

THE PRICE OF POWER

The Systemic Social Dividend

*A prudential proposal for measuring systemic risk
and financing the transition*

Giammario Battaglia

Copyright © 2026 Giammario Battaglia
All rights reserved.

First English edition — 2026

Methodological note: this volume systematically distinguishes verified data, methodological proposals and simulations. The company cases serve an analytical function and do not constitute findings of liability.

Transparency note: generative artificial intelligence tools were used in preparing and translating this manuscript to support drafting, revision, organisation and consistency checking. Final responsibility for the content remains the author's.

CONTENTS

PROLOGUE — When power changes its nature	1
PART I — THE NEW POWER	5
CHAPTER 1 — When a company becomes an infrastructure	7
CHAPTER 2 — The society of dependency	11
CHAPTER 3 — Who pays for transitions?	15
PART II — NOT A TAX ON SUCCESS	19
CHAPTER 4 — Why taxing revenue is the wrong question	21
CHAPTER 5 — The function comes before the entity	25
CHAPTER 6 — The risk that remains	29
PART III — MEASURING POWER	33
CHAPTER 7 — Six dimensions of systemic importance	35
CHAPTER 8 — The DSS Data Return	39
CHAPTER 9 — R0, R1, R2, R3, R4	43
PART IV — PAYING LESS BY PREVENTING MORE	47
CHAPTER 10 — The way out	49
CHAPTER 11 — The price of non-substitutability	53
CHAPTER 12 — The contribution	57

PART V — WHAT DOES IT COST TO PROTECT A SOCIETY?	61
CHAPTER 13 — Ten, twenty-five or forty billion?	63
CHAPTER 14 — The cost of the right to transition	67
CHAPTER 15 — The paradox of a Fund that wants to collect less	71
PART VI — THE MODEL MEETS THE WORLD	75
CHAPTER 16 — Microsoft Azure: the cloud test	77
CHAPTER 17 — OpenAI: the artificial intelligence test	81
CHAPTER 18 — NVIDIA: the test of upstream power	83
CHAPTER 19 — Three functions, three forms of power	85
PART VII — EUROPEAN LAW, PRUDENTIAL PRECEDENTS AND VALIDATION	89
CHAPTER 20 — The DSS and European law	91
CHAPTER 21 — Three prudential precedents: DORA, G-SIB and the SRF	95
CHAPTER 22 — From risk to euros: expected social loss	101
CHAPTER 23 — Attributing the need: causation, F_c and the hybrid contribution	105
CHAPTER 24 — Who really pays? Incidence and pass-through of the contribution	109
CHAPTER 25 — Calibrating the DSS: weights, aggregation and robustness	113
PART VIII — WORK, ADDITIONALITY AND EUROPEAN SOVEREIGNTY	119
CHAPTER 26 — The L dimension: measuring tasks, not predicting unemployment	121
CHAPTER 27 — Additionality: the EGF, the Just Transition Fund and the DSS	125
CHAPTER 28 — The European door: legal basis, the SRF and the model's limits	129
CHAPTER 29 — ASML: the European test of the upstream bottleneck	133

CHAPTER 30 — Geopolitical neutrality: GATS, Section 301 and the risk of an anti-US tax	137
 PART IX — GOVERNANCE, DUE PROCESS AND CONTROL	 141
CHAPTER 31 — Who decides, who manages, who supervises	143
CHAPTER 32 — The right to contest the algorithm	147
 PART X — WHO IS ENTITLED TO THE DIVIDEND	 149
CHAPTER 33 — Who receives the DSS: need and the geography of transition	151
CHAPTER 34 — The INDS: measuring need	153
CHAPTER 35 — How the DSS reaches the individual	157
 PART XI — FROM THEORY TO IMPLEMENTATION	 161
CHAPTER 36 — From model to law: who has to do what, and why	163
CHAPTER 37 — Pilot first, extend later: the roadmap	167
CHAPTER 38 — From social dividend to a new compact	169
EPILOGUE — Power must be able to diminish	173
FINAL APPARATUS — Note on sources and method	175
AUTHOR'S NOTE	179
TECHNICAL APPENDICES — APPENDIX A — Summary sheet of the DSS model	181
THE AUTHOR	187

PROLOGUE

When power changes its nature

For a long time we thought about economic power through relatively simple categories.

Whoever owned more factories produced more. Whoever had more capital could invest more. Whoever controlled a larger market share could exert more influence over prices, competitors and consumers.

The digital economy has steadily altered that geometry.

Today a company can wield enormous power not only because it owns a great deal, sells a great deal or earns a great deal, but because millions of people, businesses and institutions depend on a function it controls.

A cloud infrastructure can host the essential services of thousands of organisations. An operating system can set the terms of access to an entire digital ecosystem. A platform can become the place where firms and professionals meet their customers. A producer of compute capacity can supply an input that is hard to replace in the development of artificial intelligence. An AI model can work its way into the processes through which we write, code, analyse documents, make decisions and organise work.

The question raised by this transformation is not simply economic. It is institutional.

What happens when a private function becomes so important that the consequences of its controller's decisions can propagate across a significant part of society?

The traditional answer is usually to look at the size of the firm: how much it turns over, how much it earns, what it is worth on the stock market, how much market share it holds. These are important questions. But they are not yet the right question.

One company can be enormous and easily replaceable. Another can be smaller and control a chokepoint on which an entire supply chain depends.

Size measures size. Systemic importance measures dependency.

It is from that distinction that this book is born.

This is not a book against large companies. A company's growth is not a fault. Innovation is not behaviour to be penalised. Profit is not in itself a measure of the harm done to society. And technological success cannot automatically become the trigger for an economic obligation.

The proposal developed in these pages follows a different logic. It does not ask how much a company owns, but which function it performs. It does not ask how much it earns, but what dependencies that function creates. It does not ask how much we can extract, but how much risk can be eliminated before what remains has to be mutualised.

The shift is a radical one. Wealth may be an indicator of an operator's economic capacity. It is not, on its own, the reason why that operator should contribute. The reason must be sought in residual systemic risk.

THE PRICE OF POWER

Every innovation produces risks. The problem is not to stop change, but to decide who bears the cost of transition when change happens faster than people and institutions can adapt.

The core principle of the Systemic Social Dividend can be stated in a few words: whoever performs a systemic function must first reduce the dependencies and risks that function creates. Only the risk that cannot reasonably be eliminated can justify a mutualised contribution.

The correct sequence, then, is: function, dependency, risk, prevention, residual risk, contribution.

From this follows an apparently paradoxical consequence. A well-functioning DSS system ought to be pleased to collect less. If an operator increases portability, reduces lock-in, builds operational continuity and supports the transformation of skills, residual risk falls. And with it the contribution must fall too — at least in theory, all the way to zero.

A tax has an interest in a taxable base. The DSS has an interest in shrinking its own risk base.

This is also why the DSS should be read, before it is read as a redistributive mechanism, as an infrastructure of resilience. Portability, operational continuity, interoperability, the capacity to exit and diversification all reduce risk before the Fund ever has to step in.

In the language of the European institutions, this logic intersects with a question of autonomy of choice: a system is more sovereign not when it closes its markets, but when firms, public administrations and citizens can switch provider, absorb a shock and carry on operating without depending on a single point of failure.

The social component comes second, but it is not secondary: when prevention is not enough, the transition must be financeable without turning economic success into a fault and without dumping the entire cost on general taxation.

Why, then, *The Price of Power*? Because systemic power can arise out of millions of individually rational choices. A better service wins users; the more users adopt it, the more developers build compatible services; the more firms build their processes on it; the more data gets organised according to that architecture; the more staff are trained on it; and the higher the cost of leaving becomes.

The price of power is not a punishment for those who hold it. It is the responsibility that arises when the exercise of a private function produces consequences reaching beyond the contractual relationship between company and customer.

This book therefore tackles the most uncomfortable question: how is it actually done? How do you establish that a function is systemic? How do you measure dependency? How much weight should difficulty of substitution carry? How do you stop an authority from simply declaring a company “systemic” on the basis of a political judgement? How can the company contest that decision? How do you measure a mitigation? And, only at the end, how much should it contribute?

To answer, we will try to turn qualitative principles into observable indicators: scale, interconnection, non-substitutability, control of data and infrastructure, propagation, and the transformation of work. We will build a DSS Data Return, risk classes and mitigation procedures. And we will put the model to work on concrete cases — not to

assert that any particular company should pay, but to test whether such an obligation could be derived from data, criteria, adversarial process and residual risk.

PART I

THE NEW POWER

CHAPTER 1

When a company becomes an infrastructure

There is a difference between selling something to a great many people and being something a great many people depend on.

That difference is the starting point.

Picture two companies. The first sells a hundred billion euro of goods a year. It owns thousands of retail outlets, has millions of customers and generates large profits. But if it disappeared tomorrow, consumers could buy comparable products from any number of competitors.

The second turns over far less. Yet it supplies a technology embedded in the systems of thousands of businesses. Replacing it would take months, along with rewritten software, data migration, staff retraining and operational downtime.

Which of the two is more systemic?

Revenue would point to the first. Dependency points to the second. And it is the second measure that concerns the DSS.

1.1 Large does not mean systemic

One of the easiest mistakes is to treat *large*, *dominant*, *strategic*, *essential* and *systemic* as synonyms. They are not.

A company can be large because it serves a great many customers. It can be dominant because it holds a very high share of a market. It can be strategic because it operates in a sector that matters for economic or geopolitical autonomy. It can be essential because it supplies a service society particularly needs.

It becomes systemic when the disruption, transformation or difficulty of replacing the function it performs can produce significant effects beyond its direct relationship with its own customers.

Systemic importance therefore contains an element of propagation. In conceptual form: $\text{systemic importance} = \text{dependency} + \text{non-substitutability} + \text{propagation}$. This is not yet the quantitative formula we will use later on. It is a working definition.

1.2 The function comes before the company

A second choice follows from this. The DSS should not classify a company in the round. It should classify a function.

Microsoft is a group; Azure is an identifiable economic and infrastructural function. Alphabet is a group; Google Search, Android, Google Cloud and certain AI functions have different characteristics. Amazon spans e-commerce, logistics, advertising and cloud. NVIDIA makes several categories of hardware and software, but what may take

THE PRICE OF POWER

on systemic significance for AI is a specific combination of compute capacity, accelerators and software ecosystem.

The legal consequence matters. An authority should not be able simply to state that a company is systemic. It should be required to identify the function, the operator performing it, the relevant market or perimeter, and the characteristics of systemic importance it intends to demonstrate.

1.3 The right to leave

Here we meet a notion that will run through the whole book: substitutability.

A market is genuinely contestable not merely when competitors exist. It must be possible to reach them.

If five alternatives exist but leaving the incumbent supplier takes eighteen months, enormous investment, data loss, rewritten applications and operational risk, nominal competition may be far greater than effective competition.

The question becomes: what does it cost to leave?

The cost is not only monetary. There is a cost in time, in organisation, in information, in technology, in operations and in learning. This is why the DSS will attach great weight to the time needed to switch, the total economic cost of migration and the number of genuinely equivalent alternatives.

Freedom of choice on entry matters. In a systemic economy, freedom to exit matters just as much.

1.4 From freedom of contract to collective dependency

No company necessarily forces millions of customers to use its service. The problem can arise from precisely the opposite: millions of people choose it freely.

The better service wins users. The more users adopt it, the more developers build compatible services. The more firms build their processes on it. The more data is organised according to that architecture. The more staff are trained on it. And the higher the cost of leaving becomes.

The sequence can run: quality, adoption, ecosystem, dependency, lock-in.

There need be no abuse at any step. And yet a systemic position may emerge at the end of it. This is precisely why the DSS cannot be built as a sanctioning regime. You do not have to prove fault in order to prove dependency.

1.5 The first principle of the new model

We can therefore state the first principle we will use throughout the book:

Systemic power is not presumed from wealth. It is demonstrated by dependency.

Almost everything else follows from that proposition.

If dependency is what we have to measure, revenue and market capitalisation are not enough. New data is needed. We need to know how many parties depend on the function, how interconnected they are, how long it takes them to leave it, what leaving costs, how many alternatives genuinely exist, how quickly a disruption propagates, which infrastructures are controlled, and what transformations the function produces in work.

This is where the philosophical problem of power becomes a quantitative problem. And it is precisely here that the work of this book begins.

CHAPTER 2

The society of dependency

Dependency is one of the invisible infrastructures of contemporary society.

We notice it when something is missing. As long as a platform works, a cloud service responds, a payment system authorises transactions or an AI model returns results, the relationship between user and technology looks unremarkable. Only at the moment of disruption, of a sudden cost increase, of a change in terms or of the need to migrate does the underlying structure come into view.

A society can be technologically advanced and, at the same time, profoundly dependent on a limited number of private functions. This is not necessarily an anomaly. Specialisation produces efficiency. Economies of scale cut costs. Common standards let millions of parties interact. The cloud spares many firms the need to build their own infrastructure. Platforms lower the cost of matching supply and demand. Artificial intelligence puts within reach capabilities that only a few years ago required far greater skills and resources.

The problem arises when the efficiency obtained by concentrating a function is not matched by equally efficient instruments for exit, substitution and continuity.

Systemic dependency is therefore not the opposite of efficiency. It can be a consequence of it.

2.1 There is no single kind of dependency

To build a credible regulatory model, one mistake has to be avoided first: treating all technological dependencies as though they were alike.

Depending on a cloud infrastructure is different from depending on an operating system. Depending on an AI model is different from depending on the compute capacity needed to train or run it. Depending on a social platform is different from depending on a payment network.

We can distinguish at least five forms.

The first is **infrastructural dependency**. It occurs when essential data, applications or processes sit on infrastructures that cannot be replicated or replaced quickly.

The second is **technological dependency**. It arises when software, formats, APIs, models, development tools or proprietary architectures make migration expensive.

The third is **network dependency**. The value of the service grows with the number of others using it. Leaving therefore means losing not just a technology but a network of relationships.

The fourth is **informational dependency**. Data, histories, configurations, reputation, preferences or knowledge accumulated over time can be hard to transfer in full.

The fifth is **organisational dependency**. A company may technically be able to change supplier, yet have procedures, staff, contracts, controls and skills built around a particular ecosystem.

These forms can overlap. And it is precisely their combination that turns a commercial relationship into a potentially systemic one.

2.2 Cloud: when infrastructure becomes a service

Cloud is one of the clearest examples.

For many organisations it is no longer simply remote space in which to keep files. It is a bundle of compute capacity, databases, identity, networking, security, development tools, data analytics and AI services.

The more an organisation uses specific components of one ecosystem, the more complex the question "can I switch supplier?" becomes.

The formal answer may be yes. The economic answer may be very different.

You need to know how long migration takes, which applications have to be rewritten, which data has to be transferred, which certifications have to be repeated, which people have to be retrained and what risk of disruption has to be accepted.

For the DSS, then, measuring a cloud provider's market share is not enough. What has to be measured is the real cost of leaving.

2.3 Artificial intelligence: depending on a capability

Artificial intelligence introduces a different form of dependency again.

A company can use a model through an interface, embed it in its own products via API, hand it parts of an internal process, build agents that interact with corporate data and software, or adapt entire workflows to its availability.

In these cases the dependency is not only on the model. It is on the organisational capability built around the model.

The regulatory point is not to prevent that integration. It is to establish whether the organisation retains a reasonable possibility of replacing the technology without suffering disproportionate harm.

This makes portability central not only for data but for processes: structured prompts, evaluations, tools, integrations, configurations, application memory, agents and control procedures.

AI thus shows why substitutability cannot be measured with a simple list of competitors. Two models can compete in the market and still not be substitutable, within a given process, without substantial re-engineering.

2.4 Compute capacity: power upstream

There is also a less visible form of power, because it operates upstream.

Advanced AI systems require compute capacity, accelerators, memory, high-performance networking, software and data-centre infrastructure. Part of systemic importance may therefore be concentrated not in the service used by the end consumer, but in the inputs that make that service possible at all.

This step is essential.

If we measured systemic importance solely through the number of end users, we would risk overlooking operators with almost no direct relationship with the public who nonetheless control an input that thousands of firms would struggle to replace.

The DSS must therefore be able to look both downstream and upstream: at the platforms that organise access, and at the infrastructures that make that access possible.

2.5 Platforms and operating systems: the power of the ecosystem

Platforms add a further element: the ecosystem.

When users, developers, businesses, advertisers, service providers and content producers converge in the same environment, the value of the function no longer depends on its technical quality alone. It depends on the presence of everyone else.

In these markets, leaving can mean losing a network. For a developer, for instance, the alternative is not simply another distribution channel: it is another way of reaching the same set of users.

This is why substitutability has to be functional. A technology is an alternative only if it delivers, within reasonable limits, the same economic and social function.

The DSS should therefore ask not merely "is there a competitor?" but "can the dependent party move its activity to that competitor without losing a disproportionate share of the very function it is trying to replace?"

2.6 Payments, identity and invisible services

Some functions become systemic precisely because we stop noticing them.

Digital payments, authentication, identity, certificates, network services, name resolution, security and other intermediate layers can be embedded in thousands of everyday processes.

When they work, they are all but invisible. When they fail, the effect can spread fast.

This characteristic introduces a further dimension: speed of propagation. A risk that produces effects over months is different from a risk that halts thousands of transactions in minutes.

For this reason, in the model we will develop, propagation will be a dimension of systemic importance in its own right.

2.7 Dependency is not a fault

Talking about dependency may suggest a negative judgement on the operator that created it. That would be a mistake.

A function can become systemic precisely because it is reliable, innovative, good value and widely adopted.

The DSS must be able to distinguish dependency from abuse.

Abuse belongs to competition law and to the other forms of liability the legal order already provides. Systemic dependency, by contrast, belongs to prudential risk management.

An operator may have broken no rule and still perform a function whose unavailability would have significant collective consequences.

The response should not be punitive. It should be preventive: portability, interoperability, continuity, redundancy, transparency about dependencies and the capacity to exit.

2.8 The society of dependency

The digital transformation thus leads us towards a paradox.

The more useful a technology is, the more widely it can spread. The more widely it spreads, the more worthwhile it becomes to build processes around it. The more specialised those processes become, the higher the cost of replacement can grow.

The problem is not to halt that dynamic. It is to prevent private efficiency from producing collective fragility with no compensating instruments.

2.9 The second principle

Efficiency that creates dependency must be accompanied by the capacity to reduce, manage and reverse that dependency.

This is the step that turns the DSS from a redistributive idea into a proposal in industrial and prudential policy.

CHAPTER 3

Who pays for transitions?

Every major economic transition produces winners, losers and a vast middle ground of people and organisations who have to adapt.

The industrial revolution moved populations, trades and capital. Electrification reshaped factories and cities. Automation transformed manufacturing work. The internet changed commerce, information and communication.

Artificial intelligence adds one characteristic: the speed with which it can enter cognitive activities that are already organised digitally.

Software can be distributed far more quickly than a factory or a railway can be built. A capability can be made available in many countries simultaneously. A company can change a process without physically replacing an entire production infrastructure.

This does not mean AI will automatically produce mass unemployment. It means the speed of the transformation can itself become a risk factor.

The political problem is not to predict with certainty how many jobs will disappear. It is to build institutions capable of acting when transformation outruns the capacity to adapt.

3.1 Exposure does not mean replacement

One of the most important distinctions concerns work.

To say that an occupation is exposed to artificial intelligence is not to say that the occupation will disappear.

A job is made up of different tasks. Some can be automated, others accelerated, and others still may gain value precisely because of the technology.

A serious methodology must therefore distinguish at least between exposed work, the share of tasks that can be transformed, and actual replacement.

That distinction guards against two opposite errors: minimising every consequence because “new jobs will emerge”, or assuming that every technically automatable activity will necessarily be eliminated.

The DSS does not need to predict the future with absolute precision. It needs to measure transition risk prudently enough to fund instruments that are reversible and verifiable.

3.2 The cost of adaptation

When a technology changes an occupation, the cost of adaptation can take many forms.

New skills may have to be acquired. Working hours may have to be reduced temporarily to attend training. A company may have to reconfigure procedures and roles.

THE PRICE OF POWER

A worker may have to live through a period of lower income. An area may lose a concentration of activity before other activity emerges.

These costs are not compensation for harm.

They are transition costs.

The distinction matters because it changes the purpose of the benefit. The DSS should not pay someone because the technology exists. It should fund, within predetermined criteria, what allows people and territories to get through the transformation without being shut out of the capacity to take part in the economy.

3.3 The current asymmetry

Absent a dedicated mechanism, part of the cost of transition tends to be spread across a range of parties.

Firms invest to adapt their processes. Workers invest time and resources in training. Households absorb periods of lower income. The state funds active labour market policies, welfare and services. Regions bear the indirect effects of lost activity.

The benefits of innovation, by contrast, can concentrate more quickly.

This asymmetry does not show that anyone has done anything unlawful. It shows that markets can distribute benefits and costs on different timescales and through different mechanisms.

The DSS intervenes precisely at that temporal fracture: making resources available at the moment the transition creates the need, without waiting for the social damage to become permanent.

3.4 Why general taxation is not enough

One possible answer is simple: these costs should be borne by the state out of general taxation.

It is a legitimate answer, but an incomplete one.

General taxation funds a plurality of public functions and does not necessarily create any link between the systemic transformation that generates the risk and the resources allocated to managing it.

The DSS instead proposes a separate, ring-fenced fund. The logic resembles that of other prudential mechanisms: identify a risk, measure it, reduce it, and set resources aside for the part that remains.

This ring-fencing is also essential to the system's legitimacy. If the contribution flowed indistinctly into the general budget, it would be far harder to distinguish it from an ordinary tax and to demonstrate the relationship between risk and benefit.

3.5 Who should contribute?

The most intuitive answer would be: the companies that gain most from the transformation.

But that formula is too loose.

A company can make enormous profits without performing a systemic function. Another can have thinner margins and control a function on which entire supply chains depend.

This is why the DSS does not start from profit. It starts from function and risk.

Economic capacity remains relevant, but it enters at a different point: in the affordability of the contribution and in the caps. It must not stand in for proof of systemic importance.

3.6 The third principle

No one contributes because they are wealthy as such. Those who contribute, within sustainable limits, are those performing a systemic function — for the share of residual risk that function continues to generate.

3.7 How much should be contributed?

Here too the traditional reasoning has to be inverted.

You do not start from the rate.

You start from the need.

First you determine what it costs to fund a predetermined set of benefits: training, transition income, mobility, reskilling and territorial measures.

Then you determine which systemic sectors are causally linked to the transformations under consideration.

Finally, the sector's share of the need is apportioned among operators according to residual risk.

The structure can be set out like this:

benefit funding need → sector share → operators' residual risk → individual contribution → cap.

This stops the system becoming a revenue machine. The overall amount does not grow because company profits grow. It grows only if a documented, fundable need grows.

3.8 The right to transition

If the contribution is tied to risk, the benefit must be tied to a verifiable need.

The DSS is not a universal dividend paid out indiscriminately to everyone.

It is a right conditional on transition.

A person does not receive because a large company has paid. They receive because they are in a defined situation of vulnerability or transformation, and because a specific benefit can increase their capacity to adapt.

This separation is fundamental.

On one side, the operator's systemic responsibility is measured.

On the other, the need of the person or the territory is measured.

The two algorithms must not contaminate each other. A region's vulnerability cannot be used to inflate a particular company's contribution arbitrarily. And a company's wealth cannot be used to invent an entitlement where no need exists.

3.9 The Fund as insurance for the transformation

Perhaps the simplest way to understand the DSS is as a form of social insurance for the transformation.

THE PRICE OF POWER

Not private insurance, because the purpose is collective and systemic operators' participation derives from law.

Not an ordinary tax, because the contribution is linked to a measured risk, to a ring-fenced need and to preventive measures that can reduce it.

Not compensation, because it does not necessarily require a wrong, nor assume that every beneficiary has suffered harm attributable to a single operator.

It is a prudential mutualised mechanism.

First you try to avoid the risk. Then you fund what it was not possible to avoid.

3.10 The question that remains

At this point, though, the decisive difficulty emerges.

If we do not want simply to tax revenue, and we do not want to leave the authority a discretionary judgement, then we have to be able to measure systemic importance.

We have to turn words like scale, dependency, lock-in, propagation and the transformation of work into observable quantities.

It is the most delicate step in the entire project.

And it is why, before talking about rates and billions, the next step will be to build the method by which it is decided who enters the system and in which risk class.

Only then will we be able to return to the question in the title: what is the price of power?

PART II

NOT A TAX ON SUCCESS

CHAPTER 4

Why taxing revenue is the wrong question

When public debate encounters very large companies, the most immediate temptation is to turn their economic size into a base for levying.

The argument looks simple: if technological transformation produces social costs and some companies reap extraordinary benefits from it, part of their revenue or profits can be earmarked to fund those costs.

It is an understandable logic. But it is not the logic of the Systemic Social Dividend. The problem is not a moral one. It is methodological.

Revenue tells you how much a company sells. Profit tells you how much economic value it retains after costs. Market capitalisation expresses the market's expectations about the company's value. None of these quantities measures, on its own, how hard a function is to replace, how many activities depend on it, or how quickly a malfunction could propagate.

If we use wealth as a shortcut for measuring risk, we end up conflating two different questions: who has economic capacity, and who generates systemic dependency.

The DSS has to keep them apart.

4.1 Wealth, ability to pay and risk

Economic capacity remains important. A responsible system cannot impose on an operator a contribution incompatible with the viability of its business. Revenue, margins and other indicators can therefore help set a cap.

But the cap answers the question "how much can sustainably be asked?" It does not answer the question "why should anything be asked at all?"

The second question requires a relationship with risk.

This distinction avoids a frequent error: using revenue simultaneously as proof of systemic importance and as the measure of the contribution. The methodology would then become circular. The company would pay a lot because it is large, and would be deemed systemic because it pays a lot.

In the DSS the sequence has to run the other way: first systemic importance is proved; then residual risk is measured; finally the economic sustainability of the contribution is checked.

4.2 Why a percentage of revenue is not enough

Picture two operators with the same European revenue.

THE PRICE OF POWER

The first provides an easily substitutable service. Customers can migrate in a few days, data is portable, numerous equivalent alternatives exist and a disruption produces limited effects.

The second provides a function embedded in the processes of thousands of organisations. Migration takes months, alternatives are few and a disruption can propagate to essential services.

Applying the same rate to both would be simple. But it would not be consistent with the logic of risk.

The DSS must be able to arrive at different contributions even where revenue is identical. And it must be able to arrive, in extreme cases, at a very low contribution for an enormous company if the function is highly substitutable and residual risk is low.

4.3 Need comes before revenue

Another difference concerns the overall amount.

In an ordinary tax, revenue comes from the taxable base multiplied by the rate. In the DSS, the starting point is instead the cost of the benefits the system decides to guarantee.

Training, transition income, mobility, redeployment, occupational reskilling and territorial measures have to be defined first.

Only then is the funding need calculated.

Benefit funding need → share attributed to the sector → apportionment by residual risk → individual cap.

This ordering prevents the Fund from growing simply because companies become more valuable. The Fund should grow only when a documented need grows, or when society decides, through legislation, to expand the benefits.

4.4 A contribution that can fall

The deepest difference from a tax emerges when the operator reduces risk.

If the contribution were a percentage of revenue, improving portability and interoperability would not necessarily change the obligation.

In the DSS the opposite has to happen.

If migration time falls, if the cost of exit drops, if the number of effective alternatives rises, if operational continuity improves and if potential propagation is reduced, the residual score must fall.

And with it, the contributory burden must fall too.

The system thus creates an economic incentive to prevent. It does not reward a statement of social responsibility. It rewards a measurable result.

4.5 Not a tax on Big Tech

The phrase “tax on Big Tech” works well in communication but is conceptually inadequate as a description of the DSS.

First, because not all large technology companies must necessarily be systemic.

Second, because a systemic function can also be performed by an operator that does not fit the traditional image of Big Tech.

Third, because the contribution does not arise from the technology sector as such. The model can be applied, with appropriate parameters, to any sector in which a private function generates significant systemic dependencies.

The DSS does not identify an economic enemy. It identifies a category of risk.

4.6 The fourth principle

We can therefore set down another principle:

Economic capacity limits the contribution; residual risk justifies it and determines how it is apportioned.

That sentence sharply separates the DSS from a wealth tax, from a digital tax and from a tax on revenue.

Wealth stays in the model, but in the right place: as a sustainability constraint, not as a proxy for systemic responsibility.

CHAPTER 5

The function comes before the entity

Institutions normally regulate legal entities: companies, banks, intermediaries, operators, platforms. This is inevitable, because obligations and liabilities have to be attributed to someone.

But systemic importance arises before legal attribution.

It arises from a function.

The distinction looks abstract. In practice it is one of the model's main safeguards against arbitrariness.

If the authority started from the company's name, it would risk searching afterwards for reasons to justify the designation. If it starts from the function, it has to demonstrate which characteristics make that function significant, and only then identify who performs it.

5.1 Companies and functions do not coincide

A large group may carry on dozens of activities at once.

Some may be systemic, others not. Some may be highly substitutable, others deeply integrated. Some may produce mainly commercial effects, others bear on the continuity of infrastructures and processes.

Assigning a single score to the whole group would erase these differences.

Designation should therefore be framed as a pair: operator + systemic function.

Not "Amazon", but — if the data shows it — a particular cloud function performed by the group.

Not "Microsoft", but the Azure function, Microsoft 365, or another function identified in the investigation.

Not "NVIDIA", but the possible AI compute infrastructure function identified within a defined technical and economic perimeter.

5.2 The perimeter of the function

Defining a function means establishing at least four elements.

First: what it does.

Second: who it does it for.

Third: in which territory or market it produces the relevant effects.

Fourth: which components are indispensable for that function to be performed at all.

The perimeter must be neither too wide nor artificially narrow.

If it is too wide, we attribute to the function risks generated by different activities. If it is too narrow, we can fragment a real dependency until it becomes invisible.

The definition must therefore be reasoned and open to challenge.

5.3 The problem of conglomerates

Digital conglomerates make the problem more delicate still.

A single company may control infrastructure, software, identity, distribution, data and AI tools. The functions may reinforce one another.

The DSS must be able to observe these interactions without automatically turning the whole group into a single regulatory object.

The solution is a two-level assessment.

At the first level, the individual function is measured.

At the second, it is assessed whether intra-group dependencies exist that increase interconnection, non-substitutability or propagation.

In this way the conglomerate's integration enters the score only when it produces a demonstrable effect.

5.4 Avoiding double attribution

Separating by function creates the opposite risk, however: counting the same phenomenon twice.

If the same risk is attributed both to the cloud function and to the AI function of the same operator, the system could artificially duplicate the contribution.

A consolidation rule is therefore needed.

Each risk component must be attributed once only, or apportioned under a transparent methodology when it contributes to more than one function.

The principle is simple: a function may be assessed separately, but the same euro of need cannot be funded twice out of the same risk.

5.5 The function as a safeguard for the company

This method does not only make the contribution more precise. It also protects the operator.

The company can contest the perimeter of the function.

It can show that a service identified by the authority is severable.

It can prove that certain dependencies belong to a different market.

It can show that a risk attributed to the function in fact stems from another operator or from the customer's own organisational choices.

The function thus becomes the subject matter of the adversarial process.

The proceedings do not debate the company's reputation. They debate technical and economic facts.

5.6 From the name to the evidence

This approach also changes the language.

A serious decision should not assert that a company "has too much power".

It should state, for example, that a given function shows a high share of downstream dependencies, a median switching time above a threshold, a limited number of equivalent alternatives and a documented capacity for propagation.

Power stops being a political judgement and becomes a finding of fact.
That is the move from the name to the evidence.

5.7 The fifth principle

We can therefore state the fifth principle:

The DSS designates systemic functions and attributes the resulting obligations to the operators that perform them.

It looks like a small difference. In reality it determines the entire legal architecture of the system.

CHAPTER 6

The risk that remains

Even after a systemic function has been identified, one step is still missing.

Systemic importance is not the same thing as the contribution.

A function may present high dependencies and, at the same time, be accompanied by effective measures on portability, interoperability, continuity and redundancy. An operator may recognise the risk and reduce it before it produces social costs.

If the DSS ignored that capacity, it would turn designation into a permanent sentence.

The model must instead measure what remains after prevention.

That is residual risk.

6.1 Gross risk and residual risk

We call gross risk the exposure measured before effective mitigation is taken into account.

Residual risk is the part that persists after preventive measures have been applied and verified.

The distinction is fundamental because it makes classification dynamic.

An operator may enter the process in a high class and leave it in a lower one. It may improve over time. It may deteriorate if it removes safeguards or if the function becomes harder to replace.

The class is not a label. It is a snapshot, updated periodically.

6.2 Spending is not enough

One of the most important rules concerns how mitigation is measured.

Proving that a great deal has been invested is not sufficient.

A billion spent on a programme that does not reduce switching time cannot automatically produce a discount on the contribution.

The DSS has to measure outcomes.

If a portability procedure cuts a typical migration from a hundred and eighty days to thirty, there is an observable result.

If an open standard increases the number of genuinely usable alternatives, there is an observable result.

If a new continuity architecture reduces the number of parties potentially caught up in a disruption, there is an observable result.

Mitigation is effective when it changes the risk, not when it changes the public relations balance sheet.

6.3 Causally linked mitigation

Every measure must act on the relevant dimension of risk.

Training may reduce part of the labour transition risk, but it does not automatically eliminate technological lock-in.

Portability may reduce non-substitutability, but not necessarily the propagation of a disruption.

Redundancy may reduce propagation, but not the concentration of a critical input.

This prevents arbitrary trade-offs.

An operator should not be able to neutralise a strong infrastructural dependency by funding social initiatives unrelated to that dependency.

Mitigation must be causally linked to the risk it claims to reduce.

6.4 Independent verification

Because reducing risk yields an economic benefit for the operator, the results of mitigation must be verifiable.

The system can use independent audits, migration tests, incident simulations, standardised metrics, customer samples and certified data.

Verification must be proportionate. Not every data point requires the same intensity of scrutiny.

But the variables with the greatest bearing on the class must be capable of being reproduced or challenged.

A model that granted reductions on the basis of self-declaration would not be robust enough.

6.5 The contribution as an incentive

At this point the contribution changes in nature.

It is not merely a source of funding.

It becomes an economic signal.

If reducing the risk costs less than continuing to contribute in respect of that risk, the operator has an incentive to adopt the mitigation.

If mitigation is technically impossible or disproportionate, the contribution mutualises what remains.

The system therefore does not demand that every risk be eliminated. It demands that the reasonably preventable part be eliminated.

The rest is shared.

6.6 Can it reach zero?

Yes, in principle.

If a function passes the screening on scale and significance but, after mitigation, does not present residual risk sufficient to justify a contributory share, the contribution can be nil.

That does not make designation pointless.

The operator may remain subject to obligations on disclosure, continuity, monitoring or periodic review.

But the Fund must not collect merely because a company has been classified as systemic.

This is one of the most important tests for distinguishing the DSS from a tax.

6.7 The paradox of good regulation

Good DSS regulation ought to produce a paradox.

The more effective it is, the fewer resources it should need to collect, for a given set of benefits and a given context.

If portability improves, if dependencies fall, if systems become more resilient and if labour transition is anticipated early, part of the future need can shrink.

The Fund must not become dependent on the very risk it was created to reduce.

That is a decisive institutional safeguard.

6.8 Towards measurement

We have set out the logic: the function comes before the entity, prevention comes before mutualisation, and the contribution concerns residual risk alone. The technical question remains: how do we turn this architecture into verifiable numbers? Part III introduces the six dimensions of systemic importance, the data needed to measure them, the R0–R4 classes, and the procedure by which a decision can be made contestable and reproducible.

PART III

MEASURING POWER

CHAPTER 7

Six dimensions of systemic importance

So far we have built an argument. From this chapter on we have to build a method.

Saying that a function is “very important”, “dominant” or “hard to replace” is not enough to impose legal and economic obligations. An authority must be able to show which data it used, how it turned that data into an assessment, and why two different functions receive different classes.

The methodological proposal of this book uses six dimensions, denoted by the letters S-I-N-D-P-L:

S = Scale.

I = Interconnection.

N = Non-substitutability.

D = Control of data and infrastructure.

P = Propagation.

L = Socio-occupational transformation.

These six dimensions do not claim to be settled scientific truth. They are a regulatory grid proposed by the author: it must undergo empirical validation, public consultation, sensitivity analysis and periodic review before it could become a legislative standard.

Their purpose is more modest and, at the same time, more important: to stop the word “systemic” remaining a formula with no verifiable criteria behind it.

This architecture does not arise in a regulatory vacuum. The indicator-bucket method used for G-SIBs, the designation of critical ICT providers under DORA, and the contributory logic of the Single Resolution Fund are distinct but converging precedents. Chapter 21 sets out this prudential genealogy explicitly and, above all, the limits of transferring it to the DSS.

7.1 S — Scale

Scale measures the reach of the function.

It can be observed through end users, business users, revenue attributable to the function, geographic spread, transaction volume or another unit consistent with the sector.

Scale matters because a function used by a great many parties can produce wider effects. But it must not be sufficient on its own.

An enormous but easily substitutable function does not present the same risk as an equally large function that is almost impossible to leave.

This is why S carries weight in the model but cannot determine designation by itself.

7.2 I — Interconnection

Interconnection measures how far the function is embedded in the activities of others.

What counts is not merely the number of direct customers. It is what those customers do.

A service used by a hundred large intermediaries may have greater propagation capacity than a service used directly by a million consumers, if those hundred intermediaries are themselves essential nodes in economic networks.

I must therefore look at downstream dependencies, presence in critical sectors, integration with other essential services and concentration of relationships.

Interconnection turns the direct user into a potential multiplier of risk.

7.3 N — Non-substitutability

Non-substitutability is probably the most intuitive dimension of dependency.

It measures how difficult it is to replace the function with an equivalent alternative within reasonable time and cost.

N is not the same as the number of competitors. It must include at least switching time, migration cost, effective portability and the number of genuinely equivalent alternatives.

The presence of three competitors does not guarantee substitutability if migration takes years or if the customer loses essential components of its own process.

In the DSS, freedom to exit has to become a measurable quantity.

7.4 D — Control of data and infrastructure

D measures the operator's ability to control inputs, infrastructure, data, standards, compute capacity or other elements that are hard to replicate.

This dimension serves above all to capture systemic importance upstream.

An operator may have few end users and, at the same time, control a component without which numerous services could not be produced at all.

The measure must be adapted to the sector: compute capacity, share of a critical input, access to non-replicable data, control of standards or proprietary technologies, and the time and cost required to replicate the infrastructure.

7.5 P — Propagation

P measures how quickly and how far the effect of a disruption or malfunction can spread.

It is the dimension that distinguishes a commercial difficulty from a systemic risk.

What has to be observed is the number of parties that could be caught up directly and indirectly, the propagation time, the presence of isolation mechanisms, redundancy, and the documented effects of previous incidents.

A function may be very important and yet propagate slowly and containably. Another may generate cascading effects within minutes.

That difference has to show up in the score.

7.6 L — Socio-occupational transformation

L measures the intensity with which the function can change tasks, skills and the organisation of work.

It is the most delicate dimension, because it must not turn technological exposure into an automatic prediction of job losses.

The indicators must therefore observe diffusion across processes, the share of tasks transformed, speed of adoption, scope for reskilling, and the results of training measures.

L does not measure "how many jobs AI will destroy". It measures the transition pressure a function can help generate.

7.7 The proposed weighting

To make the grid operational I propose, as a basis for discussion, the following gross score:

$$DSS_L = 0.20S + 0.20I + 0.20N + 0.15D + 0.15P + 0.10L.$$

Each dimension is normalised on a scale from 0 to 100.

The weighting gives greater weight to scale, interconnection and non-substitutability; intermediate weight to infrastructural control and propagation; and lower — but not marginal — weight to socio-occupational transformation.

These weights should not be presented as "scientifically true" values. They are initial policy parameters. They will have to be stress-tested against real cases to check whether they produce coherent classifications and whether small variations in the weights shift the result excessively.

7.8 The rule against automatic size

To prevent scale alone from turning a large company into a systemic operator, I propose a joint requirement.

Designation requires a gross score of at least 55 out of 100 and, simultaneously, at least one of I, N, D and P equal to or above 70.

In symbols:

$$DSS_L \geq 55$$

and

$$\max(I, N, D, P) \geq 70.$$

The logic is straightforward: there must be not only size, but at least one strong form of dependency, control or propagation.

This threshold too is a proposal to be validated, not a natural fact. Chapter 25 will compare it with less compensatory aggregation rules and test its stability.

7.9 Not a sovereign algorithm

The score does not decide on its own.

It serves to discipline the decision, not to replace it.

An authority must give reasons for the data, the normalisations, any adjustments, and why it considers the information used reliable. The operator must be able to challenge both the underlying data and its transformation into a score.

THE PRICE OF POWER

The algorithm is therefore a structure of transparency.

If it is used to conceal discretion, it fails. If it forces the authority to make the steps of its decision visible, it is doing its job.

Dimension	Principal object	Proposed weight
S — Scale	Breadth and diffusion of the function	20%
I — Interconnection	Downstream dependencies and critical sectors	20%
N — Non-substitutability	Time/cost of exit and alternatives	20%
D — Data and infrastructure	Control of inputs and capabilities that are hard to replicate	15%
P — Propagation	Breadth and speed of cascading effects	15%
L — Transformation of work	Transition pressure on tasks and skills	10%

CHAPTER 8

The DSS Data Return

A quantitative model is useless if the data it needs does not exist, is not comparable, or is available only to the operator being assessed.

This is why the methodology needs a second pillar: the DSS Data Return.

The Data Return is a standardised set of information that the operator under investigation must supply to the authority for the specific function examined and for the relevant territory.

It is not an indiscriminate collection of corporate data. Every item of information must bear a relationship to one of the S-I-N-D-P-L dimensions.

Standardisation has two aims: to make different operators comparable, and to narrow the space within which the authority could pick, case by case, only the indicators that confirm its own hypothesis.

8.1 Data relating to the function, not to the group

The Data Return must respect the principle already established: the function comes before the entity.

If the investigation concerns a European cloud service, the group's global revenue is not enough. What is needed is data attributable, so far as reasonably possible, to that function and that perimeter.

Where an item of data cannot be separated out without disproportionate cost, the operator must say so and propose a verifiable allocation methodology.

8.2 Common definitions

Comparability requires a regulatory dictionary.

What does "user" mean? When is an alternative "equivalent"? How is switching cost calculated? When is an incident significant? What level of re-engineering makes a workload non-portable?

Without common definitions, two operators can supply figures that are formally correct and substantively incomparable.

The implementing regulation should therefore set out definitions, reference periods, units of measurement and sampling criteria.

8.3 The problem of confidential data

Much of the necessary information will be commercially sensitive.

Transparency does not mean full publication.

No.	Area	Data required	Investigative function
1	S	End users	Average number of end users of the function in the territory
2	S	Business users	Number of firms/organisations using the function
3	S	Revenue of the function	Revenue attributable to the function within the perimeter
4	S	Territorial diffusion	Presence and intensity of use by member state
5	I	Downstream dependencies	Number and type of organisations depending on the function
6	I	Critical sectors	Share of dependencies in health, finance, energy, public administration
7	I	Essential integrations	Third-party services/APIs/processes integrated into the function
8	N	Switching time	Median and percentiles of T_switch by representative segment
9	N	Switching cost	K_switch as a ratio of annual spend on the function
10	N	Effective portability	Q_port: functional share transferable without re-engineering
11	N	Equivalent alternatives	A_eq by customer type and requirement
12	D	Critical inputs	Control of capacity, standards, data or infrastructure
13	D	Replicability	Time and cost of reproducing the infrastructure
14	P	Significant incidents	Frequency, duration and breadth of documented effects
15	P	Isolation and redundancy	Fallback mechanisms and containment of propagation
16	L	Transformation of tasks	Indicators of exposure, adoption and re-allocation
17	—	Mitigation plan	Measures, metrics, deadlines and verified results
18	—	Audit evidence	Sources, methods and reproducibility of the data

The authority must be able to receive confidential data, verify it, and publish reasoning detailed enough to be meaningful without revealing trade secrets or security-sensitive information.

The public decision must show the logic of the score; the confidential file can hold the detail needed for verification and for the adversarial process.

8.4 Public data and operator data

The Data Return must not turn the operator into the sole source of its own assessment.

The authority must cross-check the information against public data, customers, competitors, incident reports, audits, independent studies and other legitimate sources.

Divergences must be explained.

Where data is not reliable enough, the correct response is not to invent a score. It is to declare that the component cannot yet be determined and to require further evidence.

8.5 Data quality

Every relevant variable should come with minimum metadata: reference period, source, calculation method, coverage, any estimates and the margin of uncertainty.

This makes it possible to distinguish a certified figure from a modelled estimate.

We will apply the same discipline in the case studies in this book: observed data, estimated data and missing data will always be kept separate.

8.6 Periodic updating

Systemic importance changes.

A function can grow rapidly, lose users, become more portable, or be displaced by new technologies.

The Data Return must therefore be periodic and proportionate to the risk class. The most critical functions may require more frequent updates; those close to the threshold can be reviewed at longer intervals.

The aim is not to accumulate data, but to keep the picture of risk current.

8.7 Penalties and correction

An information obligation with no consequences for false or seriously incomplete data would be ineffective.

The system must provide for a duty to correct, liability for inaccurate declarations, and powers of verification.

At the same time, the operator must be able to correct good-faith errors and to challenge the authority's interpretations.

The Data Return must produce reliability, not a procedural trap.

8.8 From data to score

Once the data is collected, one technical step remains: turning heterogeneous indicators into values from 0 to 100.

This normalisation cannot be improvised.

THE PRICE OF POWER

Each indicator needs thresholds, curves or benchmarks defined ex ante and put out to consultation. Some indicators may use absolute thresholds; others sector percentiles; others still non-linear functions.

This is precisely where the empirical work that must follow this book will be needed: applying the Data Return to real cases and checking which transformations produce stable results.

CHAPTER 9

R0, R1, R2, R3, R4

A continuous score is useful for measuring differences. But a legal system also needs discrete consequences.

I therefore propose five risk classes.

The classes are not moral judgements and do not represent levels of corporate “dangerousness”. They are regulatory bands that make it possible to link the score to proportionate obligations.

The initial classification is built on the gross score. After mitigation, the same structure is applied to the residual score to determine the actual contributory weight.

Class	Proposed score	Basic regulatory meaning
R0	< 40	Insufficient significance for the purposes of the systemic regime.
R1	40 – 54.99	Area of observation; no designation on the basis of the score alone.
R2	55 – 69.99	Moderate systemic importance, where the joint requirement is also met.
R3	70 – 84.99	High systemic importance.
R4	≥ 85	Very high systemic importance.

9.1 The designation threshold

The proposal sets 55 as the minimum quantitative threshold for designation, but couples it with the requirement already seen:

$$\max(I, N, D, P) \geq 70.$$

Accordingly, an operator with an overall score of 60 but no strong structural dimension is not automatically designated.

The threshold exists to prevent a spread of middling values from artificially producing a finding of systemic importance.

9.2 What the classes entail

The classes should bear above all on the intensity of obligations.

R0 may mean closing the investigation or ordinary monitoring.

R1 may justify observation and a simplified Data Return.

R2 may entail designation, a mitigation plan and periodic monitoring.

R3 may require more intensive obligations on continuity, portability, testing and reporting.

THE PRICE OF POWER

R4 may require the maximum level of oversight provided for by law, always subject to proportionality and the right to be heard.

The contribution, however, should not be determined by the gross class. It must await mitigation and the residual score.

9.3 From gross to residual

For each mitigable dimension, an effectiveness coefficient between zero and one is applied.

In simplified form:

$$I_R = I \times (1 - m_I)$$

$$N_R = N \times (1 - m_N)$$

$$D_R = D \times (1 - m_D)$$

$$P_R = P \times (1 - m_P)$$

$$L_R = L \times (1 - m_L).$$

Scale S is not, as a rule, reduced: having a great many users does not become less true because the operator adopts good measures.

The residual score is therefore:

$$DSS_R = 0.20S + 0.20I_R + 0.20N_R + 0.15D_R + 0.15P_R + 0.10L_R.$$

The formula is deliberately transparent. Every reduction must be traceable to a verified mitigation.

9.4 Coefficients for the relative component

To apportion the common component of costs already committed, the model can attach a coefficient k to the class or to the residual score.

The choice of coefficients is a policy decision and will have to be calibrated against the financial stress tests and the protocol in Chapter 25.

The principle matters more than the number: k_i determines a relative share of common costs already incurred. It does not, on its own, determine the overall contribution, which in the mature model also includes an absolute component linked to expected social loss.

R0 must be capable of corresponding to a k_i of zero for this relative component. That does not automatically amount to a total contribution of zero if there is a net expected social loss that is causally attributable.

9.5 Uncertainty and ranges

A score presented to two decimal places can create an illusion of precision.

In reality many indicators will be estimated.

The decision should therefore also state, where materially relevant, the uncertainty around the data, and should carry out a sensitivity analysis near the thresholds.

If a function scores 55.1 but plausible variations in the data bring it to 52, the decision calls for more caution than a case that stays above 65 in every reasonable scenario.

Regulatory precision does not consist in showing many decimal places. It consists in making uncertainty visible.

9.6 The right to be heard on the score

Before the final decision, the operator must know at least: the perimeter of the function, the data used, the normalisation methodology, the score for each dimension, any estimates, the mitigation recognised, and the overall result.

It must be able to challenge each step.

This right is not an obstacle to regulation. It is part of the model's quality.

A methodology that survives challenge is more robust than one that claims to be beyond challenge.

9.7 Review and lapse

Designation must not be permanent.

If a function loses significance, if alternatives emerge, if the market changes, or if mitigation durably removes the dependencies, the class must be able to fall and the designation must be able to lapse.

Equally, a function that is not initially systemic may cross the thresholds rapidly.

The DSS must therefore provide for periodic review, extraordinary review where material changes occur, and clear criteria for leaving the regime.

9.8 A falsifiable method

The value of the methodology will not depend on the elegance of the formula.

It will depend on its capacity to produce coherent results across different cases.

If applying the same rules to cloud, AI models and compute capacity forces us to keep changing thresholds and weights to obtain the desired result, the model will have failed.

If, on the other hand, the methodology can distinguish between different functions without knowing in advance which company ought to come out as the more systemic, we will have built something useful.

That is why the case studies will not be decorative examples. They will be genuine stress tests of the model.

First, though, we have to face another question: what can an operator do in order to pay less?

The answer is the heart of Part IV: prevent more.

PART IV

PAYING LESS BY PREVENTING MORE

CHAPTER 10

The way out

Regulation that confined itself to imposing economic obligations would turn risk into a permanent revenue base. The first question put to the operator is not “how much do you have to pay?” but “what can you do to reduce the dependency your function creates?” The DSS has to offer a way out: reducing residual risk through verifiable technical, organisational and social measures. The more a function becomes reversible, interoperable and resilient, the less it should weigh on the contributory mechanism.

10.1 Prevention comes before transfer

The principle that prevention takes priority changes the order of public intervention.

In a purely compensatory model, you wait for the cost to emerge and then collect resources to cover it.

In the DSS, by contrast, part of the cost has to be avoided beforehand.

If a dependency can be reduced by making data portable, the system must incentivise portability. If a disruption can be contained through redundancy and disaster recovery, the first answer must be resilience. If an occupational transformation can be anticipated through effective training, training comes before income support.

The Fund acts on the risk that remains, not on the risk that was reasonably avoidable.

10.2 Portability

Portability is the ability to transfer data, configurations, workloads, models, identity or other functional elements to an alternative.

But portability that is merely formal is not enough.

Downloading your data in an unusable format is not the same as being able to change service. Exporting information without being able to reconstruct relationships, permissions and configurations can leave lock-in entirely intact.

The metric must therefore be functional: what share of the function can be transferred without substantial re-engineering, within what timeframe, and with what loss of capability?

A good measure of portability has to bring the right to exit closer to the practical possibility of exercising it.

10.3 Interoperability

Interoperability reduces dependency by allowing different systems to communicate and to substitute for one another more easily.

Open standards, documented APIs, common formats and shared protocols can lower the cost of migration.

Not every standard has to be imposed by the authority. But when a function becomes systemic, the absence of interoperability can take on prudential significance.

The question is not whether the proprietary product is legitimate. The question is how far its architecture increases the risk of collective dependency and whether proportionate measures exist to reduce it.

10.4 Multi-provider and redundancy

In some sectors the best mitigation is not to replace a supplier but to stop any single supplier being a single point of failure.

Multi-provider architectures, independent backups, data replication, failover procedures and the ability to operate in degraded mode can all reduce propagation.

These measures carry costs and are not always efficient. The DSS should not impose them automatically.

But it must recognise their value where they demonstrably reduce the number of parties affected by a disruption, or the duration of its effects.

10.5 Operational continuity

Operational continuity is a particularly important form of prevention for infrastructural functions.

Service levels, recovery time, recovery point, incident response, resilience testing and disaster recovery plans can all be turned into verifiable indicators.

The aim is not to demand the complete absence of incidents. No complex infrastructure can guarantee that.

The aim is to reduce the probability, duration and breadth of the consequences.

10.6 Training and transition

Mitigation is not only technological.

Where a function accelerates the transformation of occupational tasks, the operator can help reduce risk through training, certification, access to skills and support for reskilling.

Here too what counts is the result.

The number of courses offered is not enough. What has to be observed is participation, completion, skills acquired, redeployment and other relevant outcomes.

Social spending becomes mitigation only when it reduces a measurable component of transition risk.

10.7 The mitigation plan

Every designated operator should be able to submit a DSS Mitigation Plan.

The plan identifies: the risk dimension concerned, the measure proposed, the expected result, the metric, the implementation deadline, the verification method and the duration of the effect.

The authority can approve, amend or reject the measure, giving reasons.

The plan transforms the relationship between regulator and company: not merely ex post control, but a verifiable path of risk reduction.

10.8 The right to pay less

If prevention really does take priority, the operator must have a genuine regulatory right to have the reduction it achieves recognised.

Not a discretionary concession.

When a verified mitigation lowers N, P, I, D or L under the methodology, the residual score must change.

This principle is decisive for the economic credibility of the system.

The company has to know that investing in reducing dependency may cost less than continuing to fund the risk.

10.9 The sixth principle

We can therefore state the sixth principle:

The DSS must make it economically worthwhile to reduce risk before mutualising it.

The way out is not an exception to the system. It is its main mechanism.

CHAPTER 11

The price of non-substitutability

Of all the model's dimensions, non-substitutability deserves particular attention.

A function can be very widespread without being especially risky, if those who use it can move quickly to an alternative.

Conversely, a function may have fewer users and produce a very strong dependency if leaving takes a long time, costs a great deal, or means losing essential capabilities.

The price of non-substitutability is therefore the economic, temporal and operational cost of the freedom to exit.

11.1 Switching time

The first metric is time.

Let T_{switch} denote the time needed to transfer a representative configuration to an equivalent alternative and restore an acceptable level of service.

The median is often more useful than the mean, because it reduces the influence of extreme cases. But for critical functions the distribution matters too: the fact that half the customers could migrate quickly does not eliminate the risk if a significant share would take vastly longer.

The Data Return should therefore include the median, percentiles and segmentation by customer type.

11.2 Switching cost

The second metric is economic.

We can define a ratio K_{switch} :

$K_{\text{switch}} = \text{total documented cost of migration} / \text{annual spend on the function}.$

The numerator covers more than exit fees. It must include, where relevant, consultancy, re-engineering, data transfer, temporary parallel running, training, testing, certification and the cost of disruption.

The ratio makes organisations of different sizes comparable.

A migration costing a million euro may be marginal for one customer and prohibitive for another.

11.3 Effective portability

The third element concerns what can actually be transferred.

Let Q_{port} denote the functional share transferable without substantial re-engineering.

The measure must not stop at bytes.

In a complex service what counts is data, configurations, identity, permissions, work-flows, models, application logic, dependencies and automations.

A hundred per cent portability of data can coexist with very low portability of the function.

11.4 Equivalent alternatives

The fourth element is the number of alternatives.

Let A_{eq} denote the number of operators able to replace the function within comparable requirements of capacity, security, compliance, territory, performance and timing.

Not every commercial competitor is an equivalent alternative.

For a small business there may be ten suppliers. For a critical infrastructure with specific requirements there might be two.

A_{eq} must therefore be calculated for representative segments, not as an abstract count of companies present in the market.

11.5 An index of non-substitutability

T_{switch} , K_{switch} , Q_{port} and A_{eq} can be normalised and combined into an index N from 0 to 100.

The book does not yet fix a definitive normalisation curve. It would be premature to do so without a comparable dataset of real cases.

The proposal is instead methodological: define sector benchmarks *ex ante*, publish the transformations used, and subject them to sensitivity analysis.

Non-substitutability has to be measured with several indicators, because none of them on its own describes the cost of leaving.

11.6 The migration test

For some high-risk functions, self-declarations can be supplemented by migration tests.

A representative sample of configurations can be transferred, in a controlled environment, to alternatives.

The test measures times, errors, non-portable components and costs.

It must not become a disproportionate obligation for every operator. But in R3 or R4 cases it can provide far more solid evidence of real substitutability.

11.7 Lock-in that originates with the customer

Not all switching cost is attributable to the supplier.

A customer may voluntarily choose heavily customised architectures, neglect backups, make poor use of available standards, or defer investment in portability.

The methodology must distinguish, so far as possible, the structural dependency produced by the function from the customer's own specific choices.

This matters both for fairness and for effectiveness: the DSS must not incentivise irresponsible behaviour by users.

11.8 Non-substitutability as a signal of power

Traditional economic power is often measured through prices and market shares.

In the economy of digital infrastructure there is another question: how far can a party change its relationships without compromising its own activity?

The closer the answer is to "hardly at all", the more the function acquires a form of systemic power.

Not because the operator has necessarily abused its position, but because exit has lost part of its disciplining force.

11.9 The seventh principle

We can state the seventh principle:

Effective competition is measured not only by the number of doors in, but by the cost of the doors out.

In the DSS, non-substitutability is the point where competition policy, resilience and systemic responsibility meet.

Indicator	Regulatory question	Expected effect on N
T_switch	How long does migration take?	More time = greater non-substitutability
K_switch	What is the economic weight of migration?	Higher cost = greater non-substitutability
Q_port	How much of the function is really portable?	More portability = lower non-substitutability
A_eq	How many alternatives are genuinely equivalent?	More alternatives = lower non-substitutability

CHAPTER 12

The contribution

Only now can we talk about money.

First we identified the function. Then we measured systemic importance. We allowed the operator to reduce the risk. We calculated what remains.

The contribution comes afterwards.

That sequence is what stops the DSS turning into a levy engineered backwards to reach a politically predetermined figure.

12.1 Need first

The Fund should not start from a revenue target.

It must start from the benefits.

For each period, estimates are made of the number of beneficiaries, the duration of the measures, unit amounts, administrative costs and prudential margins.

The result is the total funding need.

That need is then attributed to systemic sectors according to documented criteria.

Let F_c denote the need attributed to sector c .

It is this quantity — not the sum of companies' revenues — that determines how much the sector must fund overall.

12.2 A first apportionment formula

To understand the problem we can start from a purely relative formula, useful as a benchmark but not sufficient as a final solution.

In this first representation, each operator i is assigned a coefficient k_i derived from its class or residual score.

The relative reference share can be expressed as:

$$C_i^{\text{rel}} = F_c \times k_i / \sum(k_j).$$

The formula correctly separates amount from distribution, but it is entirely relative: it determines who bears which share of a total that has already been fixed.

For that reason it is not adopted as the DSS's final formula. In Chapters 22 and 23 it will be replaced by a structure that distinguishes costs already committed, expected social loss and the absolute component of risk.

12.3 Why the relative share is not enough

The relative benchmark brings out the problem the definitive model has to solve.

If there are only two operators and one halves its coefficient, part of the burden automatically shifts to the other.

That is coherent only if the sector's funding need is genuinely independent of mitigation.

But if mitigation also reduces expected social costs, the need must be able to fall.

The solution adopted later therefore distinguishes two effects: the relative share of costs already committed, and the absolute level of expected social loss feeding the prospective need.

The step will be formalised in Chapters 22 and 23 through expected social loss, causal attribution of F_c and a hybrid contribution.

12.4 The cap

After the theoretical share, a sustainability limit is applied.

$$C_i^{\text{bench}} = \min(C_i^{\text{rel}}, \text{Cap}_i).$$

In the benchmark, Cap_i can be built from relevant revenues, margins or a combination of economic indicators. In the definitive model the same cap applies after the hybrid contribution of Chapter 23 has been calculated.

The cap does not prove responsibility. It prevents the contribution becoming disproportionate to the operator's economic capacity.

If applying the caps makes it impossible to fund the need in full, the difference must not be concealed. It must surface as a funding gap and require an explicit decision: reduce benefits, draw on reserves, use other sources provided for by law, or recalibrate the system.

12.5 No guaranteed revenue

The DSS cannot promise a definite figure ex ante without knowing the operators, the residual classes, the caps and the need.

To say the system "will collect twenty-five billion" would be methodologically improper if that figure does not derive from real data.

We can build scenarios of ten, twenty-five or forty billion and check whether they are financially sustainable.

But a scenario is not a forecast.

That distinction will be essential in Part V.

12.6 When k_i can be zero

R0 must be able to carry a k_i of zero in the relative component of common costs.

A function under monitoring may therefore not participate in that component at all. The total contribution can nonetheless be positive if the absolute component of net expected social loss is positive.

For R1, a k_i of zero or a minimum coefficient can be provided for, but only where specific statutory conditions are met.

Higher classes can receive rising k_i coefficients for the relative share.

The progression must be calibrated through simulation, avoiding unreasonable jumps near the thresholds.

12.7 A purely illustrative benchmark

Imagine, for teaching purposes only, a sector with a need F_c of one billion euro and three operators with residual coefficients of 1, 2 and 3.

The coefficients sum to 6.

The theoretical shares would be roughly 166.7 million, 333.3 million and 500 million.

This example does not propose real coefficients or rates. It serves only to show the logic: the Fund does not take a percentage of revenue; it apportions a funding need according to risk weights.

12.8 From relativity to expected loss

Now suppose the third operator significantly reduces its risk and moves from a coefficient of 3 to a coefficient of 1.

If the need stays the same, the relative apportionment changes.

But if that mitigation also reduces the expected cost of transition or of incidents, the sector's need should be updated.

The system must therefore be designed to avoid the paradox of penalising the other operators every time one of them improves.

The solution, developed in Chapters 22 and 23, uses an absolute component tied to expected social loss, a relative component limited to costs already committed, multi-year reserves and revision of the need.

12.9 The separate Fund

The resources collected must be ring-fenced.

The contribution is easier to justify if the party paying and the citizen receiving can both verify that the money is not being absorbed indistinctly into general taxation.

The DSS Fund must therefore have separate accounts, statutory purposes, public reporting and independent audit.

Segregation of the fund is not an accounting detail. It is part of the instrument's legal nature.

12.10 The price of power is not a fixed price

We are now in a position to be precise about the book's title.

The price of power is not a rate.

It is not two, three or five per cent of revenue.

It is a relationship between function, dependency, prevention, residual risk, need and economic capacity.

In condensed form:

Systemic function → gross risk → mitigation → residual risk → expected social loss → causal attribution of the need → hybrid contribution → cap.

This chain makes the contribution explainable and contestable.

And it is only after building this chain that it makes sense to ask whether the Fund should be worth ten, twenty-five or forty billion euro.

Step	Question	Variable/result
1. Benefits and baseline	What is the eligible social cost, and what is already covered?	F_{social} ; Existing
2. Incremental cost	What cost remains relative to the counterfactual?	ΔK_h
3. Causal attribution	Which share is attributable to the sector?	$a_{c,h}$; $F_{\text{commit},c}$
4. Prospective risk	What is the net expected social loss?	EL_i^{net} ; $R_{\text{agg},c}$
5. DSS need	How much must prudentially be covered?	$F_c = F_{\text{commit},c} + \lambda_c R_{\text{agg},c}$
6. Hybrid contribution	Which common share and which absolute share?	C_i^*
7. Sustainability	Does the contribution exceed the economic limit?	Cap_i
8. Residual	What cannot be funded?	Gap_c

PART V

**WHAT DOES IT COST TO PROTECT A
SOCIETY?**

CHAPTER 13

Ten, twenty-five or forty billion?

A proposal in economic policy becomes credible when it stops confining itself to principles and agrees to be tested against numbers.

How large should the Systemic Social Dividend Fund be?

Ten billion? Twenty-five? Forty?

The right answer cannot be chosen because a figure “sounds right”. It has to follow from the benefits the system intends to guarantee, the plausible number of beneficiaries, the duration of the measures and the actual contributory capacity of the systemic sectors.

The figures used in this chapter are therefore stress scenarios, not revenue forecasts.

They serve to pose a very concrete question: if society decided to devote resources of this order of magnitude to the right to transition, would the DSS model be able to fund them without producing disproportionate obligations?

Scenario	Hypothetical annual need	Function of the scenario
Pilot	€10 billion	Test of a selective and controllable launch
Central	€25 billion	Intermediate stress test for a multi-sector system
Expansive	€40 billion	Sustainability test under intense transition pressure

13.1 A scenario is not a promise

The three figures do not describe what the DSS would collect today.

Asserting that would require complete data on designable operators, residual scores, caps, sectors, beneficiaries and benefits. That data is not yet available in sufficient form.

The figures are therefore test variables.

The distinction must be maintained throughout the book: a model can be stressed with 25 billion without claiming that 25 billion is currently owed or realistically collectable.

13.2 Building the need from the bottom up

The funding need should be built component by component.

THE PRICE OF POWER

Let B_r denote the cost of transition income, B_f that of training, B_m that of mobility and redeployment, B_t that of territorial measures, and B_a that of administration and evaluation.

In simplified form:

$$F_{\text{social}} = B_r + B_f + B_m + B_t + B_a.$$

F_{social} is the gross cost of the benefits considered in the stress test. It does not automatically coincide with what the DSS can fund: existing coverage has to be subtracted, the additionality rule applied, and only the share causally linked to systemic sectors attributed.

Let F_{DSS} denote the share of the need that passes these filters and can be linked to the mechanism set out in Chapters 22, 23 and 27.

For example, the cost of a training programme is not a percentage of GDP: it is the number of participants multiplied by average cost, adjusted where appropriate for duration, intensity and expected outcomes.

Building it this way makes the need auditable.

13.3 Pilot scenario: 10 billion

Ten billion a year represents a scenario suited to a selective launch.

It could fund a combination of targeted programmes: intensive training, temporary income top-ups, mobility, redeployment and circumscribed territorial measures.

The advantage of the pilot scenario is institutional: it allows data quality, operator behaviour, the effectiveness of mitigation and administrative capacity all to be tested before the system is widened.

The risk runs the other way: a Fund that is too small could screen beneficiaries too heavily, or deliver benefits inadequate to the scale of the transformation.

13.4 Central scenario: 25 billion

Twenty-five billion represents a more ambitious stress test.

It is not chosen as a definitive policy target. It serves to check whether a European system covering several sectors and delivering significant benefits could be funded without turning the contribution into a hidden tax.

In this scenario the presence of the caps becomes essential.

If the theoretical need exceeds what operators can sustainably fund, the model must show the funding gap rather than arbitrarily raising the coefficients.

13.5 Expansive scenario: 40 billion

Forty billion represents a scenario of intense transition pressure.

It is useful precisely because it puts the model under strain.

If reaching that figure required imposing contributions incompatible with the caps, or artificially classifying more operators as systemic, the scenario should be regarded as unfundable in the given configuration.

A serious stress test has to be capable of failing.

The credibility of the DSS also depends on its ability to say that a particular benefit cannot be funded through the Fund alone.

13.6 Aggregate contributory capacity

Once the need has been estimated, aggregate contributory capacity has to be calculated.

Let Cap_i denote the individual operator's cap.

The system's theoretical maximum capacity is:

$$C_{max} = \Sigma(Cap_i)$$

across the operators actually designated and contributing.

But even this value should not automatically be used in full. Caps are limits, not targets.

Sustainability requires looking at the incidence of the contribution on relevant revenues, on margins and on investment, as well as at any pass-through effects on prices and customers.

13.7 The funding gap

If F_DSS exceeds the contribution that is genuinely sustainable, a funding gap emerges.

$$G = F_DSS - \text{available contributions.}$$

The gap must be published.

The law can provide for several responses: reducing or prioritising benefits, drawing on reserves, public co-financing, other compatible sources, or revising the programme.

What the model must not do is conceal the gap by retroactively increasing operators' risk.

13.8 Stress tests, not stage sets

Stress tests must change our beliefs, not confirm them.

For each scenario it will be necessary to simulate at least: different numbers of operators; alternative distributions of residual classes; uncertainty over probability, severity and exposure; different causal shares; higher or lower λ_c ; overlapping scenarios; tighter or looser caps; recession; falling revenues; rising beneficiary numbers; high or low pass-through; strong effectiveness of mitigation.

If the system stays financially coherent only in the most favourable scenario, it is not robust.

If instead it shows clearly when and why it comes under strain, it becomes governable.

13.9 The eighth principle

We can state the eighth principle:

The funding need is a consequence of the benefits; sustainability is a constraint; revenue is not an objective in its own right.

This rule protects the DSS from the temptation to turn a good cause into bad arithmetic.

CHAPTER 14

The cost of the right to transition

Funding the transition requires a more precise definition of what the Fund should pay for.

“Protecting workers from artificial intelligence” is too broad a slogan to become a budget line.

The DSS has to fund benefits that are identifiable, temporary, verifiable and tied to a transition situation.

The right to transition is not the right to stay in the same job for ever. It is the right not to be abandoned in the passage from one economic structure to the one replacing it.

14.1 Transition income

A first benefit may be a temporary income top-up.

The purpose is not to replace wages indefinitely, but to create economic time for training, redeployment or reskilling.

The cost can be expressed as:

$$B_r = N_r \times \text{average amount} \times \text{average duration}.$$

The eligible population must be defined through observable criteria: loss or significant reduction of income, connection with a recognised transformation, participation in a transition pathway, and periodic verification.

The duration must be limited and adjustable.

14.2 Training that produces capability

Training is probably the most intuitive benefit and, at the same time, one of the easiest to waste.

Paying for courses is not the same as creating skills.

The Fund should favour programmes with measurable results: completion, certification, application of the skills, subsequent employment, wage progression or other appropriate indicators.

$$B_f = N_f \times \text{actual average cost per pathway}.$$

Unit cost can vary widely between digital literacy, advanced technical training and full occupational reskilling.

The budget must therefore be built by category.

14.3 Mobility and redeployment

Transition may also require geographic or sectoral mobility.

THE PRICE OF POWER

Relocation expenses, guidance, matching, recognition of skills, placement services and temporary incentives can all shorten spells of unemployment.

These benefits must be voluntary and proportionate.

The DSS must not fund the movement of people as a substitute for regional policy. It should offer one more option to those who choose to use it.

14.4 Exposed territories

Some transformations do not strike isolated individuals alone.

Where an area concentrates highly exposed activity, the simultaneous loss of income and demand can produce multiplier effects.

The Fund can therefore provide for territorial measures: skills labs, transition infrastructure, support for new business formation, partnerships with universities and firms, and guidance services.

But money must not be allocated according to an area's political ability to secure it.

A vulnerability index is needed, which we will develop later through the INDS.

14.5 Administration and evaluation

A serious Fund also has administrative costs.

Information systems, controls, outcome evaluation, audit, fraud prevention, support for beneficiaries and handling of appeals all require resources.

Hiding these costs produces unrealistic budgets.

At the same time, administration must have a ceiling and separate reporting, so that the Fund does not become an apparatus principally engaged in funding itself.

14.6 Avoiding double coverage

Many countries already have unemployment insurance, training funds and active labour market policies.

The DSS should not automatically duplicate them.

It can supplement existing benefits, fill gaps, fund additional components, or co-finance programmes that meet common standards.

Every benefit must state which additional need it covers.

Separate accounting must prevent DSS resources from simply displacing ordinary public spending without increasing actual protection.

14.7 Additionality

The principle of additionality is essential.

If a member state cut its own training spending by a billion because the DSS was funding a billion, the net benefit to society could be zero.

The Fund should therefore fund additional capacity or verifiable improvements.

This does not mean freezing public budgets for ever. It means preventing the DSS from becoming an indirect transfer from systemic companies to general taxation.

14.8 Reversible benefits

Transition is uncertain.

Benefits should therefore be, so far as possible, adjustable and reversible.

A programme that works can be expanded. One that produces no results must be capable of being scaled back or closed.

Evaluation is not a phase that follows the policy. It is part of the policy.

Every major spending line should be created with outcome indicators and a review date.

14.9 A conditional right

The right to transition is not universal in the same sense as a universal basic income.

It is conditional on a defined situation and on a relevant benefit.

This makes it possible to concentrate resources where vulnerability is greatest without turning the Fund into a generalised transfer.

The dignity of the beneficiary, however, requires simple procedures, comprehensible reasoning and the possibility of appeal.

Conditional does not mean arbitrary.

14.10 From cost to value

The cost of the right to transition should not be assessed purely as expenditure.

Faster redeployment can reduce future welfare spending, preserve skills, raise incomes and maintain the tax base.

Effective training can raise productivity.

A successful territorial intervention can prevent cumulative decline.

More mature economic stress tests will therefore have to set expected fiscal and social benefits alongside the gross cost, while staying prudent about the estimates.

The DSS funds a transition, not simply a transfer.

Item	Basic unit of cost	Outcome indicator
Transition income	Beneficiaries $\times$ amount $\times$ duration	Return to work/income after the programme
Training	Participants $\times$ cost per pathway	Certified and applied skills
Mobility/redeployment	Beneficiaries $\times$ cost of service	Time to redeployment
Territorial measures	Programmes/projects funded	Employment and additional capacity created
Administration	Operating costs of the system	Cost per beneficiary and quality of service

CHAPTER 15

The paradox of a Fund that wants to collect less

Most public funds are judged partly on their ability to command the resources they need.

The DSS introduces a further and apparently contradictory criterion: for a given level of social protection, it should seek to reduce over time the quantity of risk it has to fund.

A Fund collecting ever more while dependencies fall would have lost its own rationale.

Its success has to be measured on two sides: protecting better, and making the transfer less necessary.

15.1 The problem of the risk base

An ordinary tax can grow along with its taxable base without that representing any kind of failure.

In the DSS the relevant economic base is different.

If residual risk rises, the need to contribute may rise. But the regulatory objective is to reduce precisely that risk.

The system must therefore not have an institutional incentive to preserve the dependency from which it draws resources.

Governance of the Fund must be separated, so far as possible, from the authority that determines the scores and from decisions on mitigation.

15.2 Separating the regulator from the beneficiary of the revenue

If the same institution that assigns risk classes depended financially on the contributions those classes generate, a conflict of incentives would arise.

One possible architecture separates the assessment function, the management of the fund and the decision on benefits, with cross-checks between them.

The point is not to multiply bureaucracy.

It is to prevent whoever measures the risk from having an interest in always measuring it upwards.

15.3 Mitigation must also reduce the need

In the previous chapter we identified a mathematical problem: if the need stays fixed, reducing one operator's coefficient can shift the burden onto the others.

To avoid this, the model has to link aggregate risk to the prospective need.

The model developed in Chapters 22 and 23 distinguishes:

THE PRICE OF POWER

$F_{\text{commit},c}$ = incremental cost already incurred or committed and causally attributed to the sector;

$F_{\text{risk},c}$ = prudential prospective component;

$R_{\text{agg},c}$ = the sector's aggregate expected social loss after correcting for overlaps.

$F_{\text{risk},c} = \lambda_c \times R_{\text{agg},c}$; $F_c = F_{\text{commit},c} + F_{\text{risk},c}$.

When $R_{\text{agg},c}$ falls durably, $F_{\text{risk},c}$ falls by construction, unless the coverage ratio λ_c is explicitly and reasonably changed.

This does not cancel obligations already incurred, but it stops the Fund staying artificially large after risk has been reduced.

15.4 Multi-year reserves

A prudential fund should not automatically spend everything it collects in the same year.

A share can build reserves earmarked for shocks, transition peaks, or periods in which contributory capacity falls.

The reserves must have a ceiling.

Above that level, the excess can trigger a reduction in future contributions or another use strictly consistent with the law.

Here too the aim is to avoid accumulation for its own sake.

15.5 The regulatory give-back rule

We can imagine a simple rule: if for several cycles the Fund holds reserves above target while aggregate risk falls, the coefficients or the contributory need must automatically be reviewed.

Not necessarily refunded in cash to operators.

But regulation must return the benefit of lower risk through lower future demands.

This is what makes the promise "prevent more to pay less" credible.

15.6 Measuring success

The success of the DSS should not be communicated with the sentence "we collected 30 billion".

Better indicators would be: reduction in average switching time, increased portability, shorter incident duration, number of workers redeployed, growth in skills, shorter periods of income loss, and a shrinking of vulnerable areas.

Revenue is a means.

Resilience and transition capacity are the result.

15.7 The risk of moral hazard

Beneficiaries respond to incentives too.

If the Fund automatically covered every cost of transformation, user firms, workers or administrations might have less incentive to prepare.

Benefits must therefore be designed so as to share responsibility reasonably.

The DSS protects people through transition, but it does not remove the need to adapt.

Equally, systemic operators must not be able to shift onto the Fund costs they could prevent through economically reasonable measures.

15.8 When the Fund should shrink

In an ideal scenario, some components of the Fund can shrink over time.

If technology becomes more interoperable, if alternatives multiply, if skills spread and if the transition slows, the need for extraordinary protection falls.

The DSS must not become an institution permanently searching for a reason to exist at the same size.

It can remain as a prudential infrastructure with a smaller funding need.

15.9 Non-revenue as an outcome

We can now state one of the book's central claims:

The success of the DSS is measured not by how much it collects, but by how much risk it manages to prevent before it has to collect.

This does not mean the Fund should be small as a matter of principle.

If the transition generates enormous needs, it must be able to command adequate resources.

It means that, for a given social need, a euro of risk avoided is preferable to a euro of contribution collected.

15.10 The ninth principle

The ninth principle is therefore:

The Fund must be financially sufficient but institutionally interested in its own reduction when the risk justifying the contribution declines.

With this principle we can move from theory to the hardest stress test.

In Part VI we will apply the methodology to real and mutually different functions. Not to pass judgement on companies, but to check whether the model holds up when it meets the world.

PART VI

THE MODEL MEETS THE WORLD

CHAPTER 16

Microsoft Azure: the cloud test

The first stress test concerns a function that represents the problem of digital dependency almost perfectly: the cloud.

Microsoft Azure is not chosen because this book intends to pass judgement on Microsoft as a company. It is chosen because the cloud allows scale, interconnection, switching, operational continuity and integration with artificial intelligence all to be observed together.

In June 2026 the European Commission set out its preliminary view that Azure and AWS should be designated as gatekeepers for their respective cloud services under the Digital Markets Act. For Azure, the Commission expressly cites market position, an entrenched user base, lock-in, high switching costs, ecosystem, and the growing significance of AI tools in cloud choice.

This does not amount to a DSS designation. But it is a particularly useful piece of public evidence for checking whether the dimensions our model proposes observe phenomena that also exist in real regulatory analysis.

16.1 Observable scale

Microsoft does not publish a complete standalone income statement for Azure. That informational limitation is itself a methodological finding: a DSS Data Return would be needed to isolate the function's revenues, users and dependencies precisely.

Robust public data does nonetheless exist. In its fiscal year 2026 results published on 29 July 2026, Microsoft reported that Azure had passed 100 billion dollars in annual revenue for the first time, growing 41 per cent over the year. In the fourth fiscal quarter, Azure and other cloud services grew 43 per cent year on year, and the Intelligent Cloud segment generated 39.3 billion dollars.

These figures document scale, but they are not enough to assign a definitive S score.

16.2 Interconnection

The cloud hosts applications, databases, identity, networks, analytics tools and AI workloads. Its interconnection therefore depends not only on the number of customers but on the role those customers play.

A genuine I score would require disaggregated data on downstream dependencies, on presence in critical sectors, and on the number of essential services using Azure as infrastructure.

Public information supports a hypothesis of high interconnection, but does not allow that hypothesis to be turned honestly into a precise number.

16.3 Non-substitutability and the Data Act

The European legislature has already recognised that obstacles to switching exist in data processing services.

The Data Act introduces obligations to make the move between providers faster and smoother, to reduce contractual and technical obstacles, and to promote functional equivalence and interoperability. From 12 January 2027, the switching charges covered by the European regime are to be removed entirely.

For the DSS this is particularly interesting: a rule that lowers the cost of exit should, if effective, reduce N over time.

The model thus becomes complementary to existing regulation: it does not replace switching obligations; it measures whether those obligations have actually reduced dependency.

16.4 AI and cloud converge

By 2026 the separation between cloud and AI is increasingly blurred.

The most recent data reinforces the methodological point: Azure's scale is growing rapidly, which makes it more important still to distinguish the observable economic dimension from the data genuinely needed to measure switching, dependency and propagation.

For the DSS this raises a question of perimeter: how much systemic importance belongs to the cloud function and how much to the AI function?

The answer cannot be given by adding the two phenomena together indiscriminately. Double attribution has to be avoided and the specific dependencies identified.

16.5 The score we cannot yet calculate

The most important result of the Azure case is perhaps what we cannot do.

With public data we can document a very high scale, strong signals of interconnection, and a switching problem acknowledged by European regulation itself.

What we cannot know with sufficient precision is T_{switch} , K_{switch} , Q_{port} , A_{eq} , downstream dependencies and propagation across representative segments.

To assign, say, " $N = 82$ " today would be a simulation, not a measurement.

In the final book, any numerical score for Azure must therefore be labelled as a simulation and accompanied by a list of the missing data.

16.6 What the DSS would ask for

A DSS investigation would ask at least for: European users and business users of the function; relevant revenues; sectoral distribution; real migration times and costs; the share of transferable workloads; equivalent alternatives by customer type; significant incidents; dependencies of critical services; and the results of interoperability and switching measures.

Only then would it be possible to move from the narrative case to the quantitative one.

16.7 First conclusion

Azure shows why the DSS cannot be built from financial statements and market shares alone.

The economic dimension is visible.

Systemic dependency, by contrast, requires functional data.

And it is precisely this gap between what the market publishes and what regulation needs to know that justifies the DSS Data Return.

CHAPTER 17

OpenAI: the artificial intelligence test

The second case is harder.

OpenAI allows us to test a function that is not merely infrastructure. Its models and products can enter directly into organisations' cognitive processes: writing, coding, analysis, research, assistance, customer service, product development and software agents.

The DSS question is not whether artificial intelligence is powerful in the abstract.

It is how far a given AI function is embedded in economic processes, how substitutable it is, and what transition pressure it may produce.

17.1 Rapidly growing scale

In November 2025 OpenAI stated that it served more than a million business customers and had over seven million ChatGPT for Work seats. In February 2026 it then reported more than nine million paying business users. In April 2026 it announced that the enterprise segment accounted for more than 40 per cent of revenue and was on track to match consumer revenue by year end.

These are company figures, not statistics produced by an independent authority. They should therefore be used for what they demonstrate: very broad and growing enterprise adoption.

On their own they do not allow systemic importance to be measured.

17.2 Interconnection within workflows

The most significant characteristic is not the number of accounts, but the depth of the integration.

The enterprise report published by OpenAI describes very strong growth in use within organisational workflows and rising intensity of use.

For the DSS, however, it would be necessary to know how many processes become materially dependent on the function's continuity, which of them are critical, which have fallbacks, and how quickly they could move to an alternative model.

17.3 The substitutability of models

The AI market features numerous competing models, both proprietary and open-weight.

This might suggest high substitutability.

THE PRICE OF POWER

But real substitutability depends on the process. An application built around specific APIs, tools, agents, structured outputs, evaluations and safety procedures may require significant work to migrate.

The DSS should therefore measure switching at workload level, rather than confining itself to counting the models available.

17.4 The L dimension

OpenAI is also the case in which the L dimension becomes central.

Work's exposure to AI cannot be inferred from the number of users. What is needed is data on tasks transformed, intensity of adoption, new skills required, actual replacement and capacity for reskilling.

Public data can document use and growth. It cannot yet demonstrate causally how much job loss or job creation is attributable to the function.

Any L score must therefore be built from plural sources, not from operator data alone.

17.5 Upstream infrastructural dependency

The OpenAI case also reveals an upstream dependency.

Advanced models require enormous computational resources and cloud infrastructure. Capacity agreements and commitments can therefore link the model's systemic importance to that of the infrastructure running it.

The DSS must avoid attributing the same risk twice: once to the model provider and once to the cloud.

The right question is which part of the dependency arises from the model function and which from the infrastructural input.

17.6 What is missing

A complete score would require at least: active users and business users by territory; dependencies of critical processes; distribution by sector; times and costs of migrating to alternative models; portability of agents, configurations and evaluations; incidents and fallbacks; occupational tasks transformed; and the results of mitigation programmes.

Public data does not currently cover all these variables.

Here too, the gap must not be filled with invented numbers.

17.7 Second conclusion

OpenAI shows that the systemic importance of AI cannot be reduced either to a model's popularity or to its intelligence.

The decisive variable is how far the function is embedded in processes.

The more an organisation builds capability around a model, the more the cost of leaving can diverge from the simple cost of buying another service.

CHAPTER 18

NVIDIA: the test of upstream power

The third case shifts attention further upstream still.

NVIDIA is not chosen primarily for a service used daily by consumers. It is chosen because contemporary artificial intelligence depends on compute capacity, accelerators, networking and software.

This is the test of the D dimension: control of data and infrastructure that is hard to replicate.

18.1 Economic scale

Verified source: NVIDIA Form 10-K FY2026 filed with the U.S. Securities and Exchange Commission.

In the Form 10-K for the fiscal year ended 25 January 2026, NVIDIA reported total revenue of 215.938 billion dollars, up 65 per cent. Data Center end-market revenue was 193.737 billion, up 68 per cent.

The company describes itself as a data-centre-scale AI infrastructure business.

These figures make the economic scale of the function plain, but once again they do not automatically determine a DSS score.

18.2 Not only GPUs

The relevant function does not necessarily coincide with the individual chip.

NVIDIA describes a data centre platform comprising GPUs, CPUs, networking, interconnects, systems, libraries, software, models, APIs, SDKs and frameworks.

For the DSS the perimeter might therefore be "AI Compute Infrastructure" rather than a specific product line.

The choice of perimeter should be a matter for investigation and for the operator to contest.

18.3 Control of an input

The D dimension looks for precisely this kind of phenomenon.

An input can be systemic when it is necessary to numerous downstream operators and cannot be replicated quickly at equivalent capacity.

Measuring it takes more than knowing NVIDIA's revenues. It requires installed capacity, the actual availability of alternatives, procurement lead times, software compatibility, conversion costs and the dependency of the principal workloads.

18.4 Indirect interconnection

NVIDIA states that the customers of its Data Center platform include large cloud providers, AI model makers, enterprises, startups and public bodies.

This structure creates indirect interconnection: the upstream product enters the services other operators offer downstream.

In the DSS, I and D must therefore be kept distinct but coordinated.

D measures control of the input; I measures how much of the system depends on that input.

18.5 Software and switching

The substitutability of hardware does not depend solely on the availability of another accelerator.

Libraries, toolchains, compatibility, developer skills, networking and accumulated software all count.

Switching cost can therefore be an ecosystem cost.

A genuine Data Return would have to measure how much work is needed to move representative workloads to alternative architectures, not simply compare chip prices and performance.

18.6 Customer concentration and geography

NVIDIA's 10-K also provides information on revenue concentration and geography, but direct customers do not necessarily coincide with end users.

The company distinguishes between direct customers — including distributors, OEMs, cloud providers and model makers — and indirect customers.

This distinction is exactly the kind of problem the DSS has to confront: systemic importance does not necessarily follow the commercial invoice. It follows the functional chain of dependency.

18.7 What the DSS would ask for

For an AI Compute Infrastructure function, what would be needed at minimum is: capacity actually available by territory; the share of critical workloads; equivalent alternatives; conversion times and costs; dependency on proprietary software; concentration of downstream customers; procurement lead times; incidents or capacity constraints; and mitigation through standards and compatibility.

Without this data, a precise D score remains a simulation.

18.8 Third conclusion

NVIDIA shows why a regime limited to consumer-facing services would be incomplete.

Systemic power can sit upstream, at the point where the capacity that allows other markets to function is produced.

The DSS must be able to follow dependency along the chain.

CHAPTER 19

Three functions, three forms of power

Azure, OpenAI and NVIDIA are not three examples of the same phenomenon. That is precisely why they are useful.

Azure tests above all infrastructure, interconnection and switching.

OpenAI tests integration into processes, the substitutability of models and socio-occupational transformation.

NVIDIA tests control of an upstream infrastructural input.

If the DSS were built merely to hit "large technology companies", these cases would look alike. If it is built to measure functions, they must instead produce different profiles.

Function-case	Strong public evi- dence	DSS data still miss- ing	Dimensions to be tested
Microsoft Azure	Economic growth; position	scale; market	Real switching by seg- ment, downstream de- pendencies
OpenAI / AI models and services	Broad adoption and growth	enterprise and strong	Critical dependencies, migration between models
NVIDIA / AI Com- pute Infrastructure	Very high revenue	Data Center	Equivalent alterna- tives, conversion costs

19.1 No league table without data

The temptation would be to assign scores straight away: Azure 78, OpenAI 74, NVIDIA 82.

We will not do so as though these were real measurements.

Without a Data Return and benchmarks, such numbers would carry only an apparent precision.

What we can do instead is build explicitly declared simulated scenarios and see how the formula responds. That is legitimate because it tests the model; it does not claim to describe the companies definitively.

19.2 An illustrative simulation

Imagine three purely hypothetical profiles.

Cloud case: S 90, I 80, N 75, D 60, P 75, L 30.

THE PRICE OF POWER

AI Model case: S 85, I 70, N 60, D 45, P 65, L 80.

Compute case: S 88, I 78, N 82, D 92, P 70, L 40.

Applying the proposed weights gives gross scores of roughly 72.25, 67.50 and 77.90 respectively. The values serve solely to stress-test the formula and are obtained by applying directly the weights stated in Chapter 7.

These values are NOT scores for Microsoft, OpenAI or NVIDIA. They are synthetic profiles constructed to test the formula.

19.3 What the simulation shows

The simulation shows that different profiles can cross the threshold through different combinations.

Cloud stands out for interconnection, switching and propagation.

The AI Model can stand out on the socio-occupational dimension as well.

Compute stands out above all for infrastructural control and non-substitutability.

This is exactly what we wanted from the model: not a single route to defining systemic power.

19.4 The mitigation test

Now suppose the Cloud case reduces N from 75 to 45 through verified portability and switching, and P from 75 to 50 through resilience.

The residual score falls appreciably.

Suppose the Compute case makes workloads more portable to alternative architectures and reduces N and D.

Its score must fall too.

The simulation confirms the property we wanted: the model does not freeze power in place. It recognises reductions in dependency.

19.5 The comparability problem

The three cases also reveal a limitation.

A point of D in compute is not empirically identical to a point of L in an AI model.

Normalisation makes the dimensions combinable, but it does not remove the need for sector-specific benchmarks.

The DSS will therefore need a common methodology plus sectoral technical annexes.

The formula can be general; the raw metrics cannot be entirely so.

19.6 The missing evidence

To turn these chapters into a genuine empirical study, a pilot project would be needed.

An authority or a research group would have to collect the Data Return across a sample of operators, build benchmarks, apply different normalisation curves and observe the stability of the classes.

Only then could we know whether 55 is an adequate threshold, whether the weights are robust, and whether the R0–R4 classes really do separate different profiles.

The book proposes the architecture. Validation requires data.

19.7 From the corporate case to the general rule

Company names are useful for making the discussion concrete, but the DSS must not depend on those names.

Tomorrow Azure could lose position, OpenAI could be displaced by other models, and new architectures could reduce the role of today's dominant accelerators.

A good law has to keep working.

That is why the function comes before the entity, and the data before the designation.

19.8 The tenth principle

We can state the tenth principle:

A model of systemic responsibility is credible only if it can produce different results for different companies, and different results for the same company when the facts change.

The DSS must not know in advance who has to pay.

It must know which questions to ask.

The neutrality test is extended in Chapter 29 to ASML and advanced EUV lithography: a European case, upstream in the chain, useful for checking that the same criteria work even when the bottleneck is neither American nor a digital platform.

Essential public sources for Part VI

Microsoft Investor Relations, FY2026 Q4, 29 July 2026: Azure passed \$100 billion in annual revenue; FY2026 growth +41%; Azure and other cloud services Q4 +43%; Intelligent Cloud Q4 \$39.3 billion.

Microsoft Investor Relations, FY2026 Q4 Earnings Release and Earnings Conference Call, 29 July 2026.

European Commission, 25 June 2026: preliminary view on the possible DMA designation of AWS and Azure; reference to lock-in and switching costs.

European Commission, Data Act: rules on switching between data processing services and removal of switching charges from 12 January 2027.

OpenAI, 5 November 2025: more than 1 million business customers and over 7 million ChatGPT for Work seats; OpenAI, 27 February 2026: more than 9 million paying business users.

OpenAI, 8 April 2026: enterprise above 40% of revenue and on track to match consumer by the end of 2026.

NVIDIA, Form 10-K FY2026: total revenue \$215.9 billion; Data Center \$193.7 billion; description of the data centre platform and customer base.

PART VII

**EUROPEAN LAW, PRUDENTIAL
PRECEDENTS AND VALIDATION**

CHAPTER 20

The DSS and European law

The Systemic Social Dividend would not come into being in an empty space.

The European Union already has an articulated body of rules on digital markets, on data and on artificial intelligence.

The right question is therefore not “which new law could replace the others?” but “which risk remains uncovered, and how do we avoid duplication?”

The DSS should be designed as a complementary prudential and social layer, not as a duplicate of existing regimes.

20.1 Digital Markets Act

The DMA regulates gatekeepers and core platform services with the aim of making digital markets fairer and more contestable.

In 2026 the Commission also placed cloud and AI among the priorities of the DMA review and, on 25 June, expressed its preliminary view that AWS and Azure should be designated as gatekeepers for their respective cloud services, even though they had not crossed the ordinary quantitative thresholds.

This matters for the DSS because it shows that even current law uses a combination of thresholds and qualitative assessment.

But the objectives remain different: the DMA protects fairness and contestability; the DSS as proposed measures residual systemic risk and funds transitions.

20.2 Data Act

The Data Act addresses head-on some of the obstacles to switching in data processing services.

It provides for the removal of obstacles to switching and to multi-provider use, and establishes that from 12 January 2027 providers may not impose switching charges for the changeover process.

For the DSS this regime is a regulatory mitigation that already exists.

If the Data Act genuinely reduces the time and cost of exit, N must fall. The DSS should not charge twice for the same risk: it has to measure the result the rule produces.

20.3 AI Act

The AI Act already contains a category of general-purpose AI models with systemic risk.

For such models it lays down additional obligations, including systemic risk assessment and mitigation, incident reporting and cybersecurity. The GPAI obligations have

applied since 2 August 2025 and, from 2 August 2026, the Commission can fully enforce them, including through penalties.

The DSS uses the word “systemic” with a different perimeter.

It assesses not only model risks in the AI Act sense, but the function’s economic dependency, switching, infrastructure, propagation and socio-occupational transition.

The legislation should make this difference explicit in order to avoid conceptual overlap.

20.4 DORA and critical ICT third-party providers

DORA already regulates, within the financial sector, the risk arising from critical third-party ICT providers, and gives the European Supervisory Authorities an oversight framework with a Lead Overseer for each designated CTPP.

For the DSS this is a direct precedent on the designation of critical technological dependencies, but with a different perimeter: DORA protects the operational resilience of the financial sector; the proposed DSS looks across sectors at economic dependency, residual social risk and the funding of transition.

Data and remediation already produced under the DORA framework must be reused where relevant. Effective DORA mitigation must reduce DSS residual risk rather than generate a second obligation for the same risk.

20.5 The eleventh principle

Principle: where a European obligation already reduces a component of risk, the DSS must recognise it as mitigation; where reliable data is already available, the Data Return must reuse it so far as legally possible.

Where a European authority already holds reliable data, the Data Return should be able to reuse it rather than request it again.

The regulatory once-only principle lowers administrative costs and increases coherence.

20.6 Cooperation between authorities

Cloud, competition, data, AI, cybersecurity and labour fall within different institutional competences.

The DSS therefore needs protocols for cooperation and information exchange.

Cooperation must not erase the different purposes of the various rules. Data collected for one regime may be useful to the DSS, but the DSS decision must remain reasoned according to its own criteria.

20.7 Legal basis and level of intervention

Chapter 28 will identify a possible two-pillar architecture, but not a definitive legal solution.

For the Data Return, designation, mitigation and uniform rules on cross-border functions, Article 114 TFEU is a candidate to be tested; for the specific social and territorial actions, Article 175(3) offers concrete precedents in the EGF and the JTF.

The SRF/Portigon precedent makes it plausible, under certain conditions, for a prudential contribution not to be classified as a fiscal provision for the purposes of Article 114(2) — but the broader link between DSS contributors and social spending makes that test harder to pass. Before any legislative proposal, therefore, an independent legal-basis opinion, State aid screening and verification of international trade obligations are all required.

20.8 Compatibility, not overlap

Principle: the DSS must use existing regulation as mitigation and as a source of information, intervening only on the residual risk and the social need that remain uncovered.

This is the rule that can stop the new regime becoming a further bureaucratic layer with no added value.

CHAPTER 21

Three prudential precedents: DORA, G-SIB and the SRF

A proposal that claims to measure systemic power cannot present itself as though law and prudential regulation had never confronted analogous problems. The DSS arises in a new domain, but some of its most important insights have precise institutional precedents: the supervision of critical ICT providers introduced by DORA, the Basel Committee's indicator-based methodology for global systemically important banks, and the ex ante contribution mechanism of the Single Resolution Fund.

Acknowledging these precedents does not weaken the proposal's originality. It makes it more demanding. It means separating what the DSS can genuinely inherit from what it has to build from scratch, and it prevents us calling innovation something that is already prudential practice in other sectors.

21.1 DORA: technological criticality is already a European legal category

Regulation (EU) 2022/2554 on the digital operational resilience of the financial sector, applicable from 17 January 2025, establishes a European oversight framework for third-party ICT service providers classified as critical. The European Supervisory Authorities — EBA, EIOPA and ESMA — designate critical ICT third-party providers and assign one of the three authorities the role of Lead Overseer. On 18 November 2025 the ESAs published, under Article 31(9) DORA, the first list of 19 designated CTPPs.

The resemblance to the DSS problem is immediate. DORA does not simply ask how large a technology company is. Article 31 and Delegated Regulation (EU) 2024/1502 require assessment of the potential systemic impact of a disruption, the importance of the financial entities that depend on the provider, the support given to critical or important functions, and the degree of substitutability of the provider. In other words, European regulation has already recognised that concentration, dependency and difficulty of substitution can turn an ICT contract into a prudential problem.

The Lead Overseer does not merely observe. The DORA framework allows it to request information, conduct investigations and inspections, issue recommendations and, within the limits set by the regulation, use coercive tools and penalties to make oversight effective. This is an important precedent for the DSS Data Return too: an obligation to supply standardised information to an authority is not a theoretical invention but a technique already embedded in European law for dealing with critical technological dependencies.

21.2 Why DORA does not make the DSS redundant

It is precisely this conceptual proximity that makes the difference worth spelling out. DORA protects the operational resilience of the financial sector. Its subject matter is confined to financial entities and the ICT services they depend on; its purpose is neither to fund social transition nor to measure, across sectors, the economic cost of a systemic function for workers, territories or other industries.

The DSS must therefore treat DORA as a benchmark, not duplicate it. Where a function already falls within DORA's perimeter, information and mitigation validated by that oversight should be reused under a once-only principle. If DORA genuinely reduces concentration, operational risk or non-substitutability, the DSS score must record that reduction. The aim is to measure residual risk after existing regulation, not to charge twice for the same risk.

21.3 G-SIB: the precedent of the indicator-bucket score

The second precedent comes from global banking supervision. The Basel Committee assesses the systemic significance of global systemically important banks through an indicator-based methodology. The SCO40 framework assigns a weight of 20 per cent to each of five categories: size, cross-jurisdictional activity, interconnectedness, substitutability/financial institution infrastructure, and complexity. Banks in the sample report data to national authorities; the indicators are aggregated into a score, and banks above the threshold are allocated to buckets carrying progressively higher requirements for additional loss absorbency. The methodology also provides for a circumscribed and reasoned role for supervisory judgement.

The kinship with S-I-N-D-P-L is evident and should be stated openly. The DSS must not present its scoring as an invention without a genealogy: it transfers to a new domain the idea that systemic importance can be measured across several observable dimensions, that data can be normalised and aggregated, and that discrete classes can trigger escalating obligations.

But the transfer cannot be mechanical. G-SIBs operate within the same sector, with relatively homogeneous quantities and with a specific regulatory consequence: additional capital. The DSS compares heterogeneous technological functions and links the assessment not only to the operator's resilience but also to a possible transition fund. This is why benchmarks must be sector-specific and why the validation of the weights has to be more explicit than is necessary in a system as settled as banking.

21.4 A methodological lesson: the score must not feign neutrality

The G-SIB precedent also reinforces a caution. A composite index does not eliminate discretion: it organises it. Choosing categories, indicators, weights, denominators and thresholds means making policy decisions. The advantage lies not in turning them into mathematics, but in making them public, replicable and open to sensitivity analysis.

The DSS must therefore go one step further: accompanying the basic formula with a calibration protocol, sensitivity analysis and comparison against less compensatory forms of aggregation. The rule $\max(I, N, D, P) \geq 70$ should be regarded as a provisional

solution to be tested against alternatives, not as a natural property of systemic importance.

21.5 The Single Resolution Fund: a fund that is ex ante, ring-fenced and risk-adjusted

The third precedent concerns the most contested step in the DSS: the contribution. Regulation (EU) No 806/2014 establishes the Single Resolution Fund of the Banking Union. Institutions within the perimeter fund it through ex ante contributions; the methodology combines a basic component linked to relevant liabilities with an adjustment reflecting the risk profile. The Fund has a target level and its resources are earmarked for the bank resolution purposes laid down by law.

The parallel with the DSS is strong: a separate fund, a need or target, ex ante contribution, risk as the criterion for modulation, and sectoral mutualisation. And there is a further, particularly significant element: once the SRF's target level is reached, contribution can stop. In February 2026 the Single Resolution Board announced that, for the third consecutive year and barring a change in circumstances, it would not raise new levies, since the Fund exceeded 81 billion euro and the target was still met.

This precedent gives concrete form to one of the DSS's most important claims: a prudential fund must not desire permanent revenue. If the need is covered and the risk is adequately contained, the demand for new resources can fall or cease.

21.6 *Portigon v SRB*: why "it is not a tax" can become a legal argument

The SRF precedent is more relevant still on the question of legal basis. In its judgment in *Portigon v SRB* of 29 May 2024, the General Court of the European Union, Eighth Chamber (Extended Composition), addressed the challenge that the provisions on ex ante contributions to the resolution mechanism were fiscal in nature and therefore outside Article 114(1) TFEU. The Court rejected that approach: the rules imposing the contributions and governing their calculation do not constitute "fiscal provisions" within the meaning of Article 114(2).

The reasoning is particularly instructive. The Court gives weight to the direct earmarking of the levy for the sector's own expenditure, to the necessity of that expenditure for the sector's stabilisation, and to the insurance-like logic of the mechanism. It also stresses that the contributions do not derive from applying a simple rate to a taxable base: a final target is set, then an annual target, and the amount is apportioned among the institutions.

For the DSS this case law is a precedent, not a free pass. The resemblance is strong provided the contribution remains ring-fenced, correlated to a specific risk, directed at the stability and the transition caused by the sector, and does not flow into general taxation. But there is also a difference the book cannot conceal: the SRF funds exclusively the expenditure of the banking sector that contributes, whereas the DSS intends to fund people and territories that may lie outside the sector of the party under the obligation. The legal characterisation of the DSS will therefore depend on its ability to demonstrate a sufficiently close nexus between systemic function, risk, the share of need attributed to the sector, and the benefits funded.

21.7 The critical point the precedents make visible: F_c

This is where the problem the original architecture left too open comes into view. If the overall need is distributed across sectors through F_c, it is not enough to say the decision is “documented”. A causal methodology is required. Otherwise the discretion removed from the score returns at the moment when it is decided how much of the social cost to attribute to cloud, to AI models, to compute capacity or to some other sector.

DORA teaches how to identify functional dependencies; G-SIB teaches how to measure relative systemic importance; the SRF teaches how to tie a financial target to risk-adjusted ex ante contributions. The DSS has to add the step that none of the three precedents is called upon to solve: estimating the share of social loss or transition cost causally attributable to a function. This becomes the necessary bridge between score and euros.

21.8 A new genealogy for the DSS

After this comparison, the DSS can be described more precisely. It is not a digital tax dressed up with a few indicators. Nor is it simply DORA or Basel extended beyond their own sectors.

It is a prudential synthesis in three moves: from DORA it takes the logic of designating and overseeing critical technological dependencies; from the G-SIB framework it takes the discipline of indicators, scores and classes; from the Single Resolution Fund it takes the logic of a ring-fenced fund, of ex ante contributions modulated by risk, and of a target that can stop requiring new resources once it is reached.

The originality — and equally the thing that remains to be proved — lies in the fourth move: turning technological systemic risk into a credible measure of expected social loss and of the share of need causally attributable to the sector. It is on this ground, not on the mere designation of operators, that the DSS has to earn its economic and legal legitimacy.

21.9 Operational consequences for the model

The comparison with the three precedents yields five operational corrections. First: the DSS criteria must be explicitly calibrated and subjected to sensitivity analysis. Second: information already collected under regimes such as DORA must be reused, avoiding duplication. Third: classification into classes must be presented as a prudential technique with precedents, not as a political label. Fourth: the contribution must be linked to an absolute measure of risk and not only to a relative share. Fifth: F_c must derive from a transparent causal attribution methodology.

These steps do not dismantle the model built so far. They force it to become more rigorous. And that is precisely what should happen to a proposal that claims to turn a political principle into a European institution.

Essential sources for Chapter 21

Regulation (EU) 2022/2554 (DORA), in particular Articles 31–35; Delegated Regulation (EU) 2024/1502 on the criteria for designating critical third-party ICT providers; European Supervisory Authorities, first designation of CTPPs, 18 November 2025.

Basel Committee on Banking Supervision, Basel Framework, SCO40 — Global systemically important banks; G-SIB assessment reporting instructions 2026.

Regulation (EU) No 806/2014, Articles 69–70; Implementing Regulation (EU) 2015/81; Single Resolution Board, documentation on ex ante contributions and on the SRF target level.

General Court of the European Union, Eighth Chamber (Extended Composition), 29 May 2024, *Portigon v Single Resolution Board*, T-360/21, on the basis in Article 114 TFEU and on the non-fiscal nature, for the purposes of Article 114(2), of ex ante contributions to the SRF.

CHAPTER 22

From risk to euros: expected social loss

So far the DSS has answered a prudential question: how systemic is a function, and how much risk remains after mitigation? It has not yet answered satisfactorily the decisive economic question: what is that residual risk worth, in euros?

A score from zero to one hundred is useful for classifying, comparing and grading obligations. It is not a currency. Converting it directly into a contributory coefficient would produce false precision and would leave open the objection that the contribution is, in the end, a political convention.

To overcome this limitation I propose separating two layers sharply. The first is the systemic importance layer, built from S-I-N-D-P-L. The second is a layer of expected social loss, expressed in monetary units and built through scenarios. The score identifies where to look and how strictly to oversee a function; expected loss estimates what residual social cost that function may reasonably generate.

22.1 The prudential analogy with expected loss

In bank credit risk the Basel Committee distinguishes probability of default, loss given default and exposure. The product of these components allows expected loss to be estimated. The DSS cannot mechanically import a formula devised for credit, but it can use the same prudential logic: separating the probability of an adverse scenario from the loss conditional on its occurring.

For operator i and scenario h we define p_{ih} as the probability of the scenario, s_{ih} as the monetary social severity per unit exposed, and e_{ih} as the relevant exposure. The aggregate conditional loss of the scenario is $M_{ih} = s_{ih} \times e_{ih}$.

$$EL_i = \sum_h (p_{ih} \times M_{ih}) = \sum_h (p_{ih} \times s_{ih} \times e_{ih})$$

EL_i is the gross expected social loss attributable to the operator's scenarios, before correcting for overlaps and applying the Fund's coverage ratio. The formula is a methodological proposal of the DSS; the analogy with Basel concerns the separation of probability, severity and exposure, not any identity between credit risk and technological systemic risk.

22.2 Severity, exposure and conditional loss

M_{ih} must not be a generic estimate of "harm to society". It has to be built from the benefits and costs the DSS is actually authorised to fund. Severity s_{ih} expresses the cost per unit exposed; e_{ih} expresses how many units are actually exposed in the scenario.

For a cloud disruption scenario, for instance, s_{ih} may include only eligible and documentable unit social costs, while e_{ih} measures the parties or units of activity actually

exposed. For an occupational transformation scenario, e_{ih} may be the number of people or tasks exposed, and s_{ih} may reflect the average duration of transition, eligible income loss and the unit cost of the measures.

Severity is therefore not the same as the company's market capitalisation, nor as total macroeconomic damage. It is the portion of social cost the Fund is legally required to cover.

22.3 The DSS score does not become a rate

S-I-N-D-P-L is not converted through a table of the kind "70 points = 2 per cent of revenue". That would be precisely the step that would make the contribution resemble a parametric tax.

The DSS dimensions serve instead to build or constrain the parameters of the scenarios. Scale and interconnection help define exposure; non-substitutability bears above all on duration and severity; infrastructural control helps identify bottlenecks; propagation influences the number of indirectly affected parties; L guides scenarios of task transformation.

This relationship must be documented scenario by scenario and cannot be automatic.

DSS dimension	Function in the monetary layer	Error to avoid
S — Scale	Helps define the breadth of exposure	Using revenue or users as a measure of loss
I — Interconnection	Identifies dependencies and affected parties	Counting the same category of damage twice
N — Non-substitutability	Influences duration, cost and scope for recovery	Assuming that few competitors automatically means high loss
D — Infrastructure	Identifies critical inputs and bottlenecks	Confusing market share with irreplaceable capacity
P — Propagation	Estimates breadth and speed of indirect effects	Adding direct and indirect effects without correcting for overlaps
L — Work	Builds transition scenarios on tasks	Treating exposure as loss already incurred

22.4 Scenarios, not a single number

A systemic function can produce different risks: temporary unavailability, forced migration, an increase in the cost of exit, infrastructural concentration, rapid change in occupational tasks.

Expected loss must therefore be the sum of distinct scenarios, mutually checked to avoid double counting. Each scenario must state: the event, the time horizon, the exposed population, the probability, the conditional loss, mitigation already in place and the degree of uncertainty.

Where two scenarios share the same damage, the authority must correct for the statistical dependence or merge them into a single common scenario. Naively summing correlated events would inflate the risk.

22.5 Probability with scarce data

For rare events or new technologies, historical frequency may be insufficient. The model must not replace missing data with false precision.

Probabilities can be expressed as ranges or distributions, using observed incidents, stress tests, technical evidence, sector benchmarks and, where unavoidable, structured expert elicitation. Uncertainty has to stay visible.

The decision should therefore publish at least a central estimate and a prudential range, rather than presenting p_{ih} with decimal places that suggest a precision that does not exist.

22.6 From point to range

If p , s and e are uncertain, EL is uncertain too. A robust decision can therefore use a distribution or a range:

$$EL_i \in [EL_i^{\text{low}}, EL_i^{\text{high}}]$$

The ordinary contribution can be based on the central estimate, while a prudential reserve can use an upper percentile or an explicit margin. The choice is a policy one and must be stable over time.

This separates two concepts that are often confused: the expected value funds the predictable average cost; the prudential margin funds the uncertainty around that value.

22.7 Validating the monetary layer

Validation does not concern the composite score alone. It must also test the monetary projections.

After the event, the authority compares estimated probabilities against observed frequencies, projected severity and exposure against actual costs and populations affected, and estimated against real transition times. Systematic errors call for recalibration.

The logic is analogous to that of a prudential model: the model has to be capable of being contradicted by the data and corrected. A mechanism that does not record its own forecasting errors is not genuinely actuarial.

22.8 The twelfth principle

Principle: the DSS score classifies systemic importance; expected social loss monetises cost scenarios. Neither replaces the other.

This separation finally allows the DSS to say something more precise than “whoever has more risk pays more”: it allows an estimate of which risk, through which scenario, produces which expected cost.

CHAPTER 23

Attributing the need: causation, F_c and the hybrid contribution

The most delicate problem in the whole model remains. Even if we can estimate a function's expected loss, how do we decide which share of the social need belongs to cloud, which to artificial intelligence, which to compute capacity — and which, instead, depends on factors that have nothing to do with the DSS?

If this attribution is left to an unstructured political choice, all the discretion removed from designation comes back in through F_c . The sector would become a black box.

To avoid that outcome, the need has to be divided into a component already incurred or committed and a prospective risk component. The two components require different methods.

23.1 The counterfactual before attribution

Attributing a cost means asking what would have happened in the absence of the transformation or shock in question.

Let K_h denote the observed or committed social cost of event h , and K_h^0 the cost that would reasonably have arisen in the counterfactual. The positive difference is the incremental cost to be explained.

$$\Delta K_h = \max(0, K_h - K_h^0)$$

The counterfactual is not directly observable. It has to be estimated by the most credible method the data allows: comparison groups, difference-in-differences, matching, regression discontinuity, synthetic controls, event studies or other appropriate techniques. The European Commission's Joint Research Centre uses counterfactual impact evaluation to identify the causal effects of policies.

23.2 A hierarchy of causal evidence

Not every sector will allow the same level of causal identification. I therefore propose a hierarchy of evidence.

Level A: credible counterfactual evaluation or natural experiment.

Level B: quasi-experimental evidence, panel data or event study with adequate controls.

Level C: structural model or microsimulation calibrated on observed data.

Level D: exposure analysis and triangulation of independent sources.

Level E: structured expert elicitation, admissible only on a transitional basis, with wide ranges and a duty of review.

The weaker the evidence, the more conservative the share attributed to the sector must be.

23.3 The share attributable to the sector

For each event or class of transition h , the authority estimates a share $a_{c,h}$ between zero and one, representing the part of ΔK_h causally attributable to sector c .

Where several sectors contribute to the same cost, the shares must not sum to more than one. The unattributed part stays outside the DSS or is funded from general sources provided for by law.

$$F_{\text{commit},c} = \sum_h (a_{c,h} \times \Delta K_h) = F_{\text{common},c} + \sum_i F_{\text{direct},i}$$

$F_{\text{commit},c}$ therefore represents costs already incurred, recognised or legally committed and causally attributed to the sector. Where the evidence permits robust individual attribution, that share becomes $F_{\text{direct},i}$; the common part that cannot credibly be attributed to an individual operator remains $F_{\text{common},c}$. It is not a percentage picked from the Fund's total.

23.4 Causation with several sectors

Many transformations are joint. An AI model may depend on a cloud provider, which in turn depends on accelerators and networking. Attributing the whole cost to each level would produce triple counting.

The minimum rule is that every euro of incremental cost is attributed once only. In complex cases the authority can use methods of marginal contribution decomposition, including Shapley-type approaches, provided the underlying loss function is public and verifiable.

Where the decomposition is not robust, the answer is not to allocate a hundred per cent to everyone. It is to leave a share unattributed.

The same discipline applies to shared prospective losses. From gross expected loss EL_i , an $EL_i^{\wedge \text{net}}$ is derived after allocating common risks and removing overlaps, so that the same expected damage does not enter the sector more than once.

23.5 The prospective component

Costs already committed look to the past and the present. A prudential fund must also prepare for future risk.

Let $R_{\text{agg},c}$ denote the sector's aggregate expected social loss. The correction for overlaps and common scenarios happens beforehand, through $EL_i^{\wedge \text{net}}$.

$$R_{\text{agg},c} = \sum_i EL_i^{\wedge \text{net}}$$

The law then sets a coverage ratio λ_c between zero and one: not all expected social loss need necessarily be insured by the DSS.

$$F_{\text{risk},c} = \lambda_c \times R_{\text{agg},c}$$

$$F_c = F_{\text{commit},c} + F_{\text{risk},c}$$

23.6 Why this formula closes the relativity problem

In the earlier model, if the need F_c stayed fixed and one operator reduced its coefficient, part of the burden shifted automatically to the others.

The new structure distinguishes costs already committed from prospective risk. Future risk falls when EL_i^{net} falls. For costs already committed, by contrast, a relative apportionment remains necessary because those obligations already exist.

Let w_i denote the operator's relative share, derived from residual risk at the date the contributory cycle is set, to be applied only to $F_{common,c}$.

$$w_i = k_i / \sum_j k_j$$

$$C_i^* = F_{direct,i} + F_{common,c} \times w_i + \lambda_c \times EL_i^{net}$$

The sum of theoretical contributions coincides with F_c if $\sum_i F_{direct,i} + F_{common,c} = F_{commit,c}$ and $\sum_i EL_i^{net} = R_{agg,c}$. Directly attributable costs stay with the individual operator; only the common part is apportioned relatively; the prospective component is absolute and follows net expected loss.

23.7 Freezing the share of costs already committed

To stop a mitigation carried out mid-year retroactively shifting costs already incurred onto the other operators, w_i is frozen at the date the contributory cycle opens for $F_{common,c}$.

Mitigation produces effects on EL_i^{net} — and therefore on the prospective component — immediately or from the next assessment period. In the following cycle it can also change w_i .

The rule creates a clear incentive: reducing risk does not shift future risk onto others; it reduces the sector's prospective need.

23.8 A complete worked example

Suppose a sector has 300 million euro of committed costs, all of them common, so that $F_{direct,i} = 0$ and $F_{common,c} = 300$ million. Three operators have relative coefficients of 1, 2 and 3: their shares of the common costs are therefore 50, 100 and 150 million.

Their respective net expected social losses, assuming for the example that no residual overlaps remain, are 80, 150 and 220 million. If $\lambda = 0.40$, the prospective component is worth 32, 60 and 88 million.

The theoretical contributions therefore become 82, 160 and 238 million, for a total need of 480 million.

Operator	Share of F_{common}	EL_i^{net}	$\lambda \times EL_i^{net}$	C_i^*
A	€50m	€80m	€32m	€82m
B	€100m	€150m	€60m	€160m
C	€150m	€220m	€88m	€238m
Total	€300m	€450m	€180m	€480m

If operator C halves its EL_i^{net} from 220 to 110 million, its prospective component falls from 88 to 44 million. The total need falls from 480 to 436 million and C's contribution from 238 to 194 million. The contributions of A and B do not rise. This is the behaviour the earlier model did not guarantee.

23.9 Caps and funding gap

After the theoretical calculation, the sustainability cap already provided for in the model is applied.

$$C_i = \min(C_i^*, \text{Cap}_i)$$

$$\text{Gap}_c = F_c - \sum_i C_i$$

The gap is not automatically redistributed beyond the caps. It has to be covered through reserves, reduction or prioritisation of benefits, co-financing provided for by law, or other expressly authorised sources.

23.10 Uncertainty does not disappear

EL_i^{net} , $a_{c,h}$, the allocation of common risks and λ_c all contain choices and uncertainty.

The answer is not to hide them. The authority must publish ranges, methods, alternative scenarios and sensitivity analysis. A plausible variation in $a_{c,h}$ or λ_c that produces an enormous change in contributions is a sign of fragility in the model.

At this point the OECD/JRC protocol on composite indicators becomes relevant beyond the score as well: weights, normalisation, aggregation and assumptions must all be subjected to uncertainty and sensitivity analysis, ideally including Monte Carlo.

23.11 The thirteenth principle

Principle: no euro enters F_c without an explicit path linking it to an incremental cost, a counterfactual and an attributable share; no prospective euro is demanded without an expected social loss.

With this structure the DSS does not eliminate discretion. It makes it localised, observable and contestable: in the counterfactual, in the causal share, in the coverage ratio λ and in the expected loss parameters. That is a decisive difference, because it lets an independent reviewer know exactly where to disagree.

CHAPTER 24

Who really pays? Incidence and pass-through of the contribution

A law can determine who is obliged to pay a contribution. It cannot determine, with the same ease, who will bear the cost of it economically.

This distinction between legal incidence and economic incidence is decisive for the DSS. If a systemic operator passes a significant part of the contribution into service prices, the cost can fall on customer firms, public administrations and, indirectly, citizens. In a system created to protect people through transition, ignoring that effect would be a design error.

The question is particularly delicate precisely where the function is hard to replace. But the relationship is not mechanical: economic theory and empirical evidence show that pass-through depends on demand elasticity, competitive structure, contracts, the frequency of price revisions, cost structure, and the firm's ability to absorb the shock in its margins.

24.1 Legal incidence and economic incidence

Legal incidence identifies the party that pays the contribution to the Fund. Economic incidence identifies who suffers the loss of real income after adjustments in price, wages, margins, quantities or investment.

In the DSS the legally obliged party remains the designated operator. But the burden can be distributed among shareholders, workers, suppliers and customers. The aim of incidence analysis is to measure that distribution, not to assume it.

The distinction is essential for the political characterisation of the contribution as well: saying "the operator pays" describes a legal obligation, it does not prove who bears the final cost.

24.2 A prudential cost is not necessarily a marginal cost

The DSS, in the structure proposed after Chapter 23, contains a component relating to costs already committed and a prospective component tied to expected social loss. It is a periodic contribution levied on the operator, not a tax applied to every individual transaction.

For this reason it is not correct to assume automatically that every euro paid becomes a euro of higher marginal cost for the service. A share can be absorbed in margins, a share recovered through contract revisions, a share passed on to specific customer segments.

Incidence must therefore be estimated empirically at the level of the function and the market, not deduced from the contribution formula alone.

24.3 The pass-through coefficient

For stress-testing purposes we can define θ_i as the share of operator i 's contribution that, over the horizon considered, is recovered through higher revenue from customers relative to a counterfactual without the contribution.

In a simplified and partial representation:

$$B_{\text{customers},i} = \theta_i \times C_i$$

$$B_{\text{operator},i} = (1 - \theta_i) \times C_i$$

This decomposition is a stress-testing tool, not a universal economic law. B_{operator} is a residual category and may translate into lower margins, wages, payments to suppliers or investment. Pass-through can be incomplete, differentiated across customers and, in some markets, even higher than the initial cost as a result of markup adjustments. This is why θ has to be estimated and not imposed by definition.

24.4 Non-substitutability and elasticity: a relationship to be measured

A function that is hard to replace may face less elastic demand and therefore have a greater ability to recover costs through prices. This is the strongest objection to the DSS: the most systemic operator might have the most room to pass the contribution on.

There is, however, no monotonic rule valid in every market. Studies of cost shocks show that firms with high markups may also absorb a larger part of the shock in their margins; other contexts show greater ability to pass costs on. Market structure and pricing practices matter as much as concentration alone.

The DSS must therefore treat incidence as an independent empirical variable. N cannot automatically become a pass-through coefficient.

24.5 The signal from digital services taxes

Digital services taxes are not the DSS: they typically fall on gross revenues and have a different legal and economic basis. But they offer a useful warning.

IMF Working Paper 2026/152, *Taxing Cross-Border Services*, estimates a negative and statistically significant effect of the adoption of digital services taxes on imports of digital services relative to non-digital services. The authors note that this evidence does not cleanly identify the incidence of the tax, but that it is consistent with an increase in the domestic cost for users of imported digital services and with a consequent fall in demand. The finding does not show what the incidence of the DSS would be; it does show why the pass-through hypothesis has to be tested empirically.

The lesson is methodological: the promise "the platform pays" has to be put to a market test.

24.6 SMEs and public administrations

Pass-through matters particularly where the customers are SMEs, hospitals, universities or public administrations.

A proportionally small increase in the price of a digital infrastructure may have modest effects for a large company and significant ones for organisations with rigid budgets. Where the function is embedded in multi-year procurement contracts, the effects may emerge above all at renewal.

Incidence analysis must therefore segment at least between large firms, SMEs, the public sector and consumers where present. An average value of θ can conceal a regressive redistribution across customer categories.

24.7 A mandatory incidence test

Before setting λ_c and the sector caps, the authority should publish an Incidence Test.

The test must estimate at least: the share of the contribution potentially transferable; the timing of pass-through; the most exposed customer segments; the expected change in quantities; effects on investment and quality; and the concentration of the burden on SMEs and the public sector.

The estimate can use price differences across markets, comparable regulatory events, contracts, procurement data, renewals, margins and structural demand models. Where data is insufficient, low, central and high scenarios should be published.

Scenario	Assumed θ	Contribution C	Burden on customers	Burden on operator
Low pass-through	20%	€100m	€20m	€80m
Medium pass-through	50%	€100m	€50m	€50m
High pass-through	80%	€100m	€80m	€20m

24.8 How to respond to high pass-through

The system must not reward an operator that raises prices. Reducing the contribution ex post because it has been passed on to customers would create a perverse incentive.

If the Incidence Test signals high and socially regressive pass-through, the response must concern the design of the sector regime: gradual phasing, a more prudent λ_c , tighter caps, multi-year smoothing, greater weight on reserves, or a different mix of DSS funding and other sources.

The decision must seek to preserve the incentive to mitigate without turning the Fund's indirect beneficiaries into its principal funders.

24.9 Portability as incidence policy

Measures that reduce lock-in can also bear on the ability to pass the contribution on.

The easier it is to change supplier, the greater the competitive discipline on price can become. From this point of view, the switching and interoperability obligations of the Data Act are not merely mitigations of N: over time they can help make the market more sensitive to price increases.

But this is an effect to be measured. The DSS must not assume that every formal improvement in portability immediately produces greater demand elasticity.

24.10 Ex post monitoring

The ex ante Incidence Test must be followed by ex post monitoring.

The authority compares prices, fees, contractual terms, margins, quantities and investment before and after the contribution is introduced, controlling for other shocks.

Where it observes pass-through above its estimates, it updates the model for the next cycle. Where pass-through is lower, it avoids retaining excessively conservative assumptions out of caution.

Incidence thus becomes a calibrated parameter rather than a political argument.

24.11 The fourteenth principle

Principle: the operator's legal liability does not prove the economic incidence of the contribution.

A credible DSS has to measure who really pays, and to change the Fund's design when pass-through risks falling significantly on the very people the policy is meant to protect.

CHAPTER 25

Calibrating the DSS: weights, aggregation and robustness

The S-I-N-D-P-L formula has so far been built as an explicit policy proposal. That is preferable to concealing the choices, but it is not enough.

A composite index becomes credible when it demonstrates that its results do not depend arbitrarily on a particular set of weights, on a specific normalisation, or on a single aggregation rule.

The protocol in this chapter therefore turns the initial formula into a hypothesis to be calibrated. The main methodological reference is the *Handbook on Constructing Composite Indicators* developed by the OECD and the Joint Research Centre, which treats theoretical framework, multivariate analysis, normalisation, weighting, aggregation, uncertainty and sensitivity analysis as distinct steps in building an index.

25.1 The construct first, the weights afterwards

Calibration must not begin by asking “how much does S weigh?” It must begin by asking what systemic importance means and which dimensions are conceptually necessary to measure it.

S, I, N, D, P and L must each have operational definitions, observable indicators and a causal or prudential hypothesis explaining why they enter the model.

A statistically elegant weight does not rescue a conceptually wrong dimension. Equally, a strong correlation between two dimensions does not automatically imply that one should be dropped if they represent different mechanisms.

25.2 Multivariate analysis and redundancy

Before aggregation, the pilot dataset must undergo a statistical audit.

Correlation matrices, variance inflation factors, principal component analysis and cluster analysis can reveal redundant variables, indicators measuring the same information, or dimensions that do not behave as expected.

These tools do not by themselves determine the policy structure. They serve to make visible when the model is counting the same phenomenon twice, or when a theoretically important dimension bears no relation to the rest of the data.

25.3 Eliciting weights instead of inventing them

The initial weights of 20-20-20-15-15-10 can be retained as a working prior, but not as a final result.

One initial procedure is budget allocation: a panel of experts is given a budget — say one hundred points — to distribute across the six dimensions after definitions and scenarios have been shared with them.

The procedure must be repeated with different groups — regulators, companies, academics, representatives of users and workers — and must publish the median, the dispersion and the differences between groups.

The final weight must not be a hidden political average. It must be a reasoned choice made in the light of these distributions.

25.4 Cooke's method: where it can help and where it cannot

Cooke's Classical Model is a method of structured expert judgment that assesses experts' performance through calibration and information questions and can assign differentiated weights to their judgements.

In the DSS it may be useful above all where experts have to estimate uncertain quantities — probabilities, technical thresholds, switching ranges, scenario effects — and where seed questions with observable outcomes exist.

It must not be used to convert policy preferences automatically into "scientific truths". The fact that an expert is well calibrated on technical quantities does not confer greater democratic legitimacy on their view of how much weight the transformation of work should carry relative to, say, infrastructural continuity.

The distinction between empirical expertise and normative choice must remain explicit.

25.5 Three aggregation rules to be put in competition

The weighted sum is fully compensatory: a low value in one dimension can be offset by high values in others. The rule $\max(I, N, D, P) \geq 70$ was introduced to stop scale alone producing a designation, but it is a constraint external to the formula.

Calibration must compare at least three specifications.

A. Weighted arithmetic mean with a gate — that is, the current model.

B. Geometric aggregation of resilience, which structurally reduces compensability.

C. A non-compensatory or semi-compensatory specification used as a robustness benchmark, without assuming it should become the final model.

$$A_i = 100 \times \sum_d (w_d \times x_{id})$$

where x_{id} is the normalised risk of function i in dimension d , between zero and one.

$$G_i = 100 \times [1 - \prod_d (\max(1 - x_{id}, \varepsilon))^{(w_d)}]$$

The second formula does not compute the geometric mean of risk, but of resilience, $1 - x$. A single very high vulnerability sharply reduces overall resilience and raises G , limiting the extent to which the other dimensions can compensate. ε is a very small technical value needed to handle the case of zero resilience.

The table is a mathematical simulation, not a validation. It does, however, show a useful property: the geometric mean of resilience penalises imbalance more heavily and makes it harder to offset an extreme vulnerability with low values elsewhere.

Simulated profile	Weighted arithmetic A	Geometric of resilience G
Scale 90; other dimensions 20	34.00	47.22
N = 90; other dimensions 40	50.00	58.07
All dimensions = 65	65.00	65.00
Illustrative Cloud profile, Ch. 19	72.25	76.32

25.6 Why not adopt the geometric mean straight away

The fact that the geometric mean is used in the Human Development Index does not mean it is automatically right for the DSS.

In the HDI the geometric mean is applied to development dimensions in which low values should drag the overall index down. In the DSS we have risk dimensions, so they have to be converted into resilience or a coherent specification defined.

The choice between A and G must depend on empirical validation: the ability to distinguish known cases, the stability of the classes, consistency with independent prudential judgements, and sensitivity to assumptions.

25.7 Monte Carlo: the right question is not "what is the score?" but "how stable is it?"

For each pilot operator the system must generate thousands of plausible specifications.

Weights are drawn from distributions consistent with the elicitation; normalisation thresholds are perturbed within documented ranges; uncertain data is sampled from its respective distributions; aggregation can alternate between admissible specifications.

For each simulation, a score and a class are calculated.

The main result is then no longer merely, say, "DSS = 61". It is also "probability of R2 classification = 72 per cent; R1 = 23 per cent; R3 = 5 per cent".

A fragile class must lead to a more cautious decision than a class that is stable across 99 per cent of the simulations.

25.8 Sobol indices and the importance of assumptions

Global sensitivity analysis can then decompose the variance of the result.

First-order Sobol indices measure how much of the variance in the score or in class probability is attributable to a single uncertain input. Total indices also include interactions with the other inputs.

If the result depends almost entirely on a single threshold, on one particular imputation, or on the weight of N, the model must say so. The Joint Research Centre uses Monte Carlo techniques and variance-based sensitivity analysis precisely as a quality assurance tool for composite indicators.

Sensitivity analysis does not show that the model is "true". It shows which assumptions govern its conclusions.

25.9 Model acceptance criteria

Before the contributory phase I propose at least five tests.

THE PRICE OF POWER

First: class stability. A high proportion of cases must not change class under plausible perturbations.

Second: ranking stability, where a ranking is used.

Third: absence of unjustified dominance. No minor technical parameter should on its own explain most of the variance.

Fourth: discriminant validity. Functionally different profiles must be able to produce different configurations.

Fifth: retrospective validity. Across incidents or dependencies already observed, higher scores must be associated, at least on average, with more severe or harder-to-mitigate systemic effects.

25.10 A seven-stage calibration procedure

DSS calibration can therefore be formalised in seven stages:

- theoretical and operational definition of the six dimensions;
- collection of the pilot Data Return and quality control;
- multivariate and redundancy analysis;
- structured elicitation of weights and parameters;
- comparison between aggregation rules;
- Monte Carlo and global sensitivity analysis, including Sobol indices where appropriate;
- retrospective validation, publication of the calibration report and a reasoned choice of the regulatory specification.

The final methodology must be versioned. Every future change must state which cases change class, and why.

25.11 What remains of the ≥ 70 gate

During the pilot, the rule $\max(I, N, D, P) \geq 70$ must be retained as a benchmark, not necessarily as a definitive rule.

If geometric aggregation of resilience, or another non-compensatory specification, proves able to capture the risk of a single structural vulnerability in a stable way, the gate may become redundant.

If, on the other hand, the gate improves discriminating power without producing arbitrary discontinuities, it can be kept.

Its fate must depend on the results of the protocol, not on the author's preference.

25.12 The fifteenth principle

Principle: no weight, no threshold and no aggregation formula may be shielded from empirical testing.

The DSS becomes falsifiable when it does not merely publish its algorithm, but also publishes how far its conclusions change when the algorithm's assumptions change within plausible limits.

Essential legislative sources for Part VII

Regulation (EU) 2022/1925 — Digital Markets Act.

Regulation (EU) 2023/2854 — Data Act, in particular the rules on switching between data processing services.

Regulation (EU) 2024/1689 — AI Act, rules on general-purpose AI models and on models with systemic risk.

European Commission, 25 June 2026 — preliminary view on the possible DMA designation of AWS and Microsoft Azure.

European Commission — Guidelines on obligations for General-Purpose AI providers and the framework for applying the GPAI obligations.

PART VIII

**WORK, ADDITIONALITY AND
EUROPEAN SOVEREIGNTY**

CHAPTER 26

The L dimension: measuring tasks, not predicting unemployment

In the first version of the DSS, L was the least well grounded dimension. Talking about socio-occupational transformation without a task-based methodology exposes the model to a serious error: confusing what a technology *can* do with what firms will adopt, and confusing what is adopted with what will actually displace labour.

Recent economic literature makes it possible to build on firmer ground. The common thread is not a single prediction about how many jobs will disappear. It is the shift from treating an occupation as an indivisible block to looking at the tasks, capabilities and skills that make up the work.

The new definition of L is therefore narrower: L measures the transition pressure associated with the systemic function on tasks, skills and the scope for reallocation. It does not measure “unemployment caused by AI” and cannot be used as a shortcut for it.

26.1 The task-based framework

The framework developed by Acemoglu and Restrepo distinguishes automation that displaces labour in tasks previously performed by people from the opposite process of reinstatement, in which new technologies and new tasks can create fresh areas of comparative advantage for labour.

This distinction is essential for the DSS. A technology can increase an occupation’s exposure and, at the same time, raise the worker’s productivity without eliminating the role. L must therefore not rise automatically whenever technical exposure rises.

The relevant signal is the difficulty of the transition: which tasks change, how fast, with what balance between substitution and complementarity, and how transferable skills are towards the tasks that remain or emerge.

26.2 Eloundou et al.: exposure does not mean prediction

The study “GPTs are GPTs” assesses the correspondence between the capabilities of large language models and occupational tasks. For the United States labour market, the authors estimate that around 80 per cent of the workforce could have at least 10 per cent of their tasks exposed, and that around 19 per cent could have at least half of their tasks affected.

The most important point for the DSS is not the percentage itself. The authors state explicitly that they are not forecasting the timing of the technology’s development or adoption.

Capability exposure is therefore an indicator of potential. It is not a measure of observed substitution, nor a prediction of redundancies. In the DSS it can feed the first layer of L, but not L as a whole.

26.3 Felten-Raj-Seamans: from capabilities to occupations

Felten, Raj and Seamans construct the AI Occupational Exposure measure by linking advances in artificial intelligence applications to the abilities required by occupations. The method produces a comparable measure of exposure by occupation, sector and geography.

The value of the method lies in separating an empirical question — which abilities correlate most with AI capabilities? — from the economic conclusion about wages and employment.

For the DSS, the AIOE is therefore a useful benchmark for the potential exposure component, especially when mapped onto occupational and territorial classifications. But here too, exposure has to be combined with adoption data and observed outcomes.

26.4 Webb: patents and task content

Webb uses the similarity between descriptions of occupational tasks and patent text relating to technologies to build a different measure of exposure. His historical finding suggests that AI, unlike earlier waves of software and robotics, may also be directed at high-skill tasks.

The method matters because it reminds us that L cannot be approximated by educational attainment. High education does not automatically mean low exposure, and low education does not automatically mean high exposure.

In the DSS, personal characteristics must not become proxies for technological vulnerability. The unit of analysis remains the task and its transformation.

26.5 ESCO, O*NET and a European task infrastructure

For a European application, L must not be allowed to depend on a single United States classification.

ESCO provides a European taxonomy of occupations, skills and competences. *ONET offers a particularly granular description of tasks, abilities and work activities. The OECD has already used crosswalks between ESCO and ONET to analyse how AI exposure relates to skill demand in European vacancies.*

The pragmatic solution is a hybrid infrastructure: ESCO as the European reference for occupations and skills; O*NET and other granular databases as complementary sources on tasks; vacancy data and administrative data to observe actual skill demand.

The crosswalk must be versioned, public and subject to checks for matching errors.

26.6 The four layers of L

I propose breaking L down into four blocks, without yet fixing their weights.

L1 — potential task exposure: the share and intensity of tasks technically affected by the function's capabilities.

L2 — observed adoption: the share of firms, organisational units or workers actually using the function in the exposed tasks.

L3 — substitution/complementarity balance: evidence on the share of tasks eliminated, reduced, augmented or created.

L4 — reallocation distance: the volume of new skills, the training time and the occupational distance needed to reach accessible tasks or occupations.

Chapter 25 determines how these blocks should be normalised, weighted and aggregated. Their conceptual definition comes before their weight.

L block	Question	Possible sources	Error to avoid
L1 — Exposure	Which tasks can the technology technically transform?	ESCO/O*NET; capability studies	Equating exposure with job loss
L2 — Adoption	Is the function actually embedded in the process?	Firm surveys; usage data; vacancies	Using technological availability as proof of adoption
L3 — Substitution/complementarity	What happens to tasks after adoption?	Panels; causal studies; aggregated HR data	Counting every automation as displacement
L4 — Reallocation	How hard is it to move to accessible tasks?	Skill distance; vacancies; training	Using educational attainment or age as a proxy

26.7 A leading indicator and a lagging indicator

L must have both an anticipatory component and a verification component.

L1 and L2 are predominantly leading indicators: they make it possible to identify transition pressure before unemployment shows up.

L3 and L4, by contrast, contain data observable during and after the transformation. They also serve to validate the initial estimates.

The model must avoid circularity: if redeployment time is used to construct L, that same data cannot then be presented as independent evidence that L predicts redeployment time well. Validation requires outcomes not embedded in the index, or out-of-sample data.

26.8 From L to the need, without a causal leap

Even a high L does not automatically produce a single euro of F_c .

L identifies where transition pressure may exist. Chapter 23 still requires a counterfactual and a causal share in order to attribute costs to the sector; Chapter 22 still requires probabilistic scenarios to monetise the risk.

This separation stops an index of occupational exposure becoming, with no intermediate steps, a contributory base.

26.9 The sixteenth principle

Principle: the DSS measures the transformation of work at the level of tasks and skills, distinguishes exposure, adoption and displacement, and does not treat a technology's potential ability to perform a task as proof that the task will be eliminated.

L thus becomes a dimension of transition, not an algorithm for predicting who will lose their job.

CHAPTER 27

Additionality: the EGF, the Just Transition Fund and the DSS

The European Union does not start from zero when it funds transitions. Instruments already exist to support workers hit by restructuring and territories exposed to structural change.

The DSS therefore cannot justify itself by saying vaguely that “a Fund is needed”. It has to demonstrate which need is not already covered, which additional risk it addresses, and why a new funding source improves the capacity to act rather than duplicating existing programmes.

The two most important comparisons are the European Globalisation Adjustment Fund for Displaced Workers and the Just Transition Fund.

27.1 The EGF: from reaction to prevention

Regulation (EU) 2021/691 established the EGF to support workers and self-employed people affected by major restructuring events. In May 2026, Regulation (EU) 2026/1139 extended its scope to workers facing imminent job displacement in restructuring firms.

The consolidated regime includes digitalisation and automation among the relevant transformations. The EGF can therefore act within part of the same social space that concerns the DSS and, following the 2026 revision, can do so even before dismissal.

This makes it inadequate to describe the DSS as simply “a fund to retrain people whose jobs are at risk”. That function already exists, at least in part, in Union law.

27.2 The Just Transition Fund: the territorial dimension already exists

The Just Transition Fund, established by Regulation (EU) 2021/1056, supports the people, economies and environment of territories facing serious socio-economic challenges arising from the transition towards the Union’s climate objectives and climate neutrality.

The JTF funds, among other things, economic diversification, reskilling and upskilling, job search and active inclusion in the territories identified through territorial plans.

Here too the parallel with the DSS is plain: territory, skills and transition are not new categories. What is new about the DSS has to be sought in the trigger and in the prudential mechanism, not in the existence of the benefits.

27.3 Three different triggers

The EGF, the JTF and the DSS have to be distinguished by what activates them.

THE PRICE OF POWER

EGF: a major restructuring event and the conditions set out in the regulation, now including workers at imminent risk.

JTF: territorial vulnerability linked to the climate transition and to territorial just transition plans.

DSS: a function designated as systemic, a measured residual risk, and a causal relationship between that risk and an incremental need.

The same worker could in theory fall within the intersection of all three instruments. The rules must then establish who pays for which benefit, rather than pretending the perimeters do not overlap.

Instrument	Principal trigger	Funding	Distinctive function
EGF	Restructuring / displacement, including imminent	EU budget, mobilised under the applicable rules	Personalised reintegration services
JTF	Territories affected by the climate transition	Cohesion policy/ EU budget	Territorial diversification and reskilling
DSS	Residual risk of a designated systemic function	Dedicated prudential contribution, ring-fenced	Prevention of systemic risk and funding of transition

27.4 The additionality rule

The DSS must adopt a rule stronger than a simple declaration that it “does not duplicate”.

For each benefit *j*, let Cost_{*j*} denote total eligible cost and Existing_{*j*} the coverage already committed by the EGF, the JTF, the ESF+, national instruments or other public sources.

The share fundable by the DSS cannot exceed the residual incremental need.

$$DSS_j \leq \max(0, Cost_j - Existing_j)$$

The formula is an accounting ceiling, not a rule about priority between funds. The legislature has to determine when the DSS tops up, co-finances or acts as a last resort. In every case, the same cost cannot be reimbursed twice.

27.5 Substantive additionality, not merely accounting additionality

Avoiding double funding is not enough.

If a member state cuts ordinary training expenditure because the DSS covers the same amount, net additionality can be zero. Service levels therefore have to be compared against a baseline: number of beneficiaries, intensity, quality, timeliness and results.

The DSS must fund additional capacity or additional risk, not become a substitute source for public obligations that already exist.

27.6 A single door for the beneficiary

From the individual’s point of view, a plurality of funds must not become a plurality of procedures.

The beneficiary should make a single application, or gain access through a single service, while the administration determines behind the scenes which share falls to the EGF, the JTF, the ESF+, national instruments or the DSS.

Keeping the funds separate is an accounting and legal requirement. It must not be offloaded onto the citizen as a bureaucratic cost.

27.7 Shared data and once-only

The same logic applies to proving need.

Where data is already available in a European or national system and can lawfully be reused, the DSS should not request it again. A register of funded benefits must prevent double funding and make outcome evaluation possible.

Administrative interoperability and data protection have to be designed together: sharing what is needed, separating purposes, and limiting access.

27.8 The added value that remains to the DSS

After this comparison with existing instruments, the DSS's specificity narrows but becomes clearer.

First: it links funding to a systemic private function and to a measured residual risk.

Second: it creates an economic incentive to mitigate before contributing.

Third: it can build prudential reserves before the social cost materialises.

Fourth: it links the need to documented causation and not merely to a category of beneficiaries.

If these four characteristics do not produce observable added value relative to existing instruments, there is no good reason to create the DSS.

27.9 The seventeenth principle

Principle: the DSS must not fund what another instrument already funds, and must not displace ordinary public expenditure.

Its social legitimacy grows the more it can demonstrate additionality relative to the EGF, the JTF and the ESF+, and act on the risk that arises specifically from the function's systemic importance.

CHAPTER 28

The European door: legal basis, the SRF and the model's limits

A European proposal does not survive merely by making economic sense. It has to be capable of adoption on the basis of a competence conferred by the Treaties and through the correct legislative procedure.

For the DSS this question is existential. If the contribution were classified as a fiscal measure in the sense relevant to the Treaties, the Article 114 TFEU route would be barred by the fiscal provisions referred to in its paragraph 2, and the problem of unanimity would become central. If, on the other hand, the mechanism formed part of a harmonised internal market regime with prudential logic, the Single Resolution Fund precedent shows that the answer may be different.

But *Portigon* is not a free pass. The DSS funds social and territorial costs far broader than the stabilisation of a single sector. The resemblance to the SRF is therefore a starting point, not the conclusion.

28.1 The legal basis depends on aim and content

The Court of Justice consistently recalls that the choice of legal basis for a Union act must rest on objective factors amenable to judicial review, in particular the aim and content of the measure.

It is therefore not enough to call the DSS a “prudential contribution” in order to take it outside the fiscal sphere. Nor is it enough to declare that it promotes the internal market in order to found it on Article 114.

The legislative text has to be constructed so that its aims, obligations and financial mechanism are coherent with the competence invoked.

28.2 Article 114: what *Portigon* actually shows

In *Portigon v SRB*, concerning ex ante contributions to the Single Resolution Fund, the Union judicature held that the provisions imposing and governing those contributions do not constitute fiscal provisions within the meaning of Article 114(2) TFEU.

The reasoning gives weight to precise features: contributions paid by the operators of a sector, direct earmarking for that sector’s expenditure necessary to its functioning and stability, an insurance-like logic, and a mechanism built from a Fund target apportioned among institutions rather than from a rate applied to a taxable base.

These are precisely the features the DSS has to study if it wants to sustain a prudential characterisation.

28.3 Where the DSS is weaker than the SRF

The parallel breaks down at the most delicate point.

The SRF funds bank resolution tools, and the Court was able to describe the sectoral nexus in very tight terms. The DSS, by contrast, would fund training, transition income and territorial measures that may benefit people and entities outside the contributing sector.

The further the Fund moves from the risk and the expenditure needed to stabilise or make the sector resilient, the weaker the *Portigon* analogy becomes.

Chapter 23 on causation is therefore not only an economic requirement: it is also a possible component of the legal defensibility of the nexus between contribution and expenditure.

28.4 Article 175(3): the social and territorial precedent

On the benefits side there is another genealogy.

Both the EGF Regulation 2021/691 and the JTF Regulation 2021/1056 use, among their legal bases, the third paragraph of Article 175 TFEU, which allows specific actions to be adopted outside the funds referred to in the same title in order to pursue economic, social and territorial cohesion.

This precedent matters because the DSS has an avowedly social and territorial component. It shows that the Treaty already contains a basis used for European transition instruments.

It does not show, however, that Article 175(3) authorises, on its own, any mandatory contribution imposed on systemic operators.

28.5 A two-pillar architecture

The most prudent solution to put before a European legal service is to separate the two regulatory functions.

Pillar A — regulation of systemic functions: definition, Data Return, designation, mitigation, continuity, portability and, if sustainable in the light of the prudential nexus, the rules on the contribution. Article 114 is the natural candidate to test where the objective is to avoid internal market fragmentation and create uniform rules for cross-border operators.

Pillar B — transition benefits: eligibility criteria, coordination with the EGF/JTF/ESF+, territorial measures and additionality. Article 175(3) offers a direct precedent for specific cohesion actions.

Two pillars do not necessarily mean two separate acts. They mean that aims and competences have to be analysed separately before deciding whether a single legal basis, multiple bases or two coordinated instruments are sustainable.

28.6 The risk of Articles 113 and 115

The fiscal question cannot be settled by the name the legislature chooses.

Article 113 TFEU provides for a special procedure requiring unanimity for the harmonisation of turnover taxes, excise duties and other forms of indirect taxation to the extent necessary for the internal market. Article 115 likewise requires unanimity for

DSS component	Possible anchor to be studied	Strength	Legal risk
Data Return, designation, mitigation	Article 114 TFEU	Harmonisation of rules for cross-border functions	Must genuinely improve internal market conditions
Prudential contribution	Article 114 + SRF analogy, if the nexus holds	Portigon shows that not every levy is fiscal	Broad social spending may weaken the nexus
Social/territorial benefits	Article 175(3) TFEU	The EGF and JTF are legislative precedents	Does not automatically resolve the funding question
Specific labour provisions	Any dedicated social bases, to be assessed	Existing EU competences in social matters	Limits of conferral and reserved matters

directives approximating national provisions that directly affect the internal market where it operates outside Article 114.

If the DSS were designed as a percentage of revenue earmarked for general public expenditure, the risk of its being perceived and classified as a fiscal levy would increase. A Fund target, ring-fenced destination, risk adjustment, causation and reduction through mitigation are not merely good economic features: they are elements that distinguish the model from that architecture.

28.7 No advance certainty about characterisation

Even careful design does not allow this book to declare today that Article 114 will certainly be the correct basis.

Characterisation will depend on the final text, on the breadth of the benefits, on the relationship between contributors and beneficiaries, on the Fund's governance and on the degree of harmonisation required.

The responsible formulation is therefore: the SRF precedent makes a non-fiscal prudential route plausible, but the DSS has to be designed to pass a harder legal test precisely because it funds a cross-cutting social transition.

28.8 State aid on the expenditure side

A further problem arises where resources are not paid directly to individuals but fund firms, training bodies, infrastructure or territorial projects.

Public funding that confers a selective advantage on undertakings and affects trade may fall within Article 107(1) TFEU and require a compatibility basis, an applicable exemption or notification.

The design must therefore distinguish between benefits to individuals, purchase of services through competitive procedures, general schemes and selective aid to undertakings.

The DSS must not become, through the expenditure side, an opaque system of industrial subsidies.

28.9 Neutrality as to the operator's nationality

A regime built around the names of United States companies would be politically and legally fragile.

The criteria must apply to the function, the dependency and the risk irrespective of where the operator is established. A European, American or Asian company performing the same function within the same perimeter must be subject to the same indicators and the same procedure.

This neutrality does not by itself resolve questions of international trade law. But it is the minimum condition for showing that the DSS is a prudential instrument and not a selective levy against a geographical origin.

28.10 A realistic legislative sequence

Before a complete legislative proposal is presented, the DSS should undergo four legal checks.

First: an independent legal-basis opinion on Articles 114, 175(3) and possible combinations.

Second: analysis of the nature of the contribution in the light of *Portigon* and the case law on fiscal provisions.

Third: State aid screening of the principal categories of expenditure.

Fourth: verification of compatibility with international obligations and with the principle of non-discrimination for third-country operators.

These checks must precede, not follow, the political choice of legal basis.

28.11 The eighteenth principle

Principle: the legal basis is not the proposal's final technical footnote. It is part of the design.

The DSS has to be built so that the nexus between function, risk, contribution and the destination of the resources is economically verifiable and legally legible. The SRF precedent opens one door; the EGF and the JTF show the cohesion door. The legislative task is to work out whether and how these two doors can be connected without straining the Treaties.

CHAPTER 29

ASML: the European test of the upstream bottleneck

Until now the book's principal corporate cases have been Microsoft Azure, OpenAI and NVIDIA: three United States operators. That was a methodologically explicable choice, but a politically dangerous one. A model claiming to be neutral has to show that it works when the systemic function is performed by a European company too.

ASML is the most useful case. Not because it should be designated or should contribute to the DSS — a conclusion the public data does not support. It is useful because it tests the D dimension — control of infrastructure and inputs that are hard to replicate — at the most upstream point of the artificial intelligence chain: the capacity to produce the advanced chips on which compute, cloud and models all depend.

The analytical function is not "ASML as a company", but advanced EUV lithography capability together with the services, software and upgrades strictly necessary for its industrial use.

29.1 A European company, a global function

In 2025 ASML recorded total revenue of 32.7 billion euro, with 535 systems sold, of which 48 were EUV. In the second quarter of 2026 it reported 9.326 billion euro of revenue and raised its full-year 2026 guidance to 43–45 billion.

These figures document economic scale and the dynamics of demand. They do not on their own document systemic importance.

The qualitatively most significant characteristic is technological: ASML describes its EUV technology as one of a kind and uses it for the very high resolution lithography needed for the most advanced logic and memory nodes. The DSS therefore has to ask not what ASML is worth, but how substitutable that capability is in concrete industrial scenarios.

29.2 EUV and High-NA: the function, not the brand

The DSS perimeter should avoid two opposite errors.

If we define the function as "all of ASML", we take in DUV, metrology, software, services and activities with different risk profiles. If we narrow it to a single machine model, we may lose the dependency on the ecosystem of software, maintenance, upgrades and industrial qualification.

One possible unit of investigation is therefore "advanced EUV lithography capability", distinguishing Low-NA EUV from High-NA EUV where the data shows different profiles.

In 2026 ASML indicated Low-NA EUV capacity of around 65 systems and plans to increase it in 2027. The figure is useful for observing the scarcity of the input, but it is not equivalent to a market share, nor to a measure of the damage caused by unavailability.

29.3 D — infrastructural control

ASML is the case in which D has to prove its added value.

An input can be systemic when replacing it requires not merely purchasing an alternative good, but a different production architecture, process qualification, redesign, new skills and timescales incompatible with continuity of production.

Measuring D would require at least: actual production capacity by technology, lead times, the number of functionally equivalent alternatives, qualification times, dependencies on critical components, maintenance and upgrade capacity, and the share of advanced nodes for which the alternative would require a substantial change of process.

29.4 N — industrial substitutability is not the number of competitors

Saying a technology is “unique” is not enough to quantify N.

A chipmaker may have partial technical alternatives: multipatterning with earlier technologies, node redesign, deferring investment, moving to a different foundry, or changing the production mix. Each of these options carries different costs and timescales.

The Data Return should therefore build industrial switching scenarios: how long does it take to obtain equivalent capacity? How many wafers, and which nodes, can be produced without EUV? What does an alternative strategy cost? What performance is lost?

N has to measure concrete freedom to exit, not a reputation for technological uniqueness.

29.5 I and P — propagation along the supply chain

Dependency on an upstream input propagates through chains the supplier does not directly control.

Foundries and memory makers use lithography to produce chips; cloud providers and AI operators buy compute capacity; firms and administrations depend on the services built on that capacity.

The DSS must prevent this chain from being counted in full at every level. A delay in lithography capacity cannot be monetised once at ASML, once at the GPU maker, once in the cloud and once in the AI model if the underlying economic damage is the same.

Chapter 23 therefore makes a causal decomposition of propagation across supply chain levels essential.

29.6 Geopolitical data as evidence of criticality, not as a score

The Government of the Netherlands subjects specific semiconductor manufacturing technologies to national export controls. In 2024 it extended the licensing requirement to certain advanced DUV equipment; from 1 April 2025 it further widened the regime to other manufacturing and inspection technologies.

The official rationale refers to security risks, advanced military applications and the unique and leading role of the Netherlands in this sector.

For the DSS this is public evidence that the technology has strategic significance. It is not, however, a shortcut: export controls and DSS systemic importance serve different legal purposes. Designation must continue to depend on S-I-N-D-P-L, expected loss and functional data.

29.7 Sovereignty seen from the European side

The International AI Strategy published by the Netherlands in July 2026 describes ASML and ASM as strong elements of the European position in chip manufacturing and explicitly links AI to strategic dependencies, resilience and autonomy.

This case changes the political framing of the DSS.

The model is not only about asking foreign operators to fund European costs. It can also ask how Europe responsibly governs its own systemic strengths: what continuity guarantees it offers, how it reduces excessive dependencies, how it prevents its own infrastructural power from becoming a fragility for the European economy itself.

29.8 What the DSS would ask of ASML

A real investigation into the EUV function would require data that published accounts do not provide.

Capacity by technology family; lead times and backlog; average installation and qualification times; industrial alternatives by customer type and node; dependency on critical suppliers; availability of parts and technicians; recovery time after incidents; the effects of prolonged unavailability on downstream production; continuity and diversification measures.

Without this data we can document a plausible bottleneck; we cannot assign a real R score.

Dimension	Useful public evidence	Data missing for a real score
S	€32.7bn revenue 2025; strong growth	Revenue and capacity referable to the EUV function
I	EUV used in the advanced nodes of leading manufacturers	Downstream dependencies by customer and node
N	EUV technology described by ASML as unique to the company	Real times/costs of industrial alternatives
D	Advanced lithography capacity concentrated upstream	Equivalent capacity replicable and lead times
P	Potential propagation from fab to downstream chain	Quantitative scenarios without double counting
L	Employment effects mainly indirect	Task-based data; L must not be inferred from the sector

29.9 Why this case matters politically

ASML defuses a simple but powerful criticism: that the DSS is an anti-American digital tax under a new name.

If the same criteria can qualify a European function as potentially systemic, and if the same mitigation and safeguards apply irrespective of nationality, the model gains neutrality.

Naturally the case also creates a political difficulty: Europe would have to be prepared to apply to one of its own industrial champions the same criteria it would apply to United States operators. That is precisely the test of credibility.

29.10 The nineteenth principle

Principle: the operator's nationality does not enter the score.

A methodology that changes when the bottleneck is European is not a prudential methodology. It is selective industrial policy dressed up as risk measurement.

CHAPTER 30

Geopolitical neutrality: GATS, Section 301 and the risk of an anti-US tax

Methodological neutrality does not make geopolitics go away. A contribution imposed on cross-border digital functions can become the object of trade conflict even where the European legislature believes it has designed the measure reasonably.

The most instructive precedent is that of digital services taxes. The DSS is not a DST: it does not start from digital revenue, it does not single out a group of companies through tax thresholds, and it ties the burden to residual risk, causation, need and mitigation. But precisely for that reason it has to explain in advance which features make it different and how it avoids discriminatory effects.

The geopolitical question is therefore not “will the United States agree?” It is more demanding: can the structure of the measure be defended as origin-neutral, risk-based and consistent with the applicable international trade obligations?

30.1 The lesson of Section 301

Between 2019 and 2021 the U.S. Trade Representative conducted Section 301 investigations into the digital services taxes of several countries, including Italy. In 2021 USTR concluded, in a number of cases, that those DSTs discriminated against United States companies, were inconsistent with principles of international taxation, and burdened United States commerce.

USTR went as far as preparing additional tariff action, which was then suspended and subsequently terminated in the context of the political agreements on the transition towards the multilateral OECD/G20 solution.

These are unilateral determinations by the United States, not WTO rulings. But they show the geopolitical cost of a measure perceived as built around the characteristics of American companies.

30.2 The test the DSS has to pass

The DSS must be able to answer no to four questions.

Are the thresholds constructed so as to capture principally companies of one particular nationality?

Is an equivalent European function treated more favourably?

Is the mitigation available to domestic operators less accessible to foreign ones?

Is the need attributed to a sector because it generates risk, or because it offers an economic base that is politically easy to target?

The ASML case makes these questions testable rather than merely rhetorical.

30.3 GATS: national treatment is not only about formally identical rules

For functions falling within services covered by a WTO member's commitments, Article XVII GATS requires, in the sectors inscribed in the Schedule and subject to the conditions set out there, treatment no less favourable for foreign services and suppliers than for like domestic services and service suppliers.

The GATS text makes clear that even formally identical treatment can be less favourable where it modifies the conditions of competition in favour of domestic services or suppliers.

This is an important warning: drafting an apparently neutral rule does not guarantee that it is neutral in its effects.

The analysis must, however, be specific to the sector, the mode of supply and the commitments undertaken. This book cannot declare in the abstract that the DSS is compatible with GATS.

30.4 GATS is not the whole of the DSS's trade law

The DSS may concern functions that are services, goods, infrastructure or combinations of these.

Cloud and AI services raise issues principally in trade in services. ASML's lithography, by contrast, concerns industrial equipment and technology and calls for a different analysis, potentially involving WTO disciplines on goods, export controls and security.

The methodology must therefore avoid a second shortcut: there is no single "WTO test" valid for every sector. The trade assessment has to follow the concrete function.

30.5 Origin neutrality by design

I propose five design safeguards.

First: no S-I-N-D-P-L variable uses registered office, origin of capital or nationality.

Second: benchmarks are built by function and market, not by geopolitical grouping.

Third: an equivalent function performed by an EU operator undergoes the same screening.

Fourth: the Data Return, mitigation, caps, appeals and access to data are subject to the same conditions for domestic and foreign operators alike.

Fifth: every individual decision contains an origin-neutrality statement explaining why the outcome derives from functional data and not from the operator's origin.

30.6 Indirect discrimination and thresholds

The subtler risk is not a rule stating that "American companies pay". It is an apparently general threshold that, by design, falls almost exclusively on a foreign group.

The calibration protocol in Chapter 25 must therefore include a distributional audit by nationality, as well as by sector and size.

A concentration of outcomes on operators from a single country does not prove discrimination: it may reflect the real structure of the market. But it does require additional scrutiny of the variables producing that outcome.

The right response is more transparency, not artificial quotas by nationality.

30.7 Why the DSS must avoid the language of the digital tax

Legally and politically, calling the DSS a "tax on Big Tech" would destroy part of the book's methodological work.

That framing suggests that the entity comes before the function, that revenue comes before risk, and that economic size is the reason for the burden.

The coherent formulation is different: a prudential regime for systemic functions, and a transition Fund financed, within legally sustainable limits, by the share of residual risk causally attributed.

The difference has to emerge from the formulas and the safeguards, not from the vocabulary alone.

30.8 A Trade Impact Assessment before entry into force

Before the contributory phase, the proposal should be accompanied by a Trade Impact Assessment.

The document must map operators by origin, modes of supply, relevant WTO sectors, expected effects on trade and prices, differences in treatment, possible conflicts with international commitments, and the risk of retaliatory measures.

The analysis must include scenarios in which a trading partner challenges the measure even though the Union considers it non-discriminatory. The cost of retaliation is part of the cost of the policy and cannot be treated as an external variable.

30.9 Export controls and prudential contributions are different instruments

The ASML case also shows the other side of geopolitics.

The Netherlands uses export controls on specific semiconductor manufacturing technologies for security reasons. The DSS, by contrast, should not decide who may buy a technology, nor pursue national security objectives through the contribution.

Confusing the two instruments would create an opaque mechanism in which the prudential score becomes a geopolitical lever.

Cooperation between authorities is necessary, but the purposes must remain distinct: security and export control on one side; systemic dependency and transition on the other.

30.10 Geopolitics as a stress test of European credibility

The DSS will be credible outside Europe only if Europe accepts three disciplines.

Applying the same rules to its own operators.

Publishing the chain that links function, risk, euros and benefits.

Accepting that part of the risk goes unattributed where the causal evidence is insufficient.

These disciplines may make the model less lucrative and less aggressive in the short term. But they reduce the likelihood of its being perceived as fiscal protectionism in disguise.

30.11 The twentieth principle

Principle: the DSS must be origin-neutral, but not naive about trade effects.

The measure has to be able to survive two simultaneous forms of scrutiny: that of the prudential reviewer, who asks whether the risk is real, and that of the trading partner, who asks whether the conditions of competition are being altered in a discriminatory way.

PART IX

**GOVERNANCE, DUE PROCESS AND
CONTROL**

CHAPTER 31

Who decides, who manages, who supervises

A model that measures power must first of all limit the power of whoever does the measuring. In the DSS, governance is not an organisational detail: it is part of the model's validity. If the same body could choose the perimeter of the function, alter the parameters, determine the funding need and manage the Fund, the measurement of risk would be exposed to a structural conflict of interest.

The designation procedure must therefore be kept distinct from decisions on benefits and from financial management. Separation does not eliminate public discretion, but it makes it visible, contestable and subject to different forms of control.

31.1 Screening and opening the investigation

The first stage is screening. The authority identifies a potentially relevant function through public data, complaints, sector analysis or information from other authorities. Screening does not produce any contribution: it serves only to establish whether there is enough to open an investigation.

The decision to open must identify the function, the relevant territory and the preliminary reasons. In this way the operator knows from the outset what is being examined, and the authority cannot expand the perimeter indefinitely during the proceedings.

31.2 Data Return, case file and traceability

Once the investigation is opened, the operator receives the relevant DSS Data Return. The authority may supplement the information with independent sources, data from customers, competitors and public bodies, and material already held by other authorities, subject to confidentiality and the once-only principle.

Every variable used in the score must have a traceable provenance. The case file distinguishes observed data, estimates, inferences and scenarios. This distinction is fundamental, because verified data cannot be challenged in the same way as a modelled estimate.

31.3 Preliminary notice, hearing and decision

Before designation, the operator must receive a preliminary notice containing at least the perimeter of the function, the indicators used, the normalisation methodology, the score for each dimension, the joint requirement, the mitigation recognised and the proposed class.

The hearing must be substantive: access to the file, adequate time, the opportunity to challenge data and benchmarks, to propose alternatives and to demonstrate errors of

perimeter. The final decision must explain which evidence prevailed and which obligations follow from designation. Confidential information can be protected, but the logic of the decision must remain reviewable.

31.4 Three functions that must not coincide

The DSS contains at least three distinct institutional functions. The first assesses systemic importance and residual risk. The second determines benefits and the funding need in accordance with the law. The third physically manages the Fund's assets.

The risk assessor should not decide how many benefits to fund. The decision-maker on benefits should not alter scores in order to secure resources. The Fund manager should not be able to designate operators or expand benefits on its own initiative. This separation stops the need for revenue feeding back into the measurement of risk.

31.5 Independent control and asymmetry of resources

Alongside the separation of the three poles there must be independent control: financial audit, methodological review, outcome evaluation and judicial review are not the same thing and must be able to operate in parallel.

There is also a practical problem of asymmetry. A large operator may have more data, advisers and technical capacity than the authority assessing it. To stop the hearing becoming informational dependence on the regulated party, the authority must be able to build its own expertise, commission independent verification, document conflicts of interest and publish methodologies detailed enough to allow external replication.

31.6 Non-EU operators and the effectiveness of obligations

An information obligation with no enforcement tools would be weakest in precisely the most transnational cases. DORA offers a useful precedent: a critical ICT provider established in a third country must set up a subsidiary in the Union within twelve months of designation in order to continue supplying certain services to Union financial entities.

The DSS should not copy that solution automatically, but future legislation should provide for a legally responsible point of presence or representation in the Union, powers to request information, proportionate consequences for failure to cooperate and, in extreme cases, effects on market access for the regulated function. The concrete choice depends on the final legal basis.

31.7 Transparency, review and exit

Designation must be reviewed periodically and must be able to lapse when the facts change: loss of scale, new alternatives, greater interoperability, separation of the function, or structural mitigation. The Fund should publish aggregate data on contributions, benefits, administrative costs, reserves, funding gap and programme results.

A sunset clause and periodic review of the whole mechanism prevent weights, thresholds and institutions from becoming untouchable. The system must be designed so that a fall in risk is good news for those who administer it too.

31.8 The twenty-first principle

Principle: no contribution without a verifiable chain linking function, data, score, mitigation, residual risk and a reasoned decision.

The chain is the DSS's fundamental safeguard: it simultaneously limits the power of the operator and the power of the authority.

CHAPTER 32

The right to contest the algorithm

Every time a number produces legal consequences, that number has to be capable of being explained.

The DSS uses normalisations, weights, thresholds and coefficients. These tools can make a decision more transparent, but they can also create a new opacity if they are treated as incontestable technical truths.

The rights of the defence must therefore extend to the regulatory algorithm itself.

32.1 Knowing the formula

The general methodology must be public before it is applied.

Weights, thresholds, definitions, benchmarks and normalisation rules cannot be invented in the individual decision.

Substantive changes must follow a transparent procedure and, where appropriate, consultation and a transitional period.

32.2 Knowing your own data

The operator must know which data was used for its score and which items come from estimates or third parties.

Exceptions for confidentiality must be handled without hollowing out the rights of the defence.

Where a source cannot be disclosed in full, the authority must still provide enough information for its reliability and significance to be challenged.

32.3 Contesting the normalisation

Two authorities can start from the same switching time and assign it very different scores if they use different curves.

Normalisation is therefore a substantive regulatory choice.

The operator must be able to challenge not only the raw data but also the transformation that converts it into a value from zero to one hundred.

32.4 Sensitivity analysis

For cases close to the thresholds, the decision should include a sensitivity analysis.

What happens if an uncertain indicator varies within its plausible range? What happens if a benchmark is updated?

A designation that disappears under a minimal variation calls for more rigorous reasoning than one that is robust across every reasonable scenario.

32.5 Auditability

The model must leave a reproducible trail.

An independent reviewer, given the same data and the same version of the methodology, should be able to reconstruct the score.

Versioning of the formulas, change logs and retention of sources are therefore legal requirements, not merely technical ones.

32.6 Appeal

The final decision must be subject to judicial review under the applicable legal order.

The court need not necessarily substitute its own technical assessment for the authority's.

It must, however, be able to review competence, procedure, accuracy of the facts, coherence of the methodology, proportionality and reasoning.

32.7 No black box

The DSS does not require machine learning to work.

Indeed, in its initial phase a relatively simple and public formula may be preferable to an opaque predictive model.

Sophistication is not a value if it reduces contestability.

The system can evolve, but every increase in complexity must be justified by a demonstrable increase in the quality of the measurement.

32.8 The twenty-second principle

Principle: no regulatory algorithm may be harder to contest than the decision it produces.

With this safeguard the basic institutional architecture is complete.

The next step is to return to people: who receives the Systemic Social Dividend, under what conditions, and through which index of vulnerability.

PART X

WHO IS ENTITLED TO THE DIVIDEND

CHAPTER 33

Who receives the DSS: need and the geography of transition

Having defined who contributes, the mirror question is who receives. The DSS is neither a universal income nor civil compensation. It is a transition instrument for people and territories in which a systemic transformation produces documentable vulnerability.

Vulnerability may be individual — loss or reduction of income, accelerated obsolescence of skills, difficulty in finding new work — or territorial, where a single supply chain concentrates employment, income and opportunity in an area with few alternatives.

33.1 The link with transition

Access should not depend on the impossible proof that a particular algorithm caused a particular dismissal. Economic transformations are multi-causal. What is needed instead is an administratively verifiable link between the beneficiary's situation and a sector, occupation or territory affected by a recognised transition.

The contribution arises from the function's systemic risk; the benefit arises from the need for transition. Separating the two planes allows the Fund to act before individual harm becomes irreversible.

33.2 Prevention before unemployment

Confining the DSS to people who have already lost their jobs would be too late. Someone may still be employed yet working in a highly exposed occupation, with skills depreciating rapidly and little scope for internal redeployment. The Fund must be able to finance training, mobility and support before job loss where the indicators show a plausible and verifiable transition.

33.3 Temporary benefits and coordination with welfare

Individual benefits must have a defined duration and periodic review. Transition income, training, redeployment and mobility support meet different needs and must not turn into permanent transfers without an objective.

The beneficiary may also be entitled to national unemployment benefits or European programmes. The rule is additionality: no double payment of the same item, but the possibility of covering costs ordinary instruments do not cover, or of increasing the intensity and timeliness of inadequate provision.

33.4 When the shock is territorial

A systemic transition can strike people and places at the same time. Where much local employment depends on the same supply chain, the loss of activity reduces wages, local demand, investment and the scope for redeployment. The geography of vulnerability therefore enters the DSS not in order to attribute more blame to the operator, but to establish where the same transition costs more.

33.5 Concentration and absorptive capacity

The first territorial indicator is concentration: how much local employment and income depend on the affected activities? The second is absorptive capacity: vacancies, growth in other sectors, distance from opportunities, transport, training and the scope for remote work.

Territory must not be treated as destiny. The index has to measure precisely the levers that can change it, so that intervention funds new capacity rather than certifying permanent disadvantage.

33.6 Skills and real access to training

Differences in digital skills and in participation in training make a policy that simply offers everyone the same course regressive. The DSS must take account of available time, guidance, equipment, accessibility, financial support and distance from training opportunities.

33.7 Priority Transition Zones

The legislation could allow the temporary identification of Priority Transition Zones. The designation does not automatically allocate money to an area: it enables accelerated procedures, greater programme intensity, co-financing of training infrastructure and more frequent monitoring.

Each zone must have quantitative criteria, a limited duration, diversification objectives and an exit strategy. A vulnerable territory must not become dependent on the Fund.

33.8 The twenty-third principle

Principle: the DSS selects for need, not for worth. Economic and functional selectivity serves to concentrate resources and must never become a judgement on a person's value.

33.9 The twenty-fourth principle

Principle: where a person lives may increase the cost of transition, but policy must use that information to widen possibilities, not to entrench inequality.

CHAPTER 34

The INDS: measuring need

Distributing limited resources requires a methodology on the beneficiary side too.

I propose to call it the INDS: the Social Dividend Need Index.

The INDS must not predict a person’s future. It has to rank, on transparent criteria, the intensity of the need for transition.

As with the operator score, the weights that follow are a proposal to be validated empirically. They are not scientifically demonstrated values.

Dimension	Meaning	Proposed initial weight
V — Income vulnerability	Significant loss or risk of loss of income	30%
E — Exposure of the occupation	Intensity of the transformation of tasks	20%
T — Transferability of skills	Distance between skills held and skills in demand	20%
C — Capacity to adapt	Access to training, time, resources and opportunities	15%
R — Territorial vulnerability	Weakness of the local labour market	15%

$INDS = 0.30V + 0.20E + 0.20T + 0.15C + 0.15R$. The weighting is illustrative and will have to be validated empirically before any regulatory application.

Each dimension is normalised from 0 to 100, where a higher value indicates greater need for intervention.

The formula must be interpreted with caution: the INDS cannot become a total social score for the individual. It must use only variables relevant to transition.

34.1 V — Income vulnerability

V looks at the capacity of the person or household, according to the legislative choice, to absorb a temporary loss.

It may consider the reduction in earned income, its expected duration, the presence of other coverage and economic thresholds set by law.

The model must avoid collecting unnecessary wealth data. Informational proportionality applies to welfare too.

34.2 E — Occupational exposure

E measures how far the tasks of the occupation are affected by the transformation. Where relevant it uses the same task-based infrastructure described in Chapter 26, but on a different plane: L measures the pressure produced by the systemic function; E measures the person's exposure for the purpose of prioritising intervention.

It must not be calculated by asking the person whether they "fear AI".

It can use occupational classifications, sector data, task analysis and employer information, updated periodically.

Exposure is not equivalent to substitution: a highly exposed occupation may be transformed rather than eliminated.

34.3 T — Transferability of skills

T measures how easily the skills a person holds can be used in accessible occupations.

Someone with highly transferable skills may need only a short intervention.

Someone whose experience is concentrated in rapidly declining activities may need a more intensive pathway.

The system should give weight to real skills, including those acquired on the job, and not only to formal qualifications.

34.4 C — Adaptive capacity

C considers the concrete obstacles to transition: availability of relevant training, the time required, compatibility with personal responsibilities, digital access and the presence of realistic opportunities.

This dimension must be built without discrimination.

Age, gender, disability or other protected characteristics must not become shortcuts for presuming lower capacity. They may be taken into account, where legally permitted, only in order to remove obstacles and ensure appropriate accommodations or support.

34.5 R — Territorial vulnerability

R introduces context.

Losing a job in a dynamic labour market is not equivalent to losing it in an area where the same sector accounts for most of the employment and alternatives are scarce.

R can use the employment rate, vacancies, sectoral diversification, redeployment times, accessibility of training services and the concentration of exposed activities.

34.6 Thresholds and benefits

The INDS must not produce a single automatic benefit.

Different bands can trigger different instruments: guidance and light training; intensive training; mobility support; temporary income top-up; or combinations of these.

The thresholds must be validated on data and subject to review.

The principle is to grade the intensity of intervention, not to create rigid categories of citizens.

34.7 European data that justifies the problem

The need for a transition index does not arise in a vacuum.

Verified source: Eurostat, *Digitalisation in Europe — 2026 edition* and *Skills for the digital age*.

In 2025 only 60 per cent of EU citizens aged 16 to 74 had at least basic digital skills, against a European target of 80 per cent by 2030. In the same year, 20 per cent of EU firms were using artificial intelligence technologies.

The speed of technological adoption and the uneven distribution of skills make it plausible that the same transformation produces very different adaptive capacity across people and territories.

This data describes the context; it does not by itself establish who should receive the DSS.

34.8 Validation

Before any regulatory use, the INDS must undergo, *mutatis mutandis*, the same calibration protocol as Chapter 25: construct quality, multivariate analysis, normalisation, elicitation of weights, alternative specifications, Monte Carlo and sensitivity analysis.

It must be checked on out-of-sample data whether it does in fact identify people with longer redeployment times, greater income losses or more intensive training needs — avoiding validation of the index against the very outcomes already embedded in its components.

If it does not, indicators, normalisations, weights or even the structure of the index may have to change. The evaluation must also check for systematic errors across groups and territories, without using protected characteristics as predictive shortcuts.

The purpose of the index is to predict operational need well enough to allocate resources, not to defend the formula proposed in this book.

CHAPTER 35

How the DSS reaches the individual

A good benefit can fail through a bad procedure.

If obtaining the DSS takes months, requires documents the administration already holds and demands high digital skills, the system will implicitly select the people best able to navigate bureaucracy.

The design of delivery is therefore a substantive part of the right.

35.1 Multi-channel access

Applications must be possible digitally, but human assistance has to exist.

Public counters, employment services and accredited bodies can support those without sufficient digital skills or equipment.

A Fund created to manage the digital transition cannot exclude those who suffer from the digital divide.

35.2 Pre-filling

Where legally possible, the application should be pre-filled using administrative data already available.

The citizen checks, supplements and corrects.

This reduces errors and access costs and applies on the social side the same once-only principle proposed for operators.

35.3 An explainable decision

The individual must receive a comprehensible decision.

Communicating "INDS 63" is not enough.

What is needed is an explanation of which factors determined the result, which benefits are available, how long they last, what obligations they entail and how to challenge any errors.

35.4 Choice between pathways

Where several interventions are appropriate, the beneficiary should be able to take part in the choice.

One person may prefer training, another mobility, another a route into self-employment.

The DSS must not turn the index into an algorithm that automatically assigns a working life.

35.5 Payment

Cash benefits must be traceable, timely and separate from payments to service providers.

Training and placement can be paid directly to accredited bodies on the basis of rules and results.

Transition income, by contrast, must reach the individual on a predictable schedule.

35.6 Proportionate controls

The Fund has to prevent fraud and double payment.

But control must be proportionate to the risk and the amount.

Excessively punitive systems can generate administrative costs greater than the fraud they seek to avoid, and can deter legitimate beneficiaries.

35.7 Rapid review

Errors in welfare carry a high cost in time.

An appeal won after two years may be formally correct and socially useless.

The DSS should provide for rapid administrative review of errors in data, occupational classification, income or territory, without prejudice to judicial protection.

35.8 Data protection

The INDS must not become a permanent database of personal vulnerability.

Minimisation, purpose limitation, security, retention periods and separation between data needed for delivery and data used for statistical evaluation must all apply.

Automated decisions, where used, must comply with the safeguards laid down by the applicable legal order.

35.9 Measuring outcomes

The Fund needs to know what happens next.

Not in order to monitor the beneficiary indefinitely, but to evaluate the programmes.

Time to redeployment, subsequent income, completion of training, job stability and satisfaction can all help establish which interventions work.

Evaluation data must be used proportionately and, where possible, in aggregate form.

35.10 The twenty-fifth principle

Principle: a right is effective only if the procedure for exercising it is simpler than the problem it is meant to solve.

With the INDS and the delivery architecture, the DSS now has both sides: those who contribute and those who receive.

What remains is the final political test: how to turn this architecture into a legislative proposal that can be debated, amended, piloted and, if it works, adopted.

Essential statistical sources for Part X

Eurostat, *Digitalisation in Europe — 2026 edition*: in 2025, 20% of EU firms were using AI technologies; 60% of citizens aged 16–74 had at least basic digital skills.

Eurostat, *Skills for the digital age*, 2025 data: at least basic digital skills held by 82% of people with high educational attainment and 38% of those with low or no formal education.

Eurostat, *Adult learning statistics*: European data on adult participation in education and training.

European Commission, *European Skills Agenda*: European objectives and instruments for upskilling and reskilling.

PART XI

FROM THEORY TO IMPLEMENTATION

CHAPTER 36

From model to law: who has to do what, and why

A book can propose a regulatory architecture. It cannot pretend to stand in for the legislature. The distinction matters: drafting thirty article headings without doing the necessary evidential, legal and political work gives the impression of a finished law when what actually exists is a model still to be turned into positive law.

The move to legislation is necessary because some components of the DSS cannot be voluntary: the Data Return obligation, designation of systemic functions, investigative powers, contributions, ring-fenced allocation of resources, penalties, rights of defence and judicial protection all require a statutory basis.

36.1 Who starts the European process

For a European regime, the ordinary route starts with the European Commission, which holds the right of legislative initiative and prepares the proposal. Where the expected impacts are significant, preparation normally includes problem analysis, assessment of options, stakeholder consultation and an impact assessment. The proposal then goes to the European Parliament and the Council, which under the ordinary legislative procedure decide on an equal footing.

A national parliament, the European Parliament, a government or a network of authorities can promote the DSS politically, commission studies, pass resolutions or run experiments. But a uniform regime for cross-border operators does not emerge from a book or from a single administration: it requires a formally competent proposal and a legislative negotiation.

36.2 Why the Commission should take it up

The reason is not "to create a new tax". It is to establish whether there is a European problem that current rules do not fully solve: cross-border dependencies, exit costs, concentration risks and residual social costs of transition. If the analysis shows the problem can be better addressed with existing instruments, the DSS should stand down. If it shows a gap, then there is a reason to consider intervention.

36.3 What has to be in the legislative act

The basic act must fix the elements that cannot be altered by a technical decision: objectives, the definition of a systemic function, essential designation criteria, rights of defence, the principles of the Data Return, recognition of mitigation, the nature and

destination of the Fund, categories of benefit, limits on the contribution, governance, transparency, appeals, review and conditions for termination.

The legislation must also establish who may decide what. An authority should not be able to change the meaning of systemic importance case by case; the Fund manager should not be able to expand benefits; the political decision-maker should not be able to alter a score in order to hit a revenue target.

36.4 What must not be in primary legislation

Numerical weights, normalisation curves, detailed Data Return formats, benchmarks, Monte Carlo protocols, stress test specifications and reporting templates must all be capable of being updated as data and technology change. The correct legal level — technical annex, delegated act, implementing act or authority standard — will depend on the legal basis and on the delegation contained in the final act.

Articles 290 and 291 TFEU distinguish, respectively, delegated acts that supplement or amend non-essential elements from implementing acts intended to ensure uniform conditions of application. The book cannot choose in the abstract which instrument is correct for each parameter: it can, however, distinguish essential political elements from updatable technical ones.

36.5 Who does the technical work

The proposal would require an interdisciplinary working group: Commission services responsible for the internal market, competition, digital, employment and cohesion; supervisory authorities and regulators where relevant; independent statistical and scientific expertise; and representatives of companies, users, workers and public administrations.

The technical task is not “to write the tax” but to test the model against real data: define the perimeter of functions, verify the availability and quality of data, estimate administrative costs, test scores and expected loss, measure pass-through, check additionality and build enforcement scenarios.

36.6 The clauses that prevent fiscalisation

Legal characterisation does not depend on the name given to the contribution, but design can stop the DSS turning into a general revenue mechanism. Resources must be ring-fenced; the funding need must precede the contributory demand; verified mitigation must reduce the burden; reserves must have a ceiling; caps must be public; the funding gap must remain visible.

The legislation must also prohibit political revenue targets as a criterion for changing scores and thresholds, require a causal link between cost and sector, impose an Incidence Test where the contribution may be passed on to customers, and apply the same criteria to European and non-European operators alike.

36.7 What this book actually delivers

The book does not deliver draft articles ready for the Official Journal. It delivers a policy specification: it defines the problem, proposes a designation criterion, sets out

which data is needed, builds the bridge between risk and euros, identifies the prudential precedents, states the legal limits and proposes conditions of falsifiability.

That is the right level for an essay. The legislature can use the structure to prepare a call for evidence, an impact assessment, a pilot project or a proposal. The articles come later, once the legal basis, competent authorities, implementing instruments and validated parameters have been chosen.

36.8 The twenty-sixth principle

Principle: the DSS collects because a residual risk and a verified need exist; it does not create risk and need because it wants to collect.

Legislation serves to make that inversion difficult not merely in the author's intentions, but in the institutional architecture.

CHAPTER 37

Pilot first, extend later: the roadmap

A reform combining algorithmic classification, information obligations, contribution and social benefits should not come into force in a single day. Gradualism is not weakness: it is a mechanism for controlling regulatory risk.

The DSS has to be piloted in a way that allows it to fail. If the pilot serves only to confirm a decision already taken, it is not a test but a legitimization ceremony.

37.1 Phase zero: preparing without levying

Before the pilot, what is needed is an independent legal-basis opinion, State aid and international trade screening, cooperation protocols, consultation on the technical annexes and data infrastructure. At this stage there is no revenue target.

37.2 Four functional profiles

The pilot should use a small number of very different profiles: cloud, AI services or models, compute capacity and advanced lithography. The aim is not to designate four companies in advance, but to check whether the same architecture distinguishes an infrastructural service, a model, compute and an upstream industrial bottleneck — including at least one European case.

37.3 A year in shadow mode

For twelve months the authority can calculate gross and residual scores, class probabilities, expected social loss and theoretical contributions without imposing any payment. The results are communicated to operators and put to adversarial process.

The shadow phase must also test the monetary component: scenarios, probabilities, conditional losses, causal shares of the need, funding gap and Incidence Test. This shows what the system would have demanded before any real cost is imposed.

37.4 Social pilot and the INDS

In parallel, the INDS can be applied to transition programmes already funded from public or experimental resources. The score is compared against observable outcomes such as income loss, duration of redeployment, completion of training and job stability.

37.5 Success criteria and stop criteria

The pilot is not judged by the money it could have raised. The criteria are verifiability of data, stability of classes, accuracy of loss estimates, robustness to alternative specifi-

cations, incidence on customers, origin neutrality, administrative cost, effectiveness of mitigation, and the INDS's ability to identify real needs.

The law or act authorising the experiment must expressly provide for the possibility of not activating a component of the system. A negative result is useful information, not a political failure to be concealed.

37.6 From pilot to limited contribution

Only after validation can a limited contribution be considered, for functions and sectors that have passed the methodological, causal, incidence and neutrality tests. Prudential caps and a modest initial Fund reduce the cost of any regulatory error.

37.7 Review in the third year

The first review compares projections against outcomes: classes, class probabilities, expected loss, actual costs, causal attribution, appeals, contributions, funding gap, social results, pass-through, prices, investment and trade effects. The legislature amends what does not work.

37.8 Conditional extension

In subsequent years, extension to new sectors or new benefits is not automatic. It must depend on the evidence. The evolving European framework — the Data Act, the AI Act and other regimes — may reduce some risks: the DSS must record such mitigation and not duplicate obligations that are already effective.

37.9 The destination

After some years the DSS should be able to answer with data questions that are today theoretical: which functions produce systemic dependencies, how far mitigation reduces risk, what transition really costs, which benefits work, and what share can sustainably be funded by operators.

If the system cannot answer, it has not learned enough. If it can, it has become an infrastructure of public knowledge before it is a financial mechanism.

CHAPTER 38

From social dividend to a new compact

We began with a question: who pays the cost of power?

The simplest answer would have been to name a group of large companies and apply a rate.

This book chose a harder road.

It has tried to distinguish power from size, risk from revenue, prevention from compensation, need from universality, and politics from measurement.

The result is the Systemic Social Dividend.

38.1 What the DSS is not

It is not a tax on artificial intelligence.

It is not a robot tax.

It is not a levy on Big Tech revenue.

It is not a universal basic income.

It is not a penalty for economic success.

It is a mechanism connecting a systemic function, the residual risk that function produces, and the verified cost of a social transition.

38.2 Before redistribution: resilience and autonomy of choice

In its mature form, the DSS does not arise from the idea of distributing part of the wealth produced by technology. It arises from the problem of dependency.

Before the Fund come the ability to leave, operational continuity, interoperability, diversification and the capacity to absorb a disruption. A society that can change supplier and keep functioning is more resilient and retains greater autonomy of choice.

Redistribution comes in only when, despite these measures, a social cost remains that is causally attributable and sufficiently documented. In this sense the DSS is first a resilience policy and only second a mechanism of mutualisation.

38.3 What it asks of companies

The DSS asks first of all that they make their power less costly for society.

More portability. More interoperability. More resilience. More scope for exit. More transition instruments.

Only the risk that remains enters the contributory mechanism.

This makes investment in prevention part of the operator's economic strategy.

38.4 What it asks of the state

Of the state it asks something equally difficult: measure before you levy.

Make the rules public.

Accept challenge.

Do not mistake the need for resources for proof of responsibility.

Evaluate the results and shrink the Fund when the risk falls.

38.5 What it promises people

To people it does not promise that nothing will change.

Such a promise would be false.

It promises instead that transformation need not mean abandonment.

Time to learn, resources to get through a loss of income, the chance to find new work, and territories able to respond can turn fear of technology into the capacity to choose.

38.6 The price of power

Contemporary economic power does not show itself only when someone can set a price.

It also shows itself when millions of people, businesses and institutions build activities that become hard to move, replace or interrupt.

That power is not necessarily unlawful.

But it creates responsibility when its interruption or transformation shifts systemic costs onto society.

The price of power is responsibility for the risk that could not reasonably be eliminated.

38.7 A falsifiable proposal

The DSS does not ask to be believed.

It asks to be tested.

The weights may be wrong. The thresholds may be too high or too low. The INDS may need different variables. The caps may have to be rethought.

The strength of the proposal lies in the fact that these hypotheses can be checked against data.

A theory that accepts it might be corrected is more useful than a slogan that claims to be right already.

38.8 Conclusion

The question is not a choice between technology and people.

Technology is already a human creation, and its power can enormously widen collective possibilities.

The question is how to distribute responsibility when that power becomes infrastructure, dependency and transformation.

If those who provide a systemic function can reduce their contribution by reducing risk; if the state can spend only what it can justify; if those going through the transition can receive real tools to come out of it stronger — then the dividend is not simply money.

It is a share of power returned to the capacity to choose, to change, to withstand a shock and to cross the transition without depending on a single point of failure.

And that, in the end, is the proposal of the Systemic Social Dividend.

EPILOGUE

Power must be able to diminish

A good institution does not merely distribute the cost of risk. It tries to make that risk less necessary.

This is why the DSS contains, from the outset, its own route to reduction.

If functions become more substitutable, infrastructures more resilient, workers more able to get through change and territories less fragile, residual risk falls.

And with it, the price of power must fall too.

This is perhaps the most important difference between a tax and the mechanism proposed in these pages.

A tax may want a stable taxable base.

The DSS must want its own risk base to shrink.

The day a hugely powerful function can be replaced without trauma, an incident can be isolated without contagion, and a technological transformation can be crossed without shutting people out of the ability to choose — on that day the Fund will be able to collect less.

That will not be a defeat.

It will be its finest result.

FINAL APPARATUS

Note on sources and method

The book combines legislative and institutional sources, publicly available company data and methodological simulations. The three categories are not equivalent and are kept distinct.

The company data cited in the Microsoft Azure, OpenAI, NVIDIA and ASML cases serves to document the scale and observable characteristics of the functions analysed. It does not, on its own, constitute a designation of systemic importance.

The numerical scores used in the examples are illustrative simulations. The DSS and INDS weights, thresholds, classes and coefficients must be subjected to empirical validation before any regulatory use.

The figures of 10, 25 and 40 billion euro are stress scenarios for the funding need, not established revenue estimates.

Transparency note on drafting: generative artificial intelligence tools were used in preparing the manuscript to support drafting, revision, organisation and consistency checking. The selection of the arguments, the design of the model, the editorial decisions and final responsibility for the content remain the author's. The data and sources used for verifiable statements are kept distinct from the simulations and the methodological proposals.

Essential bibliography and institutional sources

NVIDIA Corporation, Form 10-K, fiscal year ended 25 January 2026, filed with the U.S. Securities and Exchange Commission on 25 February 2026.

OpenAI, *1 million business customers: the fastest-growing business platform in history*, 5 November 2025; *Scaling AI for everyone*, 27 February 2026; *The next phase of enterprise AI*, 8 April 2026.

Microsoft Investor Relations, FY2026 Q4 Earnings Release and Earnings Conference Call, 29 July 2026; 2025 Annual Report.

Eurostat, 11 December 2025, *20% of EU enterprises use AI technologies*, dataset isocEb_ai.

Eurostat, *Skills for the digital age*, data extracted January 2026.

Eurostat, *Digitalisation in Europe — 2026 edition*.

European Commission, *General-purpose AI obligations under the AI Act; Guidelines for providers of general-purpose AI models*. Accessed: 14 August 2026.

European Commission, 25 June 2026, *Commission reaches preliminary position that Amazon's and Microsoft's market leading cloud services should be designated under the DMA*.

European Union, Regulation (EU) 2024/1689 — Artificial Intelligence Act.

European Union, Regulation (EU) 2023/2854 — Data Act.

European Union, Regulation (EU) 2022/1925 — Digital Markets Act.

European Union, Regulation (EU) 2022/2554 — Digital Operational Resilience Act (DORA), Articles 31–35.

European Commission, Delegated Regulation (EU) 2024/1502, criteria for designating critical third-party ICT providers under DORA.

European Supervisory Authorities, 18 November 2025, first designation of critical ICT third-party providers under DORA.

Basel Committee on Banking Supervision, Basel Framework, SCO40 — Global systemically important banks; assessment methodology and reporting instructions.

European Union, Regulation (EU) No 806/2014 — Single Resolution Mechanism and Single Resolution Fund, Articles 69–70.

European Union, Implementing Regulation (EU) 2015/81 on ex ante contributions to the Single Resolution Fund.

General Court of the European Union, Eighth Chamber (Extended Composition), 29 May 2024, *Portigon v Single Resolution Board*, T-360/21.

Basel Committee on Banking Supervision, Basel Framework, CRE35 — IRB approach: treatment of expected losses and provisions; CRE31–32 — risk components.

OECD / European Commission Joint Research Centre, *Handbook on Constructing Composite Indicators: Methodology and User Guide*, OECD Publishing, 2008.

European Commission, Joint Research Centre, *Counterfactual Impact Evaluation: methodology and guidebooks for causal policy evaluation*.

Single Resolution Board, SRF levies (ex-ante contributions): target level, reporting framework, size and risk adjustment.

Commission Delegated Regulation (EU) 2015/63 on ex-ante contributions to resolution financing arrangements.

Hebous, Shafik; Brendan Crowley; Rasmi Das; Tibor Hanappi; Cory Hillier; Adam Jakubik; Eric Robert; Christophe J. Waerzeggers, *Taxing Cross-Border Services*, IMF Working Paper WP/26/152, July 2026.

Parlapiano, F. et al., *The pass-through of cost shocks to firms' prices and the impact of firms' balance sheets*, ECB research paper, 2024–2025.

Duprez, C. and Magerman, G., *Price Updating in Production Networks*, Working Paper Research No. 352, National Bank of Belgium / ECB research network: evidence on incomplete pass-through and mark-up adjustment.

European Commission Joint Research Centre, *Tools for Composite Indicators Building*, EUR 21682 EN.

European Commission Joint Research Centre, *Uncertainty and Sensitivity Analysis for policy decision making*, 2020.

United Nations Development Programme, *Human Development Report Technical Notes: the HDI as a geometric mean of normalised dimensional indices*.

Cooke, R. M., *Classical Model for Structured Expert Judgment*; applications and validation of the calibration and information scoring method.

Acemoglu, D. and Restrepo, P., *Artificial Intelligence, Automation and Work*, NBER Working Paper No. 24196, 2018.

Acemoglu, D. and Restrepo, P., *How Technology Displaces and Reinstates Labor*, NBER Working Paper No. 25684, 2019.

Eloundou, T., Manning, S., Mishkin, P. and Rock, D., *GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models*, arXiv:2303.10130, 2023.

Felten, E., Raj, M. and Seamans, R., *Occupational, industry, and geographic exposure to artificial intelligence: a novel dataset and its potential uses*, Strategic Management Journal, 42(12), 2195–2217, 2021.

Webb, M., *The Impact of Artificial Intelligence on the Labor Market*, working paper, 2019.

OECD, *Artificial intelligence and the changing demand for skills in the labour market*, OECD Artificial Intelligence Papers, 2024.

Regulation (EU) 2021/691 on the European Globalisation Adjustment Fund for Displaced Workers, as amended by Regulation (EU) 2026/1139.

Regulation (EU) 2021/1056 establishing the Just Transition Fund.

Consolidated version of the Treaty on the Functioning of the European Union, Articles 107, 113, 114, 115 and 175.

Regulation (EU) No 806/2014 establishing the Single Resolution Mechanism, Article 70 on ex-ante contributions.

ASML Holding N.V., *Annual Report 2025*: total net sales EUR 32.7 billion; 535 systems sold, including 48 EUV systems.

ASML Holding N.V., *Q2 2026 Financial Results*, 15 July 2026: Q2 net sales EUR 9.326 billion; 2026 outlook EUR 43–45 billion; Low-NA EUV capacity planning.

ASML, EUV lithography systems and lithography principles: EUV technology described by ASML as unique to the company and used for leading-edge manufacturing.

Government of the Netherlands, export controls on advanced semiconductor manufacturing equipment, measures of 7 September 2024 and 1 April 2025.

Government of the Netherlands, *The Netherlands' International AI Strategy*, 3 July 2026: strategic dependencies, resilience and the role of Dutch semiconductor manufacturing equipment.

World Trade Organization, General Agreement on Trade in Services, Article XVII — National Treatment.

United States Trade Representative, Section 301 — Digital Services Taxes: investigations and actions concerning Austria, India, Italy, Spain, Turkey and the United Kingdom, 2019–2021.

United States Trade Representative, political compromise and transition arrangements concerning digital services taxes, 2021–2024.

European Commission, *Decision-making process in the Commission*: right of legislative initiative under Article 17(2) TEU.

European Commission, *Law-making process; Planning and proposing law; Impact assessments; Better Regulation Guidelines and Toolbox*.

European Parliament, *Ordinary Legislative Procedure — Overview*: Commission proposal and joint adoption by Parliament and Council.

TFEU, Articles 290 and 291: delegated acts and implementing acts.

Regulation (EU) 2022/2554 (DORA), Article 31(12): requirement for a subsidiary in the Union for CTPPs established in third countries, within the DORA perimeter.

Register of verifiability for the principal figures

ASML 2025: total net sales €32.7 billion; 48 EUV systems sold; Q2 2026 €9.326 billion and 2026 outlook €43–45 billion — VERIFIED COMPANY DATA. ASML Investor Relations.

INDS and its weights — METHODOLOGICAL PROPOSAL. Index proposed in this volume; requires validation.

EU 2025: 60% of people aged 16–74 with at least basic digital skills — VERIFIED STATISTICAL DATA. Eurostat.

EU 2025: 20% of firms using AI — VERIFIED STATISTICAL DATA. Eurostat, dataset isoc_eb_ai.

NVIDIA FY2026 \$215.938 billion; Data Center \$193.737 billion — VERIFIED REGULATORY DATA. Form 10-K filed with the SEC.

OpenAI enterprise above 40% of revenue, April 2026 — VERIFIED COMPANY DATA. OpenAI, 8 April 2026.

OpenAI more than 1 million business customers and more than 7 million ChatGPT for Work seats (5 November 2025); more than 9 million paying business users (27 February 2026) — VERIFIED COMPANY DATA. OpenAI.

Azure above \$100 billion in annual revenue FY2026; FY2026 growth +41%; Q4 Azure and other cloud services +43% — VERIFIED COMPANY DATA. Microsoft Investor Relations, 29 July 2026.

Numerical scores in Chapter 19 — SIMULATION, ARITHMETICALLY VERIFIED. Cloud 72.25; AI Model 67.50; Compute 77.90; these are not real scores for Microsoft, OpenAI or NVIDIA.

S-I-N-D-P-L weights and DSS thresholds — METHODOLOGICAL PROPOSAL. To be calibrated and validated on real data.

Scenarios of €10/25/40 billion — SIMULATION. Stress test of the funding need; not a revenue forecast.

Task-based L architecture — METHODOLOGICAL PROPOSAL ANCHORED IN THE LITERATURE. Exposure, adoption, substitution/complementarity and reallocation distance must be calibrated and empirically validated.

EGF 2026: support for imminent displacement — VERIFIED LEGISLATIVE DATA. Regulation (EU) 2026/1139 amends Regulation (EU) 2021/691.

DSS legal basis under Articles 114/175 — LEGAL HYPOTHESIS TO BE VALIDATED. *Portigon*, the EGF and the JTF are relevant precedents but do not automatically determine the DSS's legal basis.

Notice to the reader

The Systemic Social Dividend presented in this volume is a proposal for an economic, regulatory and social architecture. The text does not attribute legal liability to the companies used as case studies and does not substitute for assessments by the competent authorities.

AUTHOR'S NOTE

This book grew out of a question that is simple and, at the same time, hard to turn into a rule: when an economic or technological function becomes so important that it shifts a significant part of the cost of its own interruption, replacement or transformation onto society, who should bear that cost?

The Systemic Social Dividend is the answer I propose.

It does not start from the idea that technology should be slowed down, nor from the conviction that economic size is in itself a fault. It starts from a different criterion: the greater the systemic dependency a function creates, the greater must be the responsibility for the residual risk that remains after the mitigation reasonably available.

This is why the model set out in the preceding pages is deliberately verifiable and correctable. The formulas, weights and thresholds do not ask for ideological agreement. They ask for data.

If the data shows that a measure does not work, it will have to be changed. If it shows that an operator has genuinely reduced risk, that operator's burden will have to fall. If it shows that a social benefit produces no results, that benefit will have to be rethought.

That is precisely the political point: turning the relationship between economic power and social responsibility from a general assertion into a measurable mechanism.

Giammario Battaglia

TECHNICAL APPENDICES

APPENDIX A — Summary sheet of the DSS model

Object of assessment: the economic or technological function, not the company considered in the abstract.

Proposed dimensions: S — scale; I — interconnection; N — non-substitutability; D — control of data or infrastructure that is hard to replicate; P — propagation; L — transition pressure on tasks and skills.

Logical sequence: identification of the function → data collection → gross score → recognition of mitigation → residual score → risk class → contribution, if any.

Fundamental rule: the contribution must not depend on economic success as such, but on the residual risk associated with the function and on the verified social need.

Methodological status: the weights, thresholds, classes and normalisation curves proposed in this volume must be empirically validated before any regulatory application.

APPENDIX B — DSS Data Return: minimum structure

- Identification of the operator and of the function assessed.
- Geographic perimeter and markets concerned.
- Revenues relevant to the function.
- Number and type of users and business users.
- Sectoral distribution of users.
- Dependencies on critical services or infrastructure.
- Downstream dependencies generated by the function.
- Actual switching times for representative categories.
- Actual switching costs.
- Technical and operational portability.
- Interoperability and standards used.
- Equivalent alternatives available.
- Concentration and infrastructural bottlenecks.
- Significant incidents and duration of their effects.
- Fallback mechanisms and operational continuity.
- Mitigation plan and related investment.
- Indicators of task and skills transformation, where relevant.
- Evidence needed for audit and for the reproducibility of the data.

APPENDIX C — INDS: summary structure

The Social Dividend Need Index ranks the intensity of the need for transition without becoming a general score for the person.

Illustrative formula proposed in this volume:

$$\text{INDS} = 0.30V + 0.20E + 0.20T + 0.15C + 0.15R$$

V — income vulnerability.

E — exposure of the occupation to transformation.

T — transferability of skills.

C — concrete capacity to adapt.

R — territorial vulnerability.

The weights are illustrative. Before any regulatory use they must undergo retrospective and prospective validation, discrimination testing, sensitivity analysis and periodic review.

APPENDIX D — Financial stress test

The scenarios of 10, 25 and 40 billion euro do not represent a forecast of DSS revenue. They are three levels of hypothetical need used to test whether the mechanism holds.

For each scenario, at least the following must be observed:

- gross funding need of the benefits;
- sustainable contribution after the caps;
- distribution of the burden across sectors;
- concentration of the burden on individual operators;
- reserves available;
- funding gap;
- effects of mitigation on aggregate risk;
- sensitivity to recession, falling revenues and rising beneficiary numbers.

Robustness condition: the model must be able to state explicitly that a scenario is not fundable in the given configuration, rather than artificially altering scores or coefficients to reach a revenue target.

Consolidated prudential formula:

$$F_c = F_{\text{commit},c} + \lambda_c \times \sum_i EL_i^{\text{net}}$$

$$C_i^* = F_{\text{direct},i} + F_{\text{common},c} \times w_i + \lambda_c \times EL_i^{\text{net}}$$

$F_{\text{direct},i}$ funds already-committed costs directly attributable to the individual operator; $F_{\text{common},c}$ is the common share of already-committed costs and is apportioned using w_i ; $\lambda_c \times EL_i^{\text{net}}$ funds the share of net expected social loss that the legislature decides to mutualise. The caps are applied afterwards, and any funding gap remains explicit.

APPENDIX E — Regulatory and institutional matrix of the DSS

This matrix is not a set of draft legislative articles. It indicates which decisions have to be taken, at what level they should be fixed, and which institutional body should take

responsibility for them. The definitive legal form depends on the legal basis and on the text that emerge from the European procedure.

Matter	Level at which to fix it	Principal responsibility	responsi-	Why
Objectives, perimeter, definitions	Legislative act	EU legislature	on a competent proposal	Essential elements affecting rights and obligations
Essential designation criteria	Legislative act	EU legislature		They limit the authority's discretion
Weights, normalisation, benchmarks	Updatable technical annexes/acts	Commission/authority	under delegation	Must be capable of recalibration on data
DSS Data Return — fields and formats	Implementing act or technical standard	Commission/competent authority		Uniform application and comparability are needed
Individual designation	Administrative decision	Competent authority		Requires investigation, adversarial process, reasons
Mitigation and verification	Reviewable technical decision	Competent authority / verifier		The contribution must be able to fall
Eligible benefits	Legislative act	Legislature		They determine the ring-fenced destination
Access criteria and INDS	Act plus technical annexes	Legislature and implementing authority		Balance between rights and calibration
Fund management	Legislative act plus governance	Dedicated body with separate accounts		Segregation and reporting
Independent control	Legislative act	Audit and evaluation bodies		Prevents conflicts of interest
Appeals	Legislative act	Courts under the applicable order		Effective judicial protection
Review and sunset	Legislative act	Legislature		The system must be able to shrink or lapse

The value of the matrix is to show that the DSS does not require a single “author of the law”. It requires a chain of responsibility: initiative and investigation, political decision on the essential elements, technical implementation, individual decisions, financial management and independent control.

APPENDIX F — Integrated falsifiability and control checklist

- ☐ Does the model clearly distinguish observed data, estimates and simulations?
- ☐ Does the score change when the facts change?
- ☐ Does verified mitigation actually reduce residual risk?
- ☐ Does the funding need derive from the benefits rather than from a revenue target?

THE PRICE OF POWER

- ☐ Do the caps make any funding gap visible?
- ☐ Are decisions reproducible using the same data and the same methodology?
- ☐ Can the operator contest the data, the normalisation and the perimeter?
- ☐ Does the INDS identify a measurable operational need without unnecessarily profiling the person?
- ☐ Do the benefits produce verifiable results?
- ☐ Can the system shrink when risk and need decline?
- ☐ Has the pass-through of the contribution been estimated for the principal customer segments?
- ☐ Have low, central and high incidence scenarios been published?
- ☐ Do the DSS weights derive from a documented elicitation protocol?
- ☐ Has multicollinearity between dimensions and indicators been checked?
- ☐ Have arithmetic and geometric aggregation of resilience been compared?
- ☐ Has a Monte Carlo uncertainty analysis been carried out?
- ☐ Have global sensitivity indices been calculated, including Sobol where appropriate?
- ☐ Has the probability of class membership under alternative specifications been published?
- ☐ Does L distinguish potential exposure, adoption, substitution/complementarity and reallocation distance?
- ☐ Does the validation of L use outcomes not already embedded in the index?
- ☐ Is every DSS benefit checked against the EGF, the JTF, the ESF+ and national instruments to avoid double funding?
- ☐ Is substantive additionality relative to ordinary public expenditure documented?
- ☐ Has a legal-basis opinion on Articles 114 and 175(3) TFEU been produced?
- ☐ Is the nexus between contributors and expenditure close enough to sustain the prudential analogy with the SRF?
- ☐ Have the expenditure categories been subjected to State aid screening?
- ☐ Are the designation and contribution criteria neutral as to the operator's nationality?
- ☐ Is there at least one European/non-US case subjected to the same DSS criteria?
- ☐ Is the perimeter of the ASML case defined as a function and not as a company in the abstract?
- ☐ Are export controls treated as evidence of strategic significance and not as automatic proof of systemic importance?
- ☐ Has a distributional audit of outcomes by operator nationality/origin been carried out?
- ☐ Is an origin-neutrality statement present in the simulated decisions?
- ☐ Is the WTO/GATS assessment specific to sector and mode of supply?
- ☐ Is a Trade Impact Assessment provided for before the contributory phase?

- ☐ Does the text avoid describing the DSS as a digital tax or a tax on Big Tech?
- ☐ Does the expected loss notation clearly distinguish probability, severity, exposure and the L dimension of work?
- ☐ Are shared expected social losses corrected for overlaps before sector aggregation?
- ☐ Are already-committed costs directly attributable to an operator separated from common costs?
- ☐ Do the stress tests distinguish gross F_{social} from F_{DSS} after additionality and causal attribution?
- ☐ Is the INDS subjected to the same calibration protocol and sensitivity analysis as the DSS score?
- ☐ Does the chapter on European law expressly integrate DORA and cross-refer to the legal-basis analysis in Chapter 28?

APPENDIX G — The twenty-six principles of the DSS

A progressive list of the principles formulated in this volume. The numbering now follows the actual order of reading and provides a summary map of the model's architecture.

- Systemic power is not presumed from wealth. It is demonstrated by dependency.
- Efficiency that creates dependency must be accompanied by the capacity to reduce, manage and reverse that dependency.
- No one contributes because they are wealthy as such. Those who contribute, within sustainable limits, are those performing a systemic function — for the share of residual risk that function continues to generate.
- Economic capacity limits the contribution; residual risk justifies it and determines how it is apportioned.
- The DSS designates systemic functions and attributes the resulting obligations to the operators that perform them.
- The DSS must make it economically worthwhile to reduce risk before mutualising it.
- Effective competition is measured not only by the number of doors in, but by the cost of the doors out.
- The funding need is a consequence of the benefits; sustainability is a constraint; revenue is not an objective in its own right.
- The Fund must be financially sufficient but institutionally interested in its own reduction when the risk justifying the contribution declines.
- A model of systemic responsibility is credible only if it can produce different results for different companies, and different results for the same company when the facts change.
- Where a European obligation already reduces a component of risk, the DSS must recognise it as mitigation; where reliable data already exists, the Data Return must reuse it so far as possible.

THE PRICE OF POWER

- The DSS score classifies systemic importance; expected social loss monetises cost scenarios. Neither replaces the other.
- No euro enters F_c without an explicit path linking it to an incremental cost, a counterfactual and an attributable share; no prospective euro is demanded without an expected social loss.
- The operator's legal liability does not prove the economic incidence of the contribution.
- No weight, no threshold and no aggregation formula may be shielded from empirical testing.
- The DSS measures the transformation of work at the level of tasks and skills, distinguishes exposure, adoption and displacement, and does not treat the technical ability to perform a task as proof that the task will be eliminated.
- The DSS must not fund what another instrument already funds, and must not displace ordinary public expenditure.
- The legal basis is not the proposal's final technical footnote. It is part of the design.
- The operator's nationality does not enter the score.
- The DSS must be origin-neutral, but not naive about trade effects.
- No contribution without a verifiable chain linking function, data, score, mitigation, residual risk and a reasoned decision.
- No regulatory algorithm may be harder to contest than the decision it produces.
- The DSS selects for need, not for worth.
- Where a person lives may increase the cost of transition, but policy must use that information to widen possibilities, not to entrench inequality.
- A right is effective only if the procedure for exercising it is simpler than the problem it is meant to solve.
- The DSS collects because a residual risk and a verified need exist; it does not create risk and need because it wants to collect.

THE AUTHOR

Giammario Battaglia is a jurist and the promoter of initiatives devoted to mediation, economic transformation and the relationship between technological innovation and social responsibility. He is the originator and co-founder of the Human Economic Forum and developed the concept of the Systemic Social Dividend as a proposal linking systemic power, risk prevention and the financing of transition.

In this volume he takes the DSS from the plane of principle to that of measurement: systemic functions, residual risk, financial stress tests, beneficiaries, governance and a possible legislative architecture.

Final note

The manuscript was constructed by distinguishing institutional sources, publicly available operator data, simulations and methodological proposals. The company cases serve an analytical function and do not constitute findings of liability. The DSS and INDS formulas are proposals to be piloted and validated.

