In *Measuring Capital in the New Economy* (pp. 11–46). University of Chicago Press.
- De Haas, R., & Popov, A. (2023). Finance and green growth. *Economic Journal*, 133(650), 637–668. <https://doi.org/10.1093/ej/ueac081>

- De Ridder, M. (2024). Market Power and Innovation in the Intangible Economy. *American Economic Review*, 114(1), 199–251. <https://doi.org/10.1257/aer.20201079>
- Demirgüç-Kunt, A., Feyen, E., & Levine, R. (2013). The evolving importance of banks and securities markets. *The World Bank Economic Review*, 27(3), 476–490. <https://doi.org/10.1093/wber/lhs022>
- Draghi, M. (2024). *The Future of European Competitiveness – Part A: A Competitiveness Strategy for Europe* (tech. rep.). European Commission.
- Feenstra, R. C., Inklaar, R., & Timmer, M. P. (2015). The next generation of the Penn World Table [TFP series RTFPNAUSA632NRUG via FRED]. *American Economic Review*, 105(10), 3150–3182. <https://doi.org/10.1257/aer.20130954>
- Garcia-Macia, D. (2017). *The financing of ideas and the great deviation* (IMF Working Paper WP/17/176). International Monetary Fund.
- Halek, M., & Eisenhauer, J. G. (2001). Demography of risk aversion. *Journal of Risk and Insurance*, 68(1), 1–24. <https://doi.org/10.2307/2678130>
- Inklaar, R., Timmer, M. P., & van Ark, B. (2008). Market Services Productivity across Europe and the US. *Economic Policy*, 23(53), 139–194. <https://doi.org/10.1111/j.1468-0327.2007.00194.x>
- Innes, R. D. (1990). Limited liability and incentive contracting with ex-ante action choices. *Journal of Economic Theory*, 52(1), 45–67. [https://doi.org/10.1016/0022-0531\(90\)90066-S](https://doi.org/10.1016/0022-0531(90)90066-S)
- Jorgenson, D. W., & Griliches, Z. (1967). The explanation of productivity change. *Review of Economic Studies*, 34(3), 249–283. <https://doi.org/10.2307/2296675>
- Kermani, A., & Ma, Y. (2023). Asset specificity of nonfinancial firms. *Quarterly Journal of Economics*, 138(1), 205–264.
- Klette, T. J., & Kortum, S. (2004). Innovating firms and aggregate innovation. *Journal of Political Economy*, 112(5), 986–1018. <https://doi.org/10.1086/422563>
- Kuvshinov, D., & Zimmermann, K. (2022). The big bang: Stock market capitalization in the long run. *Journal of Financial Economics*, 145(2), 527–552. <https://doi.org/10.1016/j.jfineco.2021.09.008>

- La Porta, R., Lopez-de-Silanes, F., Shleifer, A., & Vishny, R. W. (1998). Law and finance. *Journal of Political Economy*, 106(6), 1113–1155. <https://doi.org/10.1086/250042>
- Langfield, S., & Pagano, M. (2016). Bank bias in Europe. *Economic Policy*, 31(85), 51–106. <https://doi.org/10.1093/epolic/eiv019>
- Levine, R. (2002). Bank-based or market-based financial systems: Which is better? *Journal of Financial Intermediation*, 11(4), 398–428. <https://doi.org/10.1006/jfin.2002.0341>
- Lian, C., & Ma, Y. (2021). Anatomy of corporate borrowing constraints. *Quarterly Journal of Economics*, 136(1), 229–291. <https://doi.org/10.1093/qje/qjaa030>
- Nonnis, A., Roth, F., & Bounfour, A. (2024). *Intangible Capital in France and Germany: Is there a Measurement Issue?* (Discussion Paper No. 18). University of Hamburg, Department of Economics. <https://hdl.handle.net/10419/305807>
- Perotti, E. C., & Schwiabacher, A. (2009). The political origin of pension funding. *Journal of Financial Intermediation*, 18(4), 384–404. <https://doi.org/10.1016/j.jfi.2008.11.001>
- Rajan, R. G., & Zingales, L. (2003). The great reversals: The politics of financial development in the twentieth century. *Journal of Financial Economics*, 69(1), 5–50. [https://doi.org/10.1016/S0304-405X\(03\)00125-9](https://doi.org/10.1016/S0304-405X(03)00125-9)
- Rothschild, M., & Stiglitz, J. E. (1970). Increasing risk: I. A definition. *Journal of Economic Theory*, 2(3), 225–243. [https://doi.org/10.1016/0022-0531\(70\)90038-4](https://doi.org/10.1016/0022-0531(70)90038-4)
- Rousseau, P. L., & Wachtel, P. (2011). What is happening to the impact of financial deepening on economic growth? *Economic Inquiry*, 49(1), 276–288. <https://doi.org/10.1111/j.1465-7295.2009.00197.x>
- Scharfstein, D. S. (2018). Presidential address: Pension policy and the financial system. *The Journal of Finance*, 73(4), 1463–1512. <https://doi.org/10.1111/jofi.12710>
- Schreyer, P. (2002). Computer price indices and international growth and productivity comparisons. *Review of Income and Wealth*, 48(1), 15–31. <https://doi.org/10.1111/1475-4991.00038>
- Schumpeter, J. A. (1934). *The theory of economic development: An inquiry into profits, capital, credit, interest, and the business cycle* (R. Opie, Trans.) [Translation of *Theorie der wirtschaftlichen Entwicklung*, 1911.]. Cambridge, MA, Harvard University Press.
- Schumpeter, J. A. (1942). *Capitalism, socialism and democracy*. New York, Harper & Row.

- Stiglitz, J. E., & Weiss, A. (1981). Credit rationing in markets with imperfect information. *American Economic Review*, 71(3), 393–410.
- Syverson, C. (2011). What Determines Productivity? *Journal of Economic Literature*, 49(2), 326–365. <https://doi.org/10.1257/jel.49.2.326>
- van Ark, B., O'Mahony, M., & Timmer, M. P. (2008). The Productivity Gap between Europe and the United States: Trends and Causes. *Journal of Economic Perspectives*, 22(1), 25–44. <https://doi.org/10.1257/jep.22.1.25>
- World Bank. (2022). Global financial development database [August 2022 release. Accessed: 2026-02-16].