Where Productivity Closes

Industrial productivity read as a closure rate in the net

J. Konstapel — 19 August 2026

1. The dispute, without theory

On 13 August 2026 TNO published a research note commissioned by FNV Metaal: *Industriebeleid met aandacht voor kwaliteit van werk* (TNO 2026 17508, Dhondt & Oeij, 36 pages). It concerns the Dutch metal and electrical engineering sector: over 38,000 firms, more than 350,000 jobs, roughly 18% of Dutch exports.

The policy background is short. In December 2025 Peter Wennink published a report on Dutch earning capacity. The government adopted a growth ambition of 1.5% per year. In July 2026 it issued a productivity agenda. Its centre of gravity is the supply of labour: better foundational skills, more STEM graduates, more upskilling, more international talent, a reformed social security system. The Dutch parliament debates the Wennink report on 3 September 2026.

TNO's finding is blunt. Within five years, none of that raises productivity in this sector.

The report scores each measure on six criteria. The supply-side measures come out at −5, −3, −2, −4, −2 and 0. The demand-side measures — quality of work, technology that supports craft, workplace-bound training, works council participation in technology choices, fair division of gains — come out at +4, +4, +3, +2, +4. Two demand-side measures score −1: education as innovation engine, and regional ecosystems.

The report reads this as a supply/demand distinction. There is a sharper reading available. The sign of every score is set by two quantities, and neither of them is "supply" or "demand".

2. The net, in eight lines

This article uses one model. It is stated here in full so that nothing has to be fetched from elsewhere.

1. The smallest element is a strand. A strand is not a thing that vibrates; it is a self-resonating current.
2. A strand has two states: open, and closed on itself. A closed strand is a knot.
3. Knots hold each other in a mesh. A mesh is the loop that runs from knot to knot and back.
4. Every mesh has a period: the time a signal takes to go round it once.
5. Load raises tension in a mesh. Tension that cannot be discharged accumulates.
6. Discharge follows the weakest mesh available. If no path is open, one is made.
7. The same structure repeats at every scale. Rungs differ in winding depth, not in kind.
8. Nothing in the net is explained by one part acting on another. What survives is what returns in phase. This is selection by fit, not cause.

Two consequences matter here.

A property of a loop cannot be measured at a knot. And a change in the drawing is not a change in the knot.

3. Productivity is a closure rate

Labour productivity is defined as output per worker. That is a measure at a node.

Production is not a node event. An order enters, is prepared, is machined, is handed over, is inspected, is approved, is shipped, is paid. That is a loop. It becomes product at the moment it closes: when the work returns in phase — accepted, delivered, done.

Everything else in the loop is load. Waiting is load. Handovers are load. Rework is load. Approval queues are load. All of it consumes strand and closes nothing.

So the quantity to raise is not output per node. It is the fraction of started loops that close per unit of time, and the length of the path each loop has to travel to get there.

Call the first the closure rate. Call the second the crossing number.

4. Why more people changes nothing: the Reidemeister argument

Knot theory has a precise result about this. A knot can be drawn in infinitely many diagrams. Reidemeister proved that two diagrams show the same knot exactly when one can be turned into the other by three local moves: adding or removing a twist, sliding one strand over another, and moving a strand past a crossing. None of these three changes the knot type. The diagram gets busier or cleaner; the knot is the same knot.

To change the knot type you need a different operation. You must cut a crossing and reconnect it the other way. The minimum number of such crossing changes needed to undo a knot is its unknotting number. That number, not the length of the rope, measures how tangled it is.

Now place the measures on that distinction.

- More workers: add rope. A twist and a slide. Move I and II.
- More STEM graduates: enrich the rope before it enters the mesh. Move I.
- More labour mobility: rearrange which rope sits where. Move II.
- More international talent: import rope. Move II.
- Process innovation, work reorganisation, local control capacity: change a crossing.

Three quarters of the productivity agenda consists of moves that provably preserve the knot. This is not an indictment of the people who wrote it. It is the reason the five-year result comes out the way it does, and it would come out that way in any sector.

The Netherlands has in fact run this experiment at full scale. Between 2015 and 2025 the number of technicians rose by 400,000 to 1.9 million, a 28% increase. Employment in industry rose by nearly half a million working persons in ten years; jobs and labour volume together rose by more than 300,000 between 2014 and 2024. Over that same decade, process innovation in the Metalektro fell from 46% of firms in 2010 to 27.5% in 2021.

More rope, busier diagram, same knot.

One further figure from the same source belongs here: half of all technical positions in the Netherlands are filled by people without technical training, and a substantial share of trained

technicians move into non-technical work. Rope added at one end leaves at the other. The mesh sets the throughput, not the supply.

5. The limit is the crossing number

Every order is a path through the mesh. A crossing is any point where the work passes from one control domain to another: sales to work preparation, planning to floor, floor to quality, quality to sign-off, sign-off to shipping.

At each crossing the loop can stall, because at each crossing the authority to resolve a problem sits on the other side of the boundary from the problem.

Two operations lower the effective crossing number.

Remove the crossing: put the control where the work happens. Or make the crossing reducible: give the person standing at it the knowledge and the room to resolve it without passing it on.

The second is exactly what TNO calls workplace-bound training, and why it scores +3 with a horizon of two to three years. It is not that the worker knows more in general. It is that the loop now closes at the crossing instead of one level up.

It is also exactly what the report calls, in its own words, "reducing unnecessary coordination and bureaucracy" and "inefficient handovers". The crossing is named. It is not counted. Section 10 returns to that.

6. Why five years, and not less or more

Each loop has a period. A measure cannot show its effect before the period of the loop it touches has elapsed. This is arithmetic, not pessimism.

The report's own time horizons, taken from its two tables, sort cleanly by period:

Measure	Horizon in the report	Loop touched
Fair division of gains	immediate	return path already exists
Quality of work	1–2 years	the mesh itself
Technology supporting craft	1–2 years	the mesh itself
Works council voice in technology	1–2 years	the coupling
Workplace-bound training	2–3 years	the crossing
Education as innovation engine	3–4 years	one rung out
Regional ecosystems	3–4 years	one rung out
International talent	≥2 years	supply, short
More technical talent	≥6 years	curriculum
Foundational skills	10 years to act, 20 to show	curriculum
Individual learning accounts	5 to 40 years	the individual career

The five-year horizon does not favour or disfavour anything. It is a window. Measures whose period exceeds the window cannot appear inside it, however good they are. The productivity agenda selects measures with periods of six to twenty years and is held to account on a horizon of five.

Selection by fit applies to policy as it does to everything else in the net. Only what matches the period of the rung it acts on survives on that rung.

7. What the report's six criteria actually are

Appendix 3 scores each measure on six criteria: time to effect, cost, implementation risk, productivity potential, agency, social side effects. Each gets −1, 0 or +1.

In the terms above, those six collapse to two.

- Time to effect is the period of the loop touched.
- Productivity potential and agency together are the distance to the crossing.

The remaining three are derived. Cost, implementation risk and side effects all rise with the number of intermediaries between the actor and the crossing. That is why they move in lockstep with agency in every row of the table.

The test is the two anomalies. Education as innovation engine and regional ecosystems are demand-side measures on the report's own classification, yet they score −1. Both act one rung out from the shop floor. Both have periods of three to four years and no direct agency at the crossing. On the supply/demand distinction they are misfits. On period and distance they land exactly where they should.

The report's scoring reproduces the Reidemeister result without naming it.

8. The loading is measured, and so is the discharge

This is the part of the report with the most value beyond its own subject. Four series, same sector, overlapping years.

- Process innovation, Metalektro: 46% of firms in 2010, 27.5% in 2021. Across all sectors the fall runs from about 50% in 2010 to 28% in 2021, with a partial recovery to about 35% in 2024. Among the smallest firms, fewer than 10% report process innovation at all.
- Employer training effort: down 9% over fifteen years, and down 13% in industry. In the Metalektro the share of firms training 1–49% of staff went from 43% in 2012 to a peak of 50% in 2014, then down to 38.5% in 2021.
- Burnout complaints among Metalektro workers: 14.5% in 2010, a low of 12.3% in 2013, then 18.3% in 2022. Roughly one in five.
- Financial results: the share of Metalektro firms reporting improved profit rose from 25% in 2012 to about 41% in 2019, holding near 40% in 2021.

Add the Dutch central bank series the report cites. Since 2007 the markdown on labour has risen sharply: the gap between the extra revenue an additional worker brings and that worker's gross hourly wage has widened in the firm's favour, and net profit margins have risen with it.

Read as one picture: extraction from the mesh rises, rewiring of the mesh falls, and the difference shows up in the weakest mesh available.

No causal arrow is needed here, and none is offered. These are four readings of one state. A mesh under rising load with a falling rewiring rate has tension that must go somewhere. Where it goes is

decided by what is open. In this sector the open paths are attrition, absence, outsourcing and exit. The burnout series is the measurement of one of them.

What makes this sector useful beyond itself: this is the first case in this series where the loading curve and the discharge curve of the same rung appear in one document, in the same years, from the same instrument. That makes the Metalektro a calibrated example, in the way the FAO food price index is calibrated for the social rung.

9. Technology that does not close

One series in the report is treated as a puzzle. The share of Metalektro workers who report automation of company activities runs 6.2% in 2010, up to 11.3% in 2017, and back to 9.7% in 2022. Flat, through a decade of loud labour shortage. The report states plainly that it expected to find more and did not.

On the model it is not a puzzle.

Introducing a machine, a planning system or an AI tool adds a crossing to the mesh. A crossing can only be absorbed where there is local control capacity to close a loop across it. Where that capacity is absent, the technology is bought and stands still. It appears on the balance sheet and not in the work.

This is the same shape the report's own source [1] measures for generative AI: high adoption spend, low integration into workflow.

It is also why works council participation in technology choices scores +4 with the benefit listed as "lower mis-investment". Mutual coupling before installation decides whether the loop will close at all. One-sided coupling — the machine sets the rhythm, the floor is expected to follow — loads the mesh instead of relieving it. The same result appeared in the citizens' assembly analysis: one-sided coupling accumulates, mutual coupling discharges. Identity across scale, one rung down.

10. What is missing: the counter

The strongest part of TNO's argument rests on its weakest data, and this is worth saying plainly because it is repairable.

The demand-side thesis is carried by the Werkgevers Enquête Arbeid: employer self-report, biennial. Did you introduce a new or significantly improved process, yes or no. Did you see productivity rise, yes or no. Did 1–49% of staff attend training, yes or no. The report says so itself in Appendix 1: these are not a direct measure of labour productivity.

Every one of those is a node measure, taken by asking the node.

The mesh is not measured anywhere, in this report or in the productivity agenda. Yet the quantities that decide it are already in systems every firm runs — order entry, work orders, status changes, throughput times.

Four counters, all derivable from existing data, all readable within a quarter:

1. **Closure rate.** Share of started work items that close without returning to an earlier stage. Rework and re-approval are returns.

2. **Crossing number.** Average number of transfers between control domains per delivered item.
3. **Queue-to-hand ratio.** Time an item spends waiting divided by time an item spends being worked on.
4. **Local resolution rate.** Share of disruptions resolved at the place where they arise, without escalation.

These four are a tension meter for the shop floor. They belong to the same family as the coherence measure C(t) used elsewhere in this series: one number per mesh, comparable across firms and across rungs, moving before the outcome moves.

They also make the report's central claim testable rather than argued. The prediction is specific: firms that lowered their crossing number show productivity gains inside the five-year window; firms that added headcount or off-site training do not, regardless of size and sub-sector. TNO can run this against its own panel.

11. What the report does not reach

Three things sit just outside its frame.

Discharge that leaves the sector. The report's boundary is the firm. Tension that discharges outward — long-term incapacity, disability claims, outsourcing, insolvency — falls off its books, even though the burnout series that leads there is inside them. The rising inflow of young workers into Dutch disability schemes with mental health complaints belongs to this same reading.

Division as closure, not fairness. The report treats the distribution of productivity gains as a question of legitimacy and support. On the model it is a closure question. A loop whose yield does not return to the place where the work happened has not closed for that place. This is why fair division is the only measure in the whole table with an immediate horizon. Nothing has to be built. An existing return path is reconnected.

The mesh is deeper than the firm. ASML and a small machining shop are on different winding depths. Deep loops run for years across hundreds of crossings; shallow loops run for days across ten. One national percentage is a sum over meshes of different depth, steered from the strand side. That aggregate has no point of attachment. It can be reported. It cannot be driven.

12. Tests, and what would kill this

Five things this reading forbids. If any of them turns up, it is wrong.

1. A firm that raised headcount and off-site training substantially, changed no crossings, and showed sustained productivity gains inside five years.
2. Crossing number and closure rate that show no relation to productivity across a panel of firms.
3. A sector where supply-side measures produced measurable productivity gains on a horizon shorter than the period of the loop they act on.
4. Technology adoption proceeding normally in meshes with no local control capacity.
5. A loading period — falling rewiring under rising load — that runs for a decade with no rise in any discharge channel.

Point 5 is the one to watch. It is the strongest claim here and the easiest to check against other sectors.

13. Status ledger

Stated once, centrally, rather than woven through the text.

Element	Status
Reidemeister moves preserve knot type; crossing	Theorem. Established mathematics.
Productivity as a closure rate of the mesh	Definition within the model.
Supply measures map to Reidemeister moves, demand measures to crossing changes	Identification. Proposed here, supported by the report's own scoring.
Six criteria reduce to period and distance to crossing	Derivation from the model, checked against
Loading and discharge series read as one state	Reading of measured data. The data are TNO's; the reading is the model's.
The four counters	Instrument proposal. Not yet built, not yet
Prediction on crossing number and productivity	Testable. Not yet tested.

14. Conclusion

TNO is right that the answer lies inside the firm, and the report deserves to be read for its data as much as for its argument. The reason it is right is stronger than the reason given.

Productivity is not a property of a worker. It is the rate at which loops close in a mesh. Adding rope leaves the knot unchanged; only changing a crossing changes it. Whatever loop a measure touches, that loop's period sets when the measure can show. Both facts hold at every scale, which is why the report's own scoring reproduces them without needing the vocabulary.

What is missing is not another agenda. It is a counter for the thing everyone now agrees is the problem.

Annotated references

1. TNO (2026). ***Industriebeleid met aandacht voor kwaliteit van werk*****. TNO 2026 17508, Dhondt & Oeij, 13 August 2026. Commissioned by FNV Metaal.** The primary source for this article. Why read it? Because the two appendices are worth more than the recommendations. Appendix 1 holds the WEA and NEA series on process innovation, training effort, burnout and perceived automation; Appendix 3 holds the six-criterion scoring. Reading advice: start at Appendix 1, then Appendix 3, then the summary. The argument in the body is a compression of what those two already show.

2. Wennink, P. (December 2025). ***De route naar toekomstige welvaart. Een sterk Nederland in een relevant Europa*****.** The report that set the Dutch productivity debate and the 1.5% ambition. Why read it? To see the supply framing in its purest form, in chapter 3.2 (pp. 52–59), where the five-point National Talent Agenda is set out. Reading advice: read those eight pages against the horizon column of TNO's Table 1.

3. Draghi, M. (September 2024). ***The future of European competitiveness*****, Part A.** The European antecedent of Wennink. Why read it? Because it shows the same structure at continental scale: innovation and competitiveness foregrounded, the organisation of work barely present.

4. Brouwer, G. & de Winter, J. (2026). ***Gezonde bedrijven, gezonde groei: dynamiek, marktmacht en misallocatie in het Nederlandse bedrijfsleven*****. DNB.** The markdown series. Why read it? Figure 18 and the surrounding pages give the widening gap between marginal product and wage since 2007, alongside rising margins. Reading advice: this is the extraction side of the loading picture in section 8. Read it directly after TNO's Appendix 1.

5. Adema, Y. et al. (2025). ***Proces van creatieve destructie verzwakt in Nederland*****. CPB.** Why read it? For the declining-dynamism finding from a second, independent instrument. Where TNO sees falling process innovation, CPB sees weakening turnover of firms and resources. Two readings of the same stalled mesh.

6. Ministerie van Economische Zaken en Klimaat (2026). ***De tweede productiviteitsagenda*****.** The policy under discussion. Why read it? To check the period claim yourself: take each measure and ask how long its loop runs. Reading advice: the answer is in the measures, not in the framing.

7. Acemoglu, D. & Restrepo, P. (2019). "Automation and new tasks: how technology displaces and reinstates labor." ***Journal of Economic Perspectives*** **33(2), 3–30.** Why read it? For the distinction between automation that only displaces and automation that creates new tasks. In the terms used here: technology that removes a crossing versus technology that adds one. TNO cites it in support of works council involvement, and it carries that weight.

8. Acemoglu, D., Autor, D. & Johnson, S. (2026). ***Building pro-worker artificial intelligence*****. NBER Working Paper 34854.** Why read it? Because it makes design, rather than adoption speed, the variable that decides the outcome. Reading advice: read it as the technology chapter of the same argument, one rung up.

9. Challapally, A. et al. (2025). ***The GenAI Divide: State of AI in Business 2025*****. MIT NANDA.** Why read it? It measures the gap between AI purchased and AI integrated into workflow. That gap is the flat automation series of section 9, in a different technology and a different decade. Reading advice: the finding matters more than the framing.

10. Sevcenko, V., Wu, L., Kacperczyk, A. & Ethiraj, S. (2022). "Surplus division between labor and capital: a review and research agenda." ***Academy of Management Annals*** **16(1), 334–390.** Why read it? The literature on who captures the gains, gathered in one place. Read here as the literature on whether the loop returns to where the work happened.

11. Oeij, P., Pot, F., Van den Bossche, S. et al. (2024). ***Goed Werk: transities naar goede banen leiden*****. TNO whitepaper.** Why read it? It is the concept of quality of work that the main report leans on: autonomy, learning opportunity, voice. Reading advice: read it for the operational definitions, which the main report assumes rather than restates.

12. Breque, M., De Nul, L. & Petridis, A. (2021). ***Industry 5.0: towards a sustainable, human centric and resilient European industry*****. European Commission.** Why read it? The European frame TNO recommends for the sector. Useful as a vocabulary for the same intent, and worth reading critically: human-centricity is asserted as a value, not derived from anything.

13. Reidemeister, K. (1927/1932). ***Knotentheorie*****; and Kauffman, L. (1991).** ***Knots and Physics*****, World Scientific.** The mathematics behind section 4. Why read it? Kauffman is the accessible entry: diagrams, moves, invariants, and why a diagram change is not a knot change. Reading advice: the first three chapters of Kauffman are enough for everything used here.

14. Konstapel, J. (2026). "The Net — the vacuum from the smallest strand to the greatest knot." The founding article of this series. Why read it? It sets out the eight lines used in section 2

and the ladder of rungs on which the identity claim in sections 6 and 9 rests. Reading advice: read this if the phrase "one rung down" needs to mean something precise.

15. Konstapel, J. (2026). "Where Democracy Discharges — the citizens' assembly read as a coupling between two rungs of the net." Why read it? Same operation, one rung up: one-sided versus mutual coupling, and what happens to an assembly with no discharge path. The transfer coefficient there and the crossing number here measure the same thing on different rungs.

16. Konstapel, J. (2026). "Where War Comes From" and "Where the Exodus Comes From." Why read it? For the loading-and-discharge reading applied where the discharge is not measured in absence rates. The FAO food price index in the exodus article plays the role the burnout series plays here: a calibrated meter on a social rung.