

The Stunted O

Why sustainability reporting finds Risk two-to-three times as often as Opportunity, and the complete method for closing the gap: Ground, SPARK, PROVE, and Look Outward.

A practitioner's guide to structured Opportunity discovery and validation, built on original research across seventy European companies and the topic taxonomies that make the method repeatable.

The evidence

SPARK · discovery

PROVE · validation

The taxonomies

The business case

ABOUT THIS GUIDE

Sustainability as ambition, not defence.

Every company reporting under ESRS must identify, assess and disclose the same three kinds of material item: Impacts, Risks, and Opportunities. I, R and O, given equal billing in the architecture. Companies present them in different forms, but look closely across a real cross-section of European companies and one of the three is doing almost all the work.

It sets out the evidence for the imbalance, measured item by item rather than asserted, then the method Nemetan uses to correct it: how a company moves from a Risk-heavy disclosure to a structured, evidenced search for the Opportunity that disclosure has been quietly pointing at all along.

THE STANDARD BEHIND EVERY CLAIM

Every figure here is either drawn from a named public source, among them EFRAG, BCG, KPMG and TCFD, or from Nemetan's own comparative coding study of forty mandatory ESRS reporters and thirty voluntary VSME reporters, run under a single protocol applied identically to both sides of the ledger, and AI-assisted: human protocol, decision rules and adjudication, AI extraction and tagging. The definitions of Risk and Opportunity were fixed in writing before any company was coded and never changed; scope and edge-case handling were revised during coding, with every change dated in the decision log. The figures are directional and indicative, not population-level estimates; the full dataset is available on request.

Published August 2026 · Edition 1. This is an exploratory paper; a conclusive follow-up, grounded in a year of real client engagements, follows in 2027. · www.nemetan.com · info@nemetan.com

CONTENTS

What's inside

01	The Opportunity Gap	5
02	Why It Happens: The Missing Process	8
03	The Method at a Glance	13
04	Grounding: Start From What Has Worked	14
05	SPARK: Structured Opportunity Discovery	16
06	The Taxonomies: What, At Which Stage, and Why	18
07	PROVE: Structured Opportunity Validation	22
08	Look Outward Before You Build	24
09	Why It Shows Up on the Balance Sheet	25
10	The Evidence Base, and How to Read It	27
11	Where to Start & About the Author	30

WITH THANKS TO

The shoulders this method stands on

Nemetan's contribution is the structure: the evidence library, the topic-to-technique matching, and the discipline that keeps discovery and validation separate. The individual thinking tools that structure draws on are the work of others, and we name them here with gratitude.

TOOL OR FRAMEWORK	WHOSE WORK IT BUILDS ON
Systematic Inventive Thinking (SIT)	Jacob Goldenberg, Roni Horowitz & Amnon Levav, developed from Genrich Altshuller's TRIZ
The Business Model Navigator	Oliver Gassmann, Karolin Frankenberger & Michaela Csik, University of St. Gallen
The 9R circular-economy framework	José Potting, Marko Hekkert, Ernst Worrell & Aldert Hanemaaijer (PBL Netherlands Environmental Assessment Agency)
The AR3T mitigation hierarchy (nature & water)	The IFC Performance Standard 6 mitigation hierarchy, as carried into the Science Based Targets Network and TNFD action frameworks
The EPA pollution-prevention hierarchy (source reduction as the preferred approach)	The U.S. Environmental Protection Agency (Pollution Prevention Act, 1990)
The carbon action hierarchy (avoid, reduce, substitute, then offset residual emissions)	A Nemetan synthesis of established greenhouse-gas mitigation-hierarchy practice, not a single standardised taxonomy
Discovery-Driven Planning & the reverse income statement	Rita Gunther McGrath & Ian C. MacMillan
The pre-mortem	Gary Klein
Open innovation	Henry Chesbrough
Creating Shared Value	Michael E. Porter & Mark R. Kramer
Affinity mapping (the KJ method)	Jiro Kawakita

Full bibliographic references appear in Chapter 10. Any errors of attribution are ours; corrections are welcome at info@nemetan.com.

The Opportunity Gap

Doing less harm used to be enough. Trim the emissions, tidy the supply chain, write it up, move on. That's ending.

The people who buy from you have started to notice, and they're quietly moving their money. NYU Stern has tracked it for years: products sold on real sustainability credentials now take more than a quarter of the US consumer market and pull far more than their share of its growth. Capital is doing the same, more slowly. When NYU Stern and Rockefeller Asset Management pooled more than a thousand studies, they found that companies treating sustainability as actual strategy tend to perform better and borrow cheaper. So the ones still doing the minimum are paying more for their money, whether they've noticed yet or not.

None of that stays in the marketing department. It shows up in demand and in what capital costs you, and before long it reaches the board. A director's job is to see and manage what could damage the value of the company. There's a legal opinion, written by the senior barrister Noel Hutley SC in 2016 and quoted in boardrooms ever since, that a director who ignores a foreseeable, financially material climate or sustainability risk may be in breach of that duty. Set that against everything above, and a sustainability report heavy on risk and thin on opportunity becomes exactly the kind of blind spot a board is now expected to have looked at, and to be able to show it did.

We coded it, item by item. Forty large European companies under mandatory ESRS, and thirty mid-size companies reporting voluntarily under VSME. In the large-cap sample, Risks outnumber Opportunities 2.32 to 1. In the voluntary sample, 2.93 to 1.

In more than a third of topic assessments, Opportunity does not appear at all, even where the same company lists several Risks on the very same subject.

More disclosure would not close this. ESRS already decides which opportunities reach the page, and it decides in one direction: a company weighs whether the opportunity is one it is actually pursuing before it reports it, and nothing of the kind applies to Risk. So part of the gap is written into the standard. The part that is not written in is the more interesting half. That instruction applies identically to all forty companies in the mandatory sample, and their ratios still range from roughly one to one up to sixteen to one. It does not apply at all to the thirty voluntary reporters, where the gap is wider still. A rule that treats everyone the same cannot explain why they come out so differently.

Which points at the process rather than the paperwork. A double materiality assessment is built to assess: to take an item that is already on the table and score it for likelihood and magnitude. It is good at that. What it has no machinery for is putting the item on the table in the first place. Risk does not need it to, because that identification happened years ago, in a register somebody already owns and maintains. Opportunity has no equivalent upstream, so an exercise that only assesses will only ever weigh what somebody had already thought of.

WHAT WE MEAN BY OPPORTUNITY

*A genuine Opportunity is one where the commercial benefit and the environmental or social benefit come from the **same decision**. It is not a cost that happens to generate goodwill, and not a risk avoided dressed up as a gain created.*

That definition does real work throughout this guide. Real efficiency that lowers cost because it uses less. New revenue unlocked by doing right by workers or the planet. Market access won by moving first. Compliance programmes and defensive spend, however proactively framed, do not qualify. Avoiding a loss is not the same as creating a gain.

2.32:1

Risk-to-Opportunity,
large-cap ESRS (40
companies)

2.93:1

Risk-to-Opportunity,
voluntary VSME (30
companies)

37%

of topic assessments
with Risk but zero
Opportunity

1.55:1

EFRAG's own buried
climate ratio, 905
statements

The signal was already scattered across the published record. Buried in a section of EFRAG's 2026 State of Play, the most comprehensive study of European sustainability disclosure ever conducted, 905 assured statements deep, is a single line: companies disclose, on average, 3.1 climate Risks against 2.0 climate Opportunities. That is a 1.55 to 1 skew, in the one topic where Opportunity reporting is most mature. EFRAG never asked the question directly. The answer was sitting there anyway. Others have measured it since. KPMG went through 270 first-wave statements and found Opportunity at 14% of disclosed items against 26% for Risk, with 11% of companies identifying no Opportunity at all. We found 10%. What nobody has measured is what happens when the legal requirement disappears, which is what the second half of this study is for.

Look topic by topic and it sharpens. One large Nordic retailer in the sample makes the point better than any summary: it tags its material topics cleanly, row by row, and every row is an Impact or a Risk. Not one is an Opportunity. The process worked exactly as designed. It was simply never asked to look for the upside. At the other end sits a Southern European electricity transmission operator, disclosing more Opportunities than Risks (14 to 13), tied to regulated revenue mechanisms it can measure. Even there the discipline slips on one side only: its own narrative says the assessment identified fifteen material Opportunities, and the table itemises fourteen. No large listed company announces fifteen material Risks and then tabulates fourteen. That company is not more virtuous. It sits in a business model where finding Opportunity is unavoidable. Wherever the business model forces the question, the balance flips. Everywhere else, nobody built the process to ask.

Opportunity as a share of disclosed Risk + Opportunity items

Large-cap ESRS sample · 40 companies · bars sorted by Risk-to-Opportunity ratio ■ Opportunity ■ Risk




Figure 1. Risk versus Opportunity by sector, large-cap ESRS sample (n = 40). Green shows Opportunity's share of each sector's disclosed items; it clears the parity line only in electricity and real estate.

Opportunity is not silent. It is stunted: present and technically disclosed, yet denied the structural rigour Risk receives by default. That is the Stunted O.

Why It Happens: The Missing Process

The imbalance has two roots. One is built into the architecture of the frameworks themselves. The other is what a materiality assessment is built to do, and who is in the room while it does it.

The frameworks inherit a risk-first shape

The investor-focused disclosure frameworks descend, directly or by imitation, from a lineage built to manage financial risk first. TCFD was created by the Financial Stability Board to treat climate as a source of systemic financial risk. (It completed its work and disbanded in October 2023, its architecture absorbed into the ISSB's IFRS S2, so it endures as the foundation of investor-focused climate disclosure rather than an active body.) At the level of vocabulary it is even-handed, with six risk subcategories and five named Opportunity categories, published together in 2017. The imbalance sits one level up, in the architecture: of its four pillars, only one is named after a single side of the pair. There is a *Risk Management* pillar. There is no Opportunity Management pillar. Opportunity does appear across the other three pillars, all of which name risks and opportunities together, but it never gets a management process of its own. Risk Management asks a company to describe how climate risks are identified, assessed and managed, and how those processes are integrated into overall risk management. Nothing on the Opportunity side carries a comparable demand.

FRAMEWORK	RISK TAXONOMY	OPPORTUNITY TAXONOMY	WHERE OPPORTUNITY LIVES
TCFD (2017)	Physical + transition (6 types)	Five named categories: resource efficiency, energy source, products & services, markets, resilience	Inside Strategy; no Opportunity pillar
TNFD (2023)	Physical, transition, <i>systemic</i> (expanded)	Seven named nature-related categories: resource efficiency, capital flow & financing, markets, products & services, reputational capital, sustainable use of natural resources, and ecosystem protection, restoration & regeneration	Named within the management pillar's own disclosures, but not in its title, and left out of the disclosure on integration into overall risk management
ESRS (in force FY2024–26)	No single cross-topic taxonomy; a topic list (AR 16), plus topic-level risk categories in E1, E2, E3 and E5	No cross-topic taxonomy; opportunity categories in E2, E3, E4 and E5, but only two considerations in E1	Named in the process pillar itself (IRO), though integration into overall management is qualified “where applicable”, and a separate section filters which opportunities may be reported at all

TNFD kept the asymmetry in the name, calling its third pillar *Risk and Impact Management* rather than Risk and Opportunity Management. It did not keep it in the substance. Inside that same pillar, TNFD asks organisations to describe how they identify, assess, prioritise and manage nature-related dependencies, impacts, risks *and opportunities*, in direct operations and across the value chain. Given the chance to correct what it inherited, TNFD corrected most of it and left the risk-led name in place. One trace does survive, in an interesting spot: the disclosure asking how any of this connects to the machinery a company already runs narrows back to risk alone. Opportunity gets identified, assessed, prioritised and managed. It does not get integrated. TNFD also sets out seven named nature-related Opportunity categories, so a usable Opportunity taxonomy can plainly be built; that it published a discussion paper in November 2025 grouping those categories a second way, by business function, without reconciling the two, says something about which side of the pair is still settling.

ESRS goes furthest. The dedicated disclosure area in ESRS 2 is called Impact, Risk, and Opportunity management, all three together, and of the three frameworks here it is the only one to put the word into the name of the process area itself. Even there the symmetry stops just short. The same disclosure requirement asks how the process for impacts and risks is integrated into the overall risk management process, then asks how the process for opportunities is integrated into the overall management process, *where applicable*. Only one of the two carries the qualifier.

And one provision decides what reaches the page at all. ESRS 1 has a short section headed Reporting on opportunities, telling a company that the disclosure should be descriptive, and

that in judging what to disclose it shall consider whether the opportunity is currently being pursued and built into its general strategy, as opposed to being a general opportunity for the sector. There is no section on reporting risks, in this version of the standard or the revised one, and no equivalent restriction on disclosing a risk. A risk that has been found and not yet acted on is reportable. An opportunity in the same state is filtered.

The assessment is built to weigh, not to find

Underneath the architecture sits a distinction the whole exercise runs past. Identifying an opportunity and assessing the materiality of an opportunity are not the same cognitive task. A double materiality assessment is built for the second. It takes what is in front of it and scores it, on likelihood combined with magnitude, and it does that carefully and defensibly. Nothing in it is designed to put anything in front of it.

Risk survives that gap because it does not depend on the assessment to find it. The identification work already happened, in an enterprise risk register somebody owns, maintains and reviews, often for years before the sustainability statement existed. Opportunity has no equivalent upstream source. So an exercise that only assesses will only ever weigh what somebody had already thought of, and for one of the two categories that is almost everything, while for the other it is close to nothing.

And the room reflects it

Risk and Opportunity may appear side by side in an IRO assessment, but they are not found in quite the same way. A strong search for Opportunity needs sustainability and financial insight, but also the perspectives of people who understand customers, products, markets and innovation.

Finding material Risks calls for risk, finance and sustainability expertise. Look at who is often brought into a double materiality assessment: sustainability, finance, risk and IT, with the compliance consultants engaged to run it. Every one of those functions is excellent at what it was built to do: identify exposure, quantify downside, document a defensible position. None of them is designed as an opportunity-discovery discipline. Some companies do bring product, commercial and business-unit people to the table, and that variation is worth attending to rather than assuming away, because it is the closest thing this evidence has to a control. Ask a room assembled for the assessing job to do the finding job, with the same process, the same axes and the same instincts, and it is no surprise Risk wins better than two to one. It would be a surprise if it did not.

Our own data confirms it from the other direction. Among the thirty voluntary reporters (companies reporting under a voluntary standard like VSME, rather than a legal mandate), the single variable that predicted whether Opportunity showed up was not size or sector. It was who ran the process. That points to team and process as the leading explanation the evidence supports, though not the only one that could be at work.

The tools split by side

Look closer at that room and a deeper reason appears, one built into the instruments themselves. When the same company handles Risk, it reaches for instruments whose job is to generate possibilities: scenario analysis above all, and the wider habit of imagining in detail how a thing could go wrong. Risk gets a real search. For Opportunity, the toolkit holds ladders, the mitigation hierarchies that rank a fixed set of moves in priority order. A ladder sorts what already exists, and it was never built to invent. So the room can order its harm-reduction options neatly and still produce nothing new, because nothing in its hands has generation as its job.

Seen that way, the imbalance in the disclosure is the imbalance in the toolkit, one level down. It is a leading, structural reason rather than the only one, sitting alongside commercial confidentiality and the plain difficulty of quantifying a gain that has not launched yet. It is also the one a method can fix, because it is a missing capability rather than a failure of will. The chapters that follow are, in effect, the tool that was missing.


That missing discipline now has a name. **Sustainability Opportunity Discovery** is the systematic search for candidate Sustainability Opportunities, conducted before and separately from any assessment of their materiality, and producing an auditable record of what was searched, what was found and what was set aside. Its definition, its boundaries against adjacent practice, and the requirements any method must meet are published separately and cited at the end of this guide. Everything from here is one implementation of those requirements.

Who runs the process decides whether Opportunity appears

Voluntary VSME sample · 30 companies · item counts by facilitator type

■ Risk

■ Opportunity



Only one group reached parity, the structured, human-led external process. Automated platforms did worse than the mandatory-reporter average.

Figure 2. Risk versus Opportunity by facilitator type, voluntary VSME sample (n = 30). Externally-facilitated structured assessments were the only group to reach parity; automated platforms produced Risk at seven times Opportunity, and companies with no formal process disclosed nineteen Risks and not a single Opportunity.

THE MOST ACTIONABLE FINDING IN THE RESEARCH

The skew behaves like a **facilitator-quality floor**, the outcome a mature but risk-adjacent process reliably produces by default, which a genuinely Opportunity-probing engagement, human-led and structurally designed to treat the two categories as separate exercises, can meaningfully improve on at any company size. Structure changes the outcome, and that is the entire basis for the method that follows.

The Method at a Glance

If facilitator quality is the strongest predictor of whether Opportunity gets developed, ahead of company size, resources or sector, then the fix is a repeatable method any organisation can apply, rather than a plea for more disclosure or another framework.

The method runs in four movements. It starts before ideation, with a step most processes skip, and it does not end when a shortlist appears. It ends when a validated Opportunity has been tested against outside evidence and matched to whoever has already solved a piece of it.




Figure 3. The Nemetan method, end to end. Discovery (SPARK) and validation (PROVE) are deliberately separate disciplines, because a promising idea and a fundable one are not the same thing.

Two disciplines carry the weight. **SPARK** is the discovery method: five stages that take a company from its own disclosed data to a short, defensible list of candidate Opportunities, using an ideation tool matched to the specific topic. **PROVE** is the validation method: five moves that put each candidate through a real test before anyone gets attached to it. Between and beneath them sit the **taxonomies**, the topic structure and the discovery-lens library that decide which technique is used where. Chapter 6 sets those out in full, because they are what makes the method repeatable rather than personal to whoever happens to run it.

START HERE

Two free tools place you before you read another page. The **Gap Check** scores your most recent disclosure on the Risk-to-Opportunity balance measured in this guide and benchmarks it against everyone who has taken it, in about ninety seconds. The **Maturity Matrix** gives you two scores rather than one, how far your general sustainability practice has travelled and, separately, how far your *opportunity* practice has, with the gap between them being the finding. If you want something to circulate first, the **Executive Brief** carries the headline evidence in a form you can share across the organisation.

Discovery and validation are two different jobs, so the method runs them as two separate disciplines instead of one hopeful pass.

Grounding: Start From What Has Worked

Fixing the staffing gap only works if the team that walks in has a method. Ours starts before ideation, with a step most processes skip.

Grounding solves a specific problem. Ask almost any company for proof that a sustainability move actually paid off, in real numbers, and you hit a wall. The world is drowning in generic claims and desperately short of usable case studies: named companies, verifiable data, documented benefits you could take to a board. That scarcity is a large part of why Opportunity work stalls. Nobody wants to be first, and almost nobody can show they were even second.

So Nemetan built the evidence it could not find: a proprietary database of more than 2,000 sustainability Opportunity precedents, which we believe is the largest of its kind anywhere, each one tied to a real company, a real mechanism, and evidence you can check. Every case clears a strict bar before it counts: a named company, a quantified *achieved* result (past tense, not projected) on both the business and the environmental side, from the same move, recent enough to matter, and honest about its limitations.

2,000+

verified opportunity precedents in
the library

5

independent tests each case must
pass

3-level

tagging: theme → sector →
mechanism

Grounding puts it to work. Before a room generates a single idea, candidate concepts for the topic at hand are surfaced and matched, with an AI tool built for the job, to the closest real precedents in the database. The team starts from what has already worked somewhere else, rather than an empty whiteboard. A company beginning its search from a matched set of proven precedents is starting years ahead of one working from scratch, and the advantage grows the more the library is used. The same library returns later, at the Verify step of PROVE, to test a surviving concept against real precedent.

Why Grounding cannot be skipped

The load-bearing stage of this method is Apply, where Systematic Inventive Thinking generates the concept. SIT's templates are operators. Subtraction removes a component from an existing configuration; Task Unification gives a new job to an element already present. Each acts on something, and none can fire on an empty input.

That is the mechanical reason Grounding comes first. What SIT produces is bounded by what it is handed. Hand it a matched set of screened precedents and the templates work on configurations that have already succeeded somewhere. Hand it nothing and they work on whatever the room remembered.

The principle is not ours. Goldenberg, Mazursky and Solomon published it in 1999: a finite set of named templates outperforms the open brief.

A discovery process that begins from a blank page is therefore not Sustainability Opportunity Discovery. That is a stated requirement of the published definition, not a house preference. It is not a claim that Grounding produces better ideas than another precedent source, which we have not tested.

WHY STRUCTURE IS THE REAL PRODUCT

One verified case is a research task. Two thousand of them, tagged consistently by theme and sub-theme, by sector, and above all by the underlying *mechanism* that made each one work, become something you can query. A loose folder tells you only that something happened somewhere once. A tagged library tells you which mechanism applies to a challenge like yours, and where the precedent sits that proves it works.


This is also why a generic AI tool, on its own, cannot produce what this work needs. Ask a general model for sustainability case studies and it returns plausible, unverified, untagged prose. Call that **noise**: no evidence bar, no mechanism structure, no way to tell a proven-at-scale result from a hopeful one. A curated, graded, mechanism-tagged library turns the same territory into a grounded **signal** you can build a board-ready case on. The AI does the fetching and matching inside that structure. On its own it has no way to tell a proven result from a plausible-sounding one.

SPARK: Structured Opportunity Discovery

Grounding gives the room a starting point. SPARK gives it a method: five deliberate stages that separate naming a possibility from building a solution, and match each possibility to the technique built for its topic.

SPARK runs in five stages

From a company's own disclosed data to a short, traceable shortlist



S: Spot the gap

Benchmark a company's own disclosed data against real external standards, rather than internal impression, to find where a genuine, evidenced gap actually sits, an opportunity scan across the company's own material topics.

P: Pose the question

Hold a strict line between naming a possibility and building a solution. Every structured room converges on an answer too early unless something deliberately stops it; this stage is that stop.

A: Apply the tool · THE LOAD-BEARING STAGE

Match each possibility to the right kind of tool. A ladder points to where the opportunity is likely to sit, an inventive technique generates the concept, and a business-model pattern shapes how it captures value. This stage lives or dies on getting the match right.

R: Refine

Filter every candidate for logical consistency, commercial viability, and real precedent. This is the opportunity triage that turns a wide field of raw ideas into a short, defensible list.

K: Keep

Produce a final shortlist where every surviving concept is fully traceable back to the exact data and evidence that produced it.

WHY "APPLY" CARRIES THE MOST WEIGHT

Not every ESRS topic offers the same starting point. The environmental topics all arrive with a mature ladder of some kind. Social topics have a value-creation search framework instead, and governance has neither. Running one generic ideation exercise across all ten topics is part of how the gap gets built in the first place. Matching the tool to the topic is the thing a single undifferentiated brainstorm cannot do, and it depends on the discovery-lens taxonomy behind it.

When there is no assessment to start from: Affinity Mapping

SPARK assumes a starting point, a company that already has a materiality assessment. Our research shows that assumption fails for a large share of companies the moment the mandate disappears: two of every five voluntary reporters (like VSME) disclosed no itemised content at all, yet several had genuinely run a process whose output simply never got structured. Affinity Mapping closes that gap. It takes whatever a company already has, from internal trackers and committee notes to survey responses and supplier correspondence, and clusters it into coherent themes, each mapped to its nearest ESRS topic, so discovery can begin without a formal filing already in hand.

The Taxonomies: What, At Which Stage, and Why

Three taxonomies decide which technique the method uses, where, and why. Between them they keep Opportunity discovery from depending on whoever happens to run it.

1 · The topic taxonomy: the scanning layer

Everything is organised around the ten ESRS topics. This is the index the whole method runs against: it is where *Spot the gap* benchmarks, what the evidence library is tagged to, and how a candidate is filed. It is deliberately the same structure a company's own disclosure already uses, so nothing has to be translated.

ENVIRONMENT	SOCIAL	GOVERNANCE
E1 · Climate change & energy	S1 · Own workforce	G1 · Business conduct (broader governance sits in ESRS 2)
E2 · Pollution	S2 · Workers in the value chain	
E3 · Water & marine resources	S3 · Affected communities	
E4 · Biodiversity & ecosystems	S4 · Consumers & end-users	
E5 · Resource use & circular economy		

2 · The discovery-lens taxonomy: the engine of “Apply the tool”

This is the proprietary core. It is recorded three levels deep, and SPARK's third stage draws on it to match a technique to a topic. Read each entry below as those three levels: the **Method** (the column it sits in), the **Framework** (in bold), and its **Components** (the individual moves beneath it). The original source is credited in grey. The left column holds the where-to-look search lenses; the right holds the how-to-build methods that generate the concept and shape how it captures value.

METHOD · SEARCH LENSES (WHERE TO LOOK)

9R Framework (R0–R9), circular economy (E5)
R0 Refuse · R1 Rethink · R2 Reduce · R3 Reuse · R4 Repair · R5 Refurbish · R6 Remanufacture · R7 Repurpose · R8 Recycle · R9 Recover
Potting, Hekkert, Worrell & Hanemaaijer, PBL.
Called “9R” after its highest rung R9; because the strategies are numbered from R0, the ladder in fact lists ten.

Carbon action hierarchy, climate (E1)
Avoid · reduce · substitute · offset residual emissions
Nemetan synthesis of GHG mitigation-hierarchy practice

EPA pollution-prevention hierarchy, pollution (E2)
Source reduction · recycling · treatment · disposal
U.S. EPA, source reduction first

AR3T + TNFD opportunity categories, water & biodiversity (E3/E4)
Avoid · Reduce · Restore & Regenerate · Transform, plus TNFD's nature-related opportunity categories
Mitigation-hierarchy logic (IFC Performance Standard 6), aligned with the SBTN and TNFD action frameworks

Creating Shared Value, social topics (S1–S4)
Reconceive products & markets · redefine value-chain productivity · enable local clusters
Porter & Kramer

METHOD · GENERATION & VALUE CAPTURE (HOW TO BUILD)

Systematic Inventive Thinking (SIT), generation
Subtraction · division · multiplication · task unification · attribute dependency
Goldenberg, Horowitz & Levav, from Altshuller's TRIZ

Business Model Navigator, value capture
e.g. pay-per-use · take-back / circular recovery · by-product & asset valorisation · certified/ trademarked differentiation (from 55 patterns)
Gassmann, Frankenberger & Csik, University of St. Gallen

The search lens tells the room where to look; the inventive engine and the business-model patterns do the building. Circular and resource questions start from the 9R hierarchy, climate from the carbon mitigation hierarchy, biodiversity from AR3T, and the social topics from Creating Shared Value's three routes. Governance has no lens of its own, so there the room goes straight to SIT and the business-model patterns, applied to the governance arrangement itself.

WHY A CONSTRAINT IS THE ENGINE, NOT THE BRAKE

These are libraries, not inspiration. The inventive templates and the business-model patterns are each a finite, named set of moves drawn from what has already worked, the same way the evidence library is a set of proven outcomes. Generation here is recombination across those libraries applied to your own situation, never a blank page. That is what keeps it systematic and teachable: you can run every template and every pattern and know nothing was missed, and an idea can be named by the move that produced it rather than left to float. The constraint is what makes the search productive.

3 • The framework taxonomies: diagnostic context

The third taxonomy is not one we apply so much as one we read the landscape through. TCFD's five named Opportunity categories, TNFD's own nature-related Opportunity categories, and ESRS's topic-based approach explain *why* the gap exists: TCFD and TNFD both show an Opportunity taxonomy can be built, while ESRS gives Opportunity equal billing but no cross-topic discovery taxonomy of its own. TCFD's categories seed the climate prompt set at *Pose the question*. They are covered in Chapter 2; here they simply mark where the method borrows structure and where it has to supply its own.

Which taxonomy fires at which stage

STAGE	TAXONOMY IN PLAY	WHY HERE
Grounding	Evidence library, tagged by theme, sector and mechanism: the discovery-lens taxonomy, applied automatically as a first-pass classification	To match a topic to proven precedents by mechanism before any idea is generated
SPARK • Spot	ESRS topic taxonomy + external benchmarks (TPI, CDP, WBCSD, EFRAG)	To benchmark disclosed data against the right external standard
SPARK • Pose	Possibility prompts: TCFD opportunity categories (climate) and Creating Shared Value routes (social)	To frame “how might we” questions without reaching for a solution yet
SPARK • Apply	Discovery-lens taxonomy (lenses, SIT, business-model patterns)	To match the ideation technique to the specific topic
SPARK • Refine / Keep	Topic + mechanism tags	To triage candidates and keep every survivor traceable to evidence
PROVE • Verify	Evidence grade (A / B) on each supporting case	To separate proven-at-scale results from illustrative ones when testing a claim

The discovery-lens taxonomy therefore fires twice. Because the taxonomies and the precedent library sit together in one structured database, the tool applies them automatically: first to classify and retrieve matching precedents at Grounding, then to select the ideation technique at Apply.

Which technique fires against which topic

Applied inside SPARK Stage 3 ("Apply the tool"). Green is a mature ladder; teal is the social value-creation search lens.

	ESRS TOPIC	MATCHED DISCOVERY TECHNIQUE	STARTING POINT
E1	Climate change & energy	Carbon mitigation hierarchy	Mature ladder
E2	Pollution	EPA pollution-prevention	Mature ladder
E3	Water & marine resources	AR3T Action Framework	Mature ladder
E4	Biodiversity & ecosystems	AR3T Action Framework	Mature ladder
E5	Resource use & circular economy	9R Framework (R0–R9)	Mature ladder
S1	Own workforce	Creating Shared Value	Value-creation search lens
S2	Workers in the value chain	Creating Shared Value	Value-creation search lens
S3	Affected communities	Creating Shared Value	Value-creation search lens
S4	Consumers & end-users	Creating Shared Value	Value-creation search lens
G1	Business conduct	SIT · Business Model Navigator	No ladder, go to generation

Only the environmental topics come with a ladder. Social topics use a value-creation search lens; governance has neither, and goes straight to generation. That uneven coverage is itself part of the argument.

Figure 4. The topic-to-technique map applied inside SPARK Stage 3 ("Apply the tool"). The environmental topics come with a ladder, the social topics with a value-creation search lens, and governance with neither, so a method that runs one technique everywhere builds the gap in by design.

PROVE: Structured Opportunity Validation

Generating a shortlist of candidate Opportunities is the easy half. The harder, more valuable half is knowing which ones are credible enough to justify further investment or a real test. That is a different discipline, so SPARK is paired with a separate validation method: PROVE.

PROVE runs in five moves

It borrows openly from people who have done this for a living



P: Pre-mortem

Imagine the concept has already failed, and write down why, before anyone in the room has built enough attachment to defend it. The technique is Gary Klein's, and the discipline it enforces is honesty before investment.

R: Reverse the income statement

Rita McGrath and Ian MacMillan's discovery-driven discipline: work backward from a real target figure to the assumptions that would have to hold for it to be real.

O: Obligations

Name the assumptions the case depends on, out loud, and rank them on three axes: how essential each is to success, how uncertain it is, and how cheaply it can be tested. Test the ones that are essential, uncertain and cheap first.

V: Verify

Test the assumptions that matter against genuine external evidence, not internal conviction. This is where the evidence grade on each supporting case earns its place, and where the first look outward begins, scanning for anyone who has already solved a piece of it.

E: Evaluate

With kill criteria set in advance, reach a plain decision: go, adjust, or kill. Stopping a concept that does not hold up counts as the process doing its job.

WHY THE DECISION DISCIPLINE MATTERS

We would rather tell you honestly that something does not apply to your business than manufacture enthusiasm for a shortlist that needs it. Kill criteria set *before* the evidence arrives are exactly how a CFO already builds and defends a serious business case, which is why PROVE tends to land as a discipline already trusted rather than an unfamiliar one.

Kill criteria set before the evidence arrives are what let a team say go, adjust, or kill and mean it.

Look Outward Before You Build

There is one more discipline inside validation, and it is the one most companies skip. Once an Opportunity is validated as worth pursuing, the next question is not *how do we build this*. It is *who has already solved a piece of it*.

Open innovation, deliberately looking beyond your own walls for existing solutions, partners and technology, is one of the most powerful levers in sustainability strategy and one of the most underused. Henry Chesbrough named the discipline twenty years ago, and the sustainability field is only now catching up to it. In PROVE it runs in two passes: an early scan during Verify, which often reshapes which assumptions even matter, and a deeper partner search once a concept has survived the decision. Startups are often years ahead of large incumbents on exactly the kind of narrow technical problem a validated Opportunity tends to reduce to: a novel material process, a monitoring technology, a circular logistics model. Universities are sitting on adjacent research further along than most corporate innovation functions realise, for the simple reason that nobody has asked.

This extends PROVE's rigour rather than departing from it. A startup's solution can be genuinely well-evidenced and still be a poor fit for a particular company's scale, context or supply chain. Those are two separate, honest judgments, evidence quality and situational fit, and keeping them apart is what turns an exciting pitch into a defensible business case.

For a company that has just discovered an Opportunity its own materiality assessment was never built to find, looking outward for who has already solved part of it is the fastest, lowest-risk path from a validated idea to something real.

Why It Shows Up on the Balance Sheet

There is a version of this argument that stops at “good for the world.” That version, however true, will not move a CFO or a board on its own. Here is the version that will.

Value created through genuine shared-value Opportunity, not compliance spend dressed up as strategy, can affect the **value drivers** lenders and insurers weigh: cash flow, asset resilience, transition exposure, and evidence of management capability. Whether it moves the terms on offer depends on materiality, execution, and each institution's own methodology. A company that can point to real, evidenced Opportunity stands on firmer ground than one that can only point to Risk mitigated. One caveat carries the whole thing: an *identified* Opportunity does not move the number, only a *productised, launched* one does. Read these as value pathways rather than guaranteed effects, each dependent on execution, scale, and how credibly the result is evidenced. Discovery is the start of the work, not the end of it.

Insurance and underwriting

Insurers already price climate and transition exposure into premiums. A company actively creating value through better resource use, more resilient operations and de-risked dependencies is a measurably different underwriting proposition than one merely clearing a minimum.

Stranded assets

An asset strands when the market moves faster than the business model built around it. A company that spent its materiality process only cataloguing Risk, and never systematically searching for the Opportunity on the other side of the same transition, is far likelier to be caught holding one. The least-exposed companies are the ones that found the Opportunity the transition created, and moved first.

WHAT'S ACTUALLY MISSING

None of this financial logic is new; it is close to the working consensus in sustainable finance. What is missing is a repeatable method that actually produces the Opportunity, evidenced well enough to put in front of a board, rather than an aspiration. That is the gap this method is built to fill.

What counts as a genuine Opportunity, throughout, holds a hard line: the commercial benefit and the environmental or social benefit come from the *same* decision. Real efficiency that lowers cost because it uses less. New revenue unlocked by doing right by workers or the planet. Market access won by moving first. Compliance programmes, insurance, and defensive spend

do not qualify, however proactively dressed up. Avoiding a loss is not the same as creating a gain, and that one distinction is the whole argument.

The Evidence Base, and How to Read It

Every claim in this guide is either drawn from a named public source or from original research we ran ourselves: an item-by-item coding of forty large European ESRS statements and thirty mid-size VSME statements, under a single protocol applied identically to both sides of the ledger. The definitions of Risk and Opportunity were fixed in writing before any company was coded and never changed; scope and edge-case handling were revised during coding, with every change dated in the decision log.

The coding was AI-assisted. The protocol, the decision rules and the adjudication of edge cases were human; the extraction and item-level tagging were done with AI tooling working to that protocol, cross-checked against companies' own stated totals wherever they published one. This is the same approach EFRAG took in its 2025 review of 656 sustainability statements, and it is disclosed here for the same reason.

One limit shapes all of it. We counted what companies disclosed, which is not the same as what they found. ESRS filters which opportunities may be reported, having a company weigh whether an opportunity is currently being pursued and built into strategy, and applies nothing similar to risks. So the disclosed record understates what any company's assessment may have surfaced, by an amount we cannot measure. The imbalance could arise in the search, in what survived assessment, in what was carried into strategy, in what is being pursued, or in what was judged reportable. We observe the last of those.

We are upfront about what else it does and does not mean. The samples are real and independently verifiable, but neither was gathered through a statistically representative frame. Read the figures as directional and indicative, not as precise population-level estimates. The full dataset, which includes company-level tallies, an item-level register annotated by topic and evidentiary tier, a company-by-topic matrix, and a numbered coding-decision log, is available on request.

Is the headline an artefact of sector mix? No.

Because the sample is purposive, the natural question is whether the 2.32 : 1 headline is an artefact of which sectors were included. Recomputing under six exclusion scenarios shows it holding in a narrow band, from just over 2 : 1 up to 2.32 : 1, and climbing close to 3 : 1 without reaching it when the two genuinely opportunity-rich sectors are removed.

SCENARIO	N	RISK	OPP.	RATIO
Full sample	40	455	196	2.32 : 1
Core 6 stratified sectors only	30	322	157	2.05 : 1
Exclude Pharmaceuticals	38	421	192	2.19 : 1
Exclude Pharma + Food & Beverage	36	405	191	2.12 : 1
Exclude the 4 two-company probes	32	334	161	2.07 : 1
Exclude opportunity-rich pair (Elec. + C&RE)	31	371	127	2.92 : 1

The honest reading states the full-sample 2.32 : 1 as primary and this band as its robustness check. The near-balance in utilities and real estate is doing real work, and even removing both of those sectors leaves the ratio short of three to one.

What was checked

The coding was then re-read against source. Twenty-six of the forty large-cap statements were opened and checked page by page, and eighty-seven page references were corrected. One count moved: a chemicals group had disclosed four typed Opportunities where our register held two, and both missing items sat on the same page as four Risks we had captured correctly. That correction took the headline from 2.35 to 1 down to 2.32 to 1, narrowing the very gap this guide describes. Every item now carries a page reference into the company's own statement, on both sides of the ledger, and no page number has been estimated. No independent second coder re-coded either sample; that check can be run on request.

The regulatory moment gives it urgency

The EU's Omnibus package reshaped the rules in two steps. Directive (EU) 2026/470, in force since 18 March 2026, narrowed CSRD scope to companies exceeding both an average of 1,000 employees and €450 million in net turnover; separately, revised ESRS adopted by the European Commission on 3 July 2026 cut the mandatory data points by more than 60% and the total by more than 70%. That delegated act was still within the European Parliament and Council scrutiny period at the time of writing, and applies to financial years beginning on or after 1 January 2027. The first-wave companies reporting under the fuller original standard now form a bounded historical cohort, making this a data-rich moment unlikely to persist in its current form. At the same time, a large and growing population of smaller companies, many now outside mandatory scope, will still be asked for sustainability information by customers and lenders, and some will run a structured materiality exercise for the first time to answer that request, precisely the population where our findings show outcome quality depends most heavily on facilitator choice.

Selected sources

Research & frameworks: EFRAG, State of Play 2025 & 2026 (656 and 905 assured statements). · Boston Consulting Group, *Sustainability Reporting in Europe: From Compliance to Strategy*, 2026. · KPMG, *Understanding Double*

Materiality, 2025. · Deloitte, *A Practical Approach to Assess Financial Materiality*, 2024. · TCFD, 2017 Recommendations. · TNFD Recommendations & LEAP guidance. · CDP, climate risk & opportunity analysis. · Nemetan, original comparative coding study of 40 ESRS and 30 VSME reporters. Full methodology and dataset available on request.

Thinking tools credited in this guide: Goldenberg, Horowitz & Levav, Systematic Inventive Thinking (from Altshuller's TRIZ). · Goldenberg, Mazursky & Solomon, "Toward Identifying the Inventive Templates of New Products: A Channeled Ideation Approach," *Journal of Marketing Research* 36(2), 1999. · Goldenberg, Mazursky & Solomon, "The Fundamental Templates of Quality Ads," *Marketing Science* 18(3), 1999. · Gassmann, Frankenberger & Csik, *The Business Model Navigator*, University of St. Gallen, 2014. · Potting, Hekkert, Worrell & Hanemaaijer, *Circular Economy: Measuring Innovation in the Product Chain* (the 9R framework), PBL, 2017. · Science Based Targets Network, AR3T Action Framework. · U.S. EPA, Source Reduction / Pollution Prevention Act, 1990. · McGrath & MacMillan, "Discovery-Driven Planning," HBR 1995. · Gary Klein, "Performing a Project Premortem," HBR 2007. · Chesbrough, *Open Innovation*, 2003. · Porter & Kramer, "Creating Shared Value," HBR 2011. · Kawakita, the KJ (affinity) method.

Where to Start

The ratios in this guide are worth applying to your own materiality output, not just reading about. Two free tools do it in minutes.

- **The Gap Check** scores your most recent disclosure on the same Risk-versus-Opportunity balance measured here, and benchmarks it against every organisation that has taken it. About ninety seconds.
- **The Maturity Matrix** gives you two scores rather than one, how far your general sustainability practice has travelled and, separately, how far your opportunity practice has, with the gap between them being the finding.

From there, the right first engagement depends on where you already sit. A company with a completed materiality assessment can move straight to structured discovery against its own disclosed data; a company without one starts by turning what it already has, from scattered documents to trackers and notes, into a usable starting point. Either way, the entry point is a conversation, not a commitment.

A NOTE ON THIS EDITION

This is Edition 1, deliberately exploratory: a first, evidenced attempt at a method we intend to pressure-test through a year of real client engagements before publishing a conclusive follow-up in 2027. If running this method surfaces something that sharpens or complicates what is here, we want to hear it.

About the author



Mike Kelly

Founder, Nemetan

Mike Kelly has spent more than 25 years building companies, and learning the same lesson each time. The best opportunities are usually the ones nobody thought to look for. He is a Chartered Director with the Institute of Directors and is completing an MSc in Sustainability Management at Erasmus University's Rotterdam School of Management. He also holds a BSc in Management from Trinity College Dublin, where he specialised in marketing, and studied disruptive innovation at MIT's Sloan School of Management under Hal Gregersen, co-author of *The Innovator's DNA*. That is the side of his training that made the gap in this report register as a problem rather than a curiosity. A process built to assess what is already in front of it does not look broken from the inside. It looks thorough.

He founded Nemetan to give the Opportunity side of sustainability the rigour Risk has always had, in a field where companies report their Risks two to three times as often as their Opportunities. He is a GRI-certified ESRS Professional with a Certificate in Green and Sustainable Finance from the Chartered Banker Institute.

[Connect with Mike on LinkedIn](#)

CITE THIS RESEARCH

Kelly, M. (2026). *The Opportunity Gap: Risk and Opportunity Asymmetry in European Sustainability Disclosure under ESRS and VSME* (Version 1.0). Nemetan. <https://doi.org/10.5281/zenodo.22124385>

Kelly, M. (2026). *Sustainability Opportunity Discovery: Definition, Scope and Method Requirements* (Version 1.0). Nemetan. <https://doi.org/10.5281/zenodo.22125068>

ORCID: orcid.org/0009-0008-4522-1786

© Michael Kelly trading as Nemetan 2026. This guide may be shared in full and unaltered, with attribution to Nemetan. To reproduce extracts or adapt the material, contact info@nemetan.com. Findings draw on original research by Nemetan together with named public sources; the underlying coded dataset is available on request.



Opportunity, given the same rigour Risk has always had.

Most companies treat their sustainability disclosure as a filing to survive once a year. The evidence in this guide says it is something else entirely: the raw material for a structured search for the value it has been quietly pointing at all along.

- Benchmark your Risk-to-Opportunity balance with the free Gap Check
- Place yourself on the free Maturity Matrix in minutes
- Talk to us about a structured Opportunity engagement