iMORSE

The Australian National Competitive Grants Program under Reform

A system-dynamics account of the ARC's 2027 changes and what they are likely to deliver

Andrew Flitman
iMORSE
July 2026

Executive Summary

The Australian Research Council is reshaping its National Competitive Grants Program from 2027. It will remove salaried mid-career and senior fellowships and redirect the money to projects, stand up two dedicated schemes for the next generation - Initiate for early-career researchers and Lead & Mentor for the senior researchers who train them - lift dedicated Indigenous funding from about one per cent to five, and cut fifteen schemes to seven. This report tells the story of a system-dynamics model built to test what those changes are likely to deliver, with every one of the seven new schemes represented.

The model is calibrated to documented ARC and government data from 2010 and reproduces it: the rising fellowship share, the near-flat-to-declining real budget, the senior-skewed workforce, the growth in industry collaboration, and low Indigenous participation. It was then validated - the two simulation engines agree to machine precision, stocks stay sane under extreme settings, every lever moves the model, and the backcast tracks the observed history.

The finding separates into two mechanisms.

Removing fellowships is what recovers funded project volume, lifting it from around 1,800 with no reform to around 2,470 by the mid-2030s - but that same removal raises mid-career precarity.

The reform's early-career ambition is carried by a different mechanism: the Initiate and Lead & Mentor schemes lift early-career researchers from around 2,180 to around 3,440 and grow the engaged workforce from around 11,300 to around 13,900, close to the ARC's own 'up to 20 per cent more researchers' claim.

With those schemes switched off, early-career numbers stay flat and the workforce keeps contracting - so the early-career gain rests on the schemes working, not on the money shuffle. Indigenous capability grows strongly under the funding step, though it lags the money by years because capability is a slow-building stock.

This is delivered as an interactive simulator so we can test the reforms, and the contested assumptions behind them, for ourselves.

Introduction

The intent of this report is to explain the model. It sets out what the model of the National Competitive Grants Program does, how it is built from feedback loops, what evidence informed it, and what it was validated against. It is written for an academic and policy audience.

This is not the technical manual. The equations, units and full parameter set live in the accompanying Model Specification. The point is to make the behaviour of the funding system legible, so that the reforms can be judged on how the system responds over time rather than on a static allocation.

The Problem

Since 2004 the NCGP has grown from six schemes to fifteen, and the share of its funding going to fellowships has risen from under ten per cent to forty-six. Over the same period the average cost of

a project climbed from about 440,000 dollars in 2015 to 720,000 by 2024, well ahead of inflation, while the real budget stayed flat and then fell. The funded workforce skewed senior. Indigenous researchers remained about one per cent of the total.

These are not independent facts. A near-flat budget is a shared pool. As fellowships took up more of it and each project cost more, fewer projects could be funded, and the funded workforce they support contracted. Discovery success rates fell from twenty-three per cent in 2020 to sixteen in 2024, and the number of projects funded fell with them.

This is a dynamic problem. It is about accumulation and feedback over time, not a one-off sum. The research workforce is a set of stocks that fill and drain with long delays. Grants tie up budget for years. Reputation and capability compound. The method for problems of this shape is system dynamics, developed by Jay Forrester and set out by John Sterman, which represents a system as stocks, the flows that change them, and the feedback loops that connect them.

How the Model Sees the System

The model is built from ten feedback loops. Five reinforce and five balance.

- The first reinforcing loop is workforce regeneration. Senior and mid-career researchers supply the mentoring and supervision that fill the PhD pipeline, which produces early-career researchers, who age into mid-career and senior roles and mentor in turn. Run with seniors retained it renews the workforce. Run in reverse, when they leave, it hollows it out. Its delays are long: four years for a PhD, five to mid-career, ten more to senior.
- The second is success breeding success. Researchers with a strong record win more, produce more, and burnish the record, concentrating funding on a proven few.
- The third is collaboration maturity, the loop behind the rise in industry cash from a third to two thirds of collaborative funding.
- The fourth builds Indigenous capability, slowly, as funded researchers mentor the next cohort.
- The fifth drifts funding toward fewer, larger grants.

The balancing loops are the limits.

- One is the shared-budget commons: every scheme draws on one flat pool.
- One is absorptive capacity: money converts to projects only as fast as the eligible workforce allows.
- Another is peer-review throughput.
- The fourth is administrative burden.
- The last is labour-market leakage and precarity, which drains the workforce toward the outside economy and is the mechanism of the feared mid-career cliff.

Read together, these balancing loops are also familiar archetypes (per system dynamics theory). Limits to Growth, as the reform's ambitions meet slow supply. Tragedy of the Commons, as cohorts contend for one budget. Success to the Successful, as funding concentrates. And Fixes that Fail, as removing fellowships to help early-career researchers undermines the mentoring that produces them - which is precisely why the reform pairs removal with the Lead & Mentor scheme, the model's countermeasure to that trap.

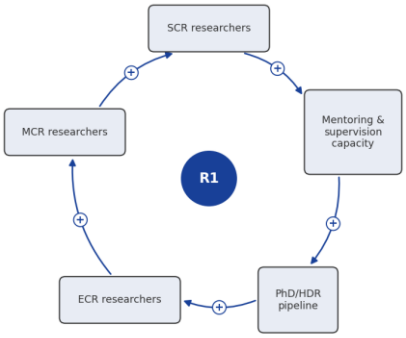

R1 - Workforce regeneration (seed corn) (reinforcing loop)

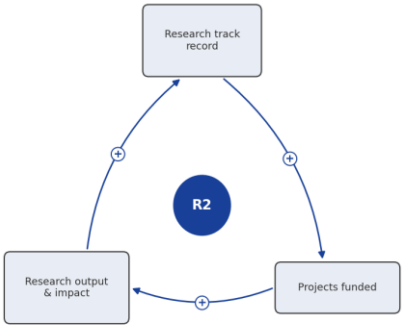

R2 - Success breeds success (Matthew effect) (reinforcing loop)

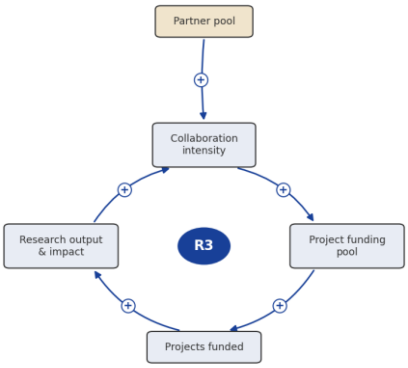

R3 - Collaboration maturity (reinforcing loop)

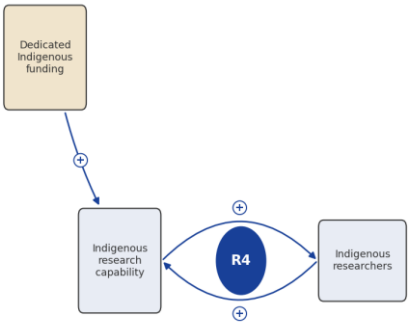

R4 - Indigenous capability build (reinforcing loop)

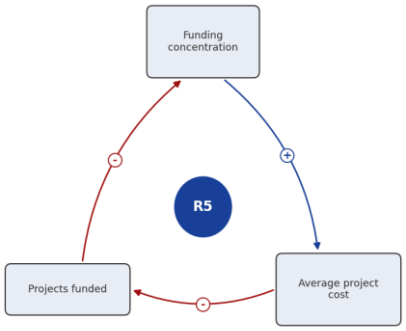

R5 - Grant-size concentration (reinforcing loop)

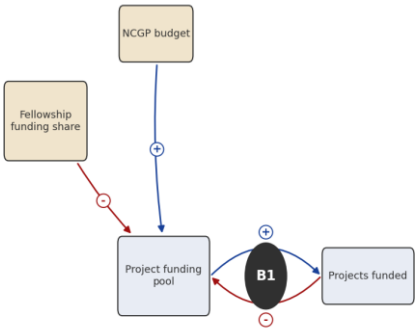

B1 - Shared-budget commons (balancing loop)

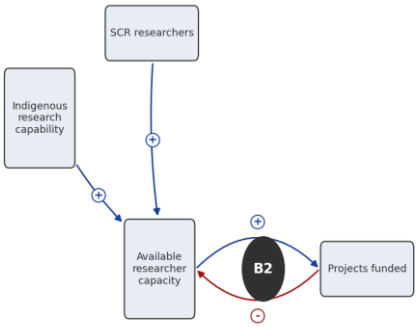

B2 - Absorptive-capacity constraint (balancing loop)

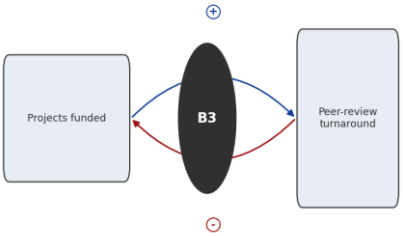

B3 - Peer-review throughput (balancing loop)

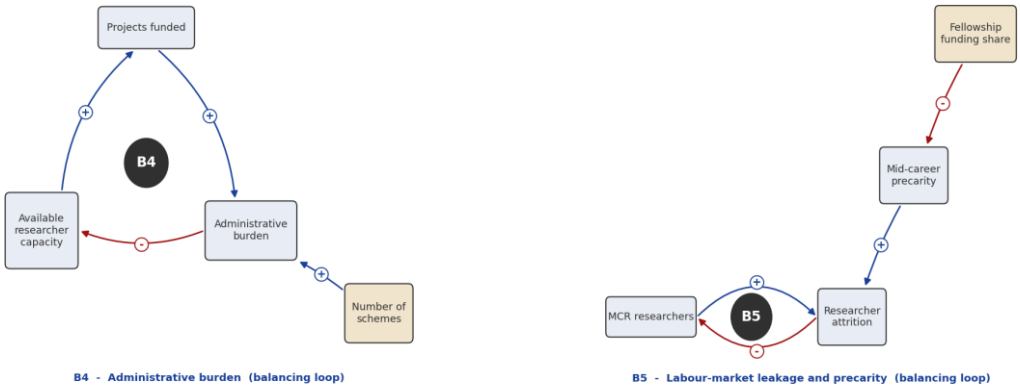

Figure 1. The ten feedback loops. The model combines these into a single integrated system.

From Diagram to Simulation

The loops name what influences what. The simulation says what accumulates. The stocks are the things with memory: the three workforce cohorts, the PhD pipeline that feeds them, the active project portfolio, the Indigenous researcher pool, the reputation and knowledge stocks, the cost level and the collaboration level. Each would keep its value if every flow stopped. That is the test of a stock.

What sits outside the model is deliberate. The total ARC budget, the cost of research, the national supply of PhD graduates and the pull of the outside labour market are given to the model, not generated by it. The reforms are exogenous switches - fellowship removal, the Initiate early-career channel, the Lead & Mentor mentoring channel, the Indigenous funding step and scheme simplification - because they are decisions, not outcomes. So are the two contested assumptions.

The delays are where the behaviour comes from. The workforce pipeline turns over across decades. Grants commit budget for years. Capability builds slowly. A model without these delays would show the reforms landing at once. The whole point is that they do not.

Evidence and Validation

The calibration is anchored to documented sources and flagged by confidence. The observable outcomes are firm or approximate: the career-stage shares from the review, the average project cost, the industry-collaboration share, the Indigenous count, the fellowship-share trajectory, the budget series from ARC annual reports, and the estimated return of about three dollars and thirty-two cents per dollar invested. The behavioural relationships - the loop strengths and adjustment times - are structural assumptions, not measured, and they are the model's main uncertainty. The two most consequential are exposed as levers rather than fixed.

Validation ran in two stages.

The first asked whether the model is built correctly. A Python reference and the JavaScript engine that ships in the simulator agree to machine precision, so the tool the reader uses is the tool that was tested. Stocks never go negative under extreme settings, and every lever moves the model.

The second asked whether it reproduces reality. The backcast tracks the observed career shares, cost, collaboration and Indigenous numbers within tolerance, and it did so while being driven by the real historical inputs. Figure 2 shows this.

The model was then exported to the vendor-neutral XMILE standard and reproduced exactly under PySD, so it does not depend on any single tool to survive. One objective caution reads the backcast correctly: the workforce shown is the pool engaged with the NCGP, which contracts under the funding squeeze, and is not the whole university academic workforce, which grew.

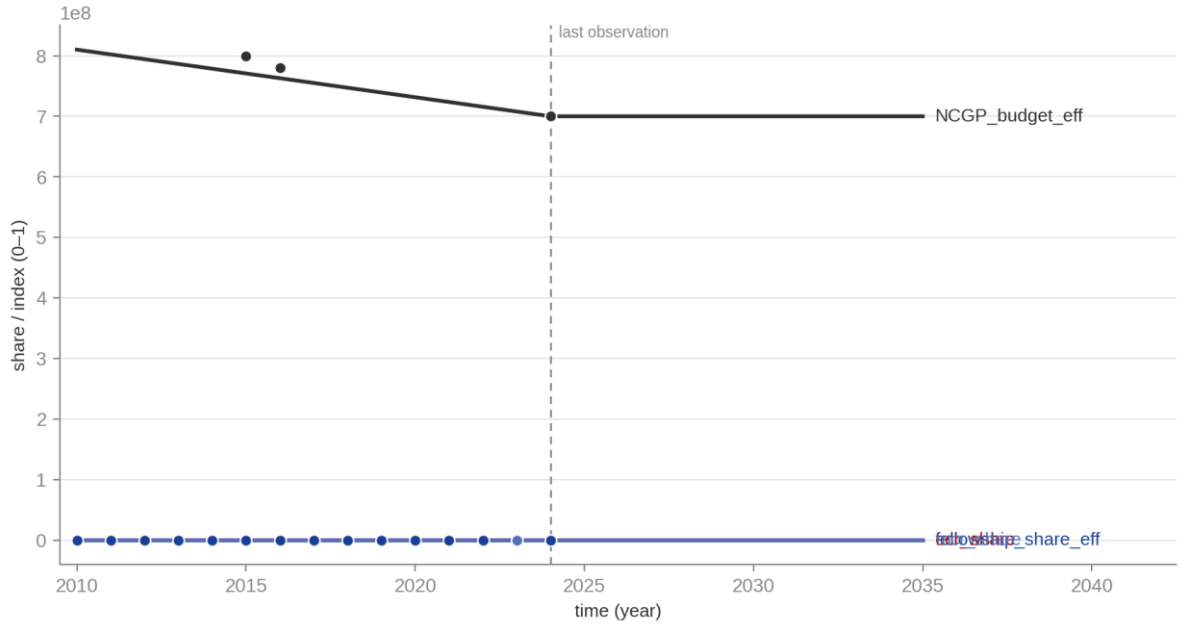

Figure 2. The calibrated backcast against observed history, with the last observation marked and the projection to its right.

What the Model Says

First, the pre-reform path was unsustainable for projects. A rising fellowship share and a falling real budget against rising costs squeeze funded project volume down through the period. The reform is a response to a real decline, not an improvement on a healthy baseline.

Second, it is fellowship removal, specifically, that recovers project volume - lifting funded projects from about 1,800 to about 2,470. The Indigenous and scheme reforms, on their own, do not; keep the fellowships and project volume stays near the squeezed baseline at about 1,680.

Third, the early-career recovery comes from the schemes, not the money. The Initiate and Lead & Mentor schemes lift early-career researchers from about 2,180 to about 3,440 and grow the workforce to about 13,900. Switch those two channels off, leaving the fellowship money still redirected to projects, and early-career numbers stay flat at about 2,180 - the headline early-career ambition rests almost entirely on the schemes recruiting and mentoring at the assumed scale.

Fourth, the recovery carries a cost the headline misses. Removing salaried senior and mid-career fellowships raises mid-career precarity sharply; the schemes help the early-career cohort but do not undo that. Precarity is displaced up the career pipeline rather than removed.

Fifth, the Indigenous ambition is reachable but slow. The five-fold funding step lifts Indigenous researchers from about 300 to about 460, under any variant of the reform, but the rise lags the money by years because capability cannot be bought at once.

Boundary and Limitations

The model is aggregate. It does not resolve individual disciplines, institutions or grants, and it treats medical research and the wider innovation system as outside its boundary. The economic return is carried as a light index, not a macroeconomic forecast. The levels it reports are illustrative and rest on behavioural parameters that are assumptions rather than measurements.

Read shapes, not points. The calibration is firm on the observable outcomes and structural on the behavioural relationships, and the objective reading is of the directions, the turning points and the gaps between scenarios, not the third decimal place.

Conclusion

The reform recovers funded projects, and it does so by removing fellowships, which is also what unsettles the mid-career workforce. Indigenous capability grows, slowly. Whether the trade is worth making is a judgement, and the model exists to inform it rather than to settle it.

The next step is to put the simulator in front of the board and let it test the reforms, and the assumptions behind them, for itself.

References

Australian Research Council (2026). National Competitive Grants Program Policy Review Final Report.

Australian Research Council. Annual Reports 2014-15, 2015-16, 2023-24.

Australian Research Council. Discovery Projects Selection Reports, 2020-2024.

ACIL Allen (2023). Impact assessment of ARC-funded research.

Universities Australia (2020). Research workforce statistics.

Australian Department of Education. Higher Degree by Research completions time series.

Reserve Bank of Australia. Inflation data.

Australian National Audit Office. Audit of the ARC's administration of the NCGP.

Forrester, J. W. (1961). Industrial Dynamics. MIT Press.

Sterman, J. D. (2000). Business Dynamics: Systems Thinking and Modeling for a Complex World. McGraw-Hill.

Senge, P. M. (1990). The Fifth Discipline. Doubleday.

This model was validated against the documented history of the Australian National Competitive Grants Program from 2010 to 2024: the fellowship-share trajectory and career-stage shares from the ARC NCGP Policy Review, the average project cost, the industry cash-contribution share, the Indigenous researcher count, and the NCGP budget series from ARC annual reports. It passed cross-engine parity to machine precision, extreme-condition and live-control checks, backcast reproduction

of the fit targets within tolerance, bounded forward and extreme extension, and an XMILE-to-PySD durability proof reproducing the reference exactly.

Appendix: the integrated causal loop diagram

The ten loops above combine into one system. The integrated diagram below shows the whole feedback structure, with same-direction influences in blue and opposite-direction influences in red, and the exogenous drivers as dashed boxes.

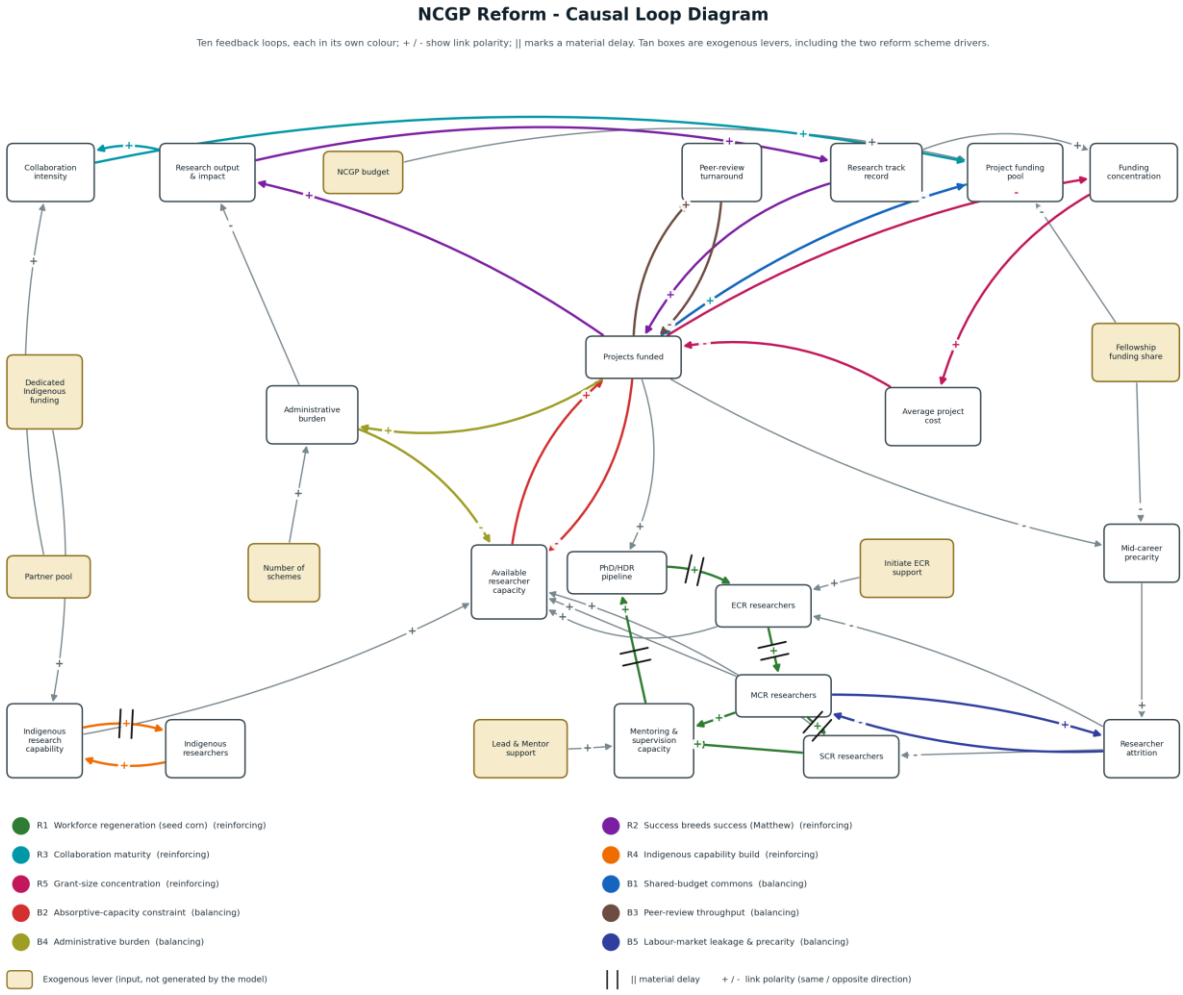

The integrated causal loop diagram of the NCGP model.