ABM and assessing SES resilience…
…what could possibly go wrong?

Bruce Edmonds

Centre for Policy Modelling, UK.

UiT Norway University of the Arctic, Tromsø.

Umeå University, Sweden.

Part 1

Foundations and Motivation

Some key questions

In trying to avoid disasters in a complex Socio-ecological System (SES)…

- Are there generic measures of resilience that will help us?
- How can we know what to measure?
- How can we know what level of perturbation needs to be tested?
- How can we know what kinds of perturbations we need to worry about?
- How can ABM help us in all this?

SES modelled as an algorithm

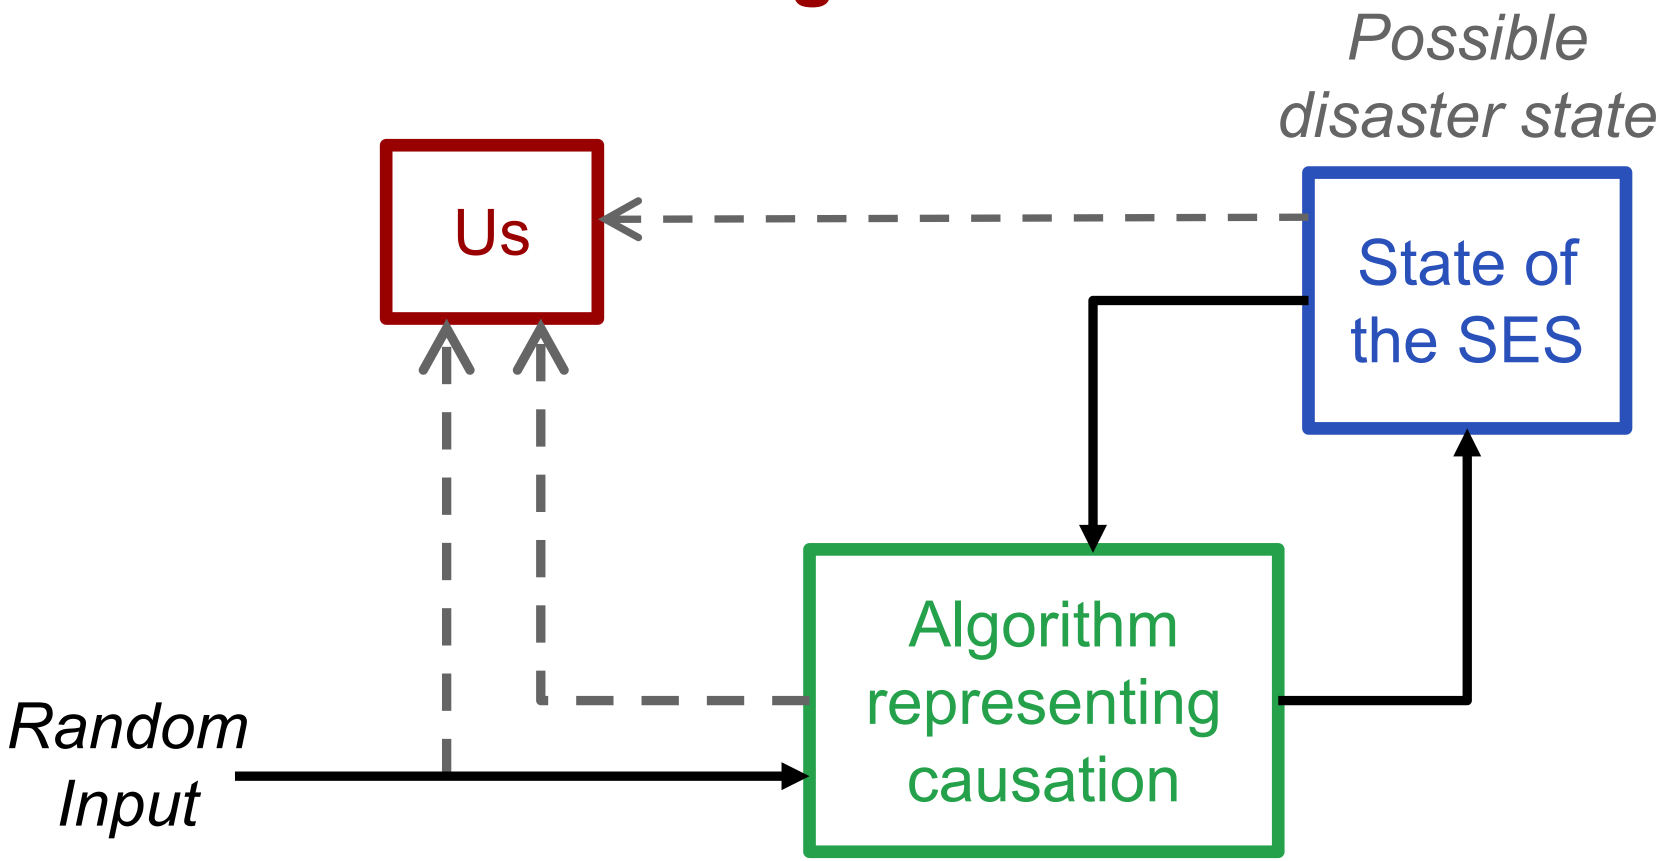

- The most general model of a SES is an algorithm
- Here there is some state of the world, operated on by a program with an input that we can change

Consequences of this…

Even if we have *complete* knowledge of:

- The current state of the SES system
- The algorithm (representing the causation)
- The random input into the algorithm

…is there a method or algorithm that will allow us to predict if the SES algorithm will make the state of the SES reach a disaster?

In general, the answer is ***No****,*
this is impossible.

And even some apparently simple systems are Turing complete

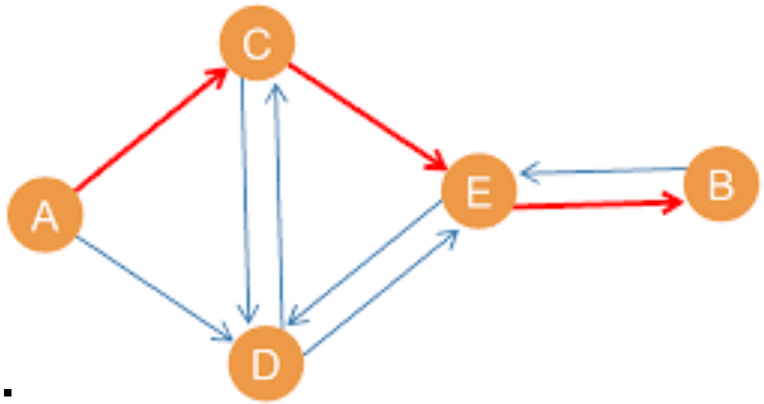

Imagine a class of SES composed of a set of locations which pass units of something between them.
Some units ‘rain’ down on them each tick.
Each nodes has a fixed set of patterns of passing these units to each other (if they have any left)
But which of these patterns they follow might depend on whether another location has run out of these units or not.

Such systems can do any computation

SES modelled as a learning/exploration process finding peaks in a landscape

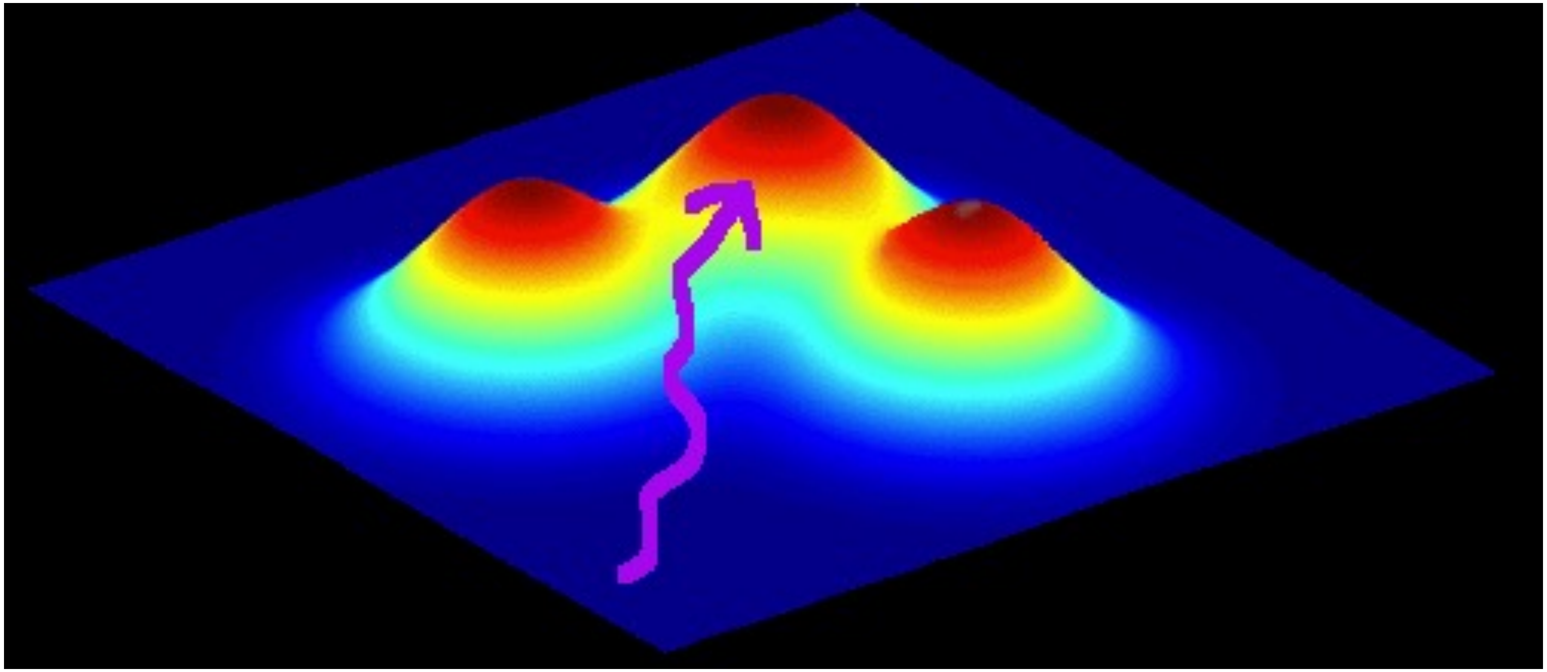

- Which learning algorithm is best, in general, for navigating such landscapes?

No Free Lunch Theorems

- Again, the answer is negative in general
- In an absolute/abstract sense, all learning algorithms are equally good/bad
- If a learning approach is good on a certain kind of fitness landscape, it will be bad on another set

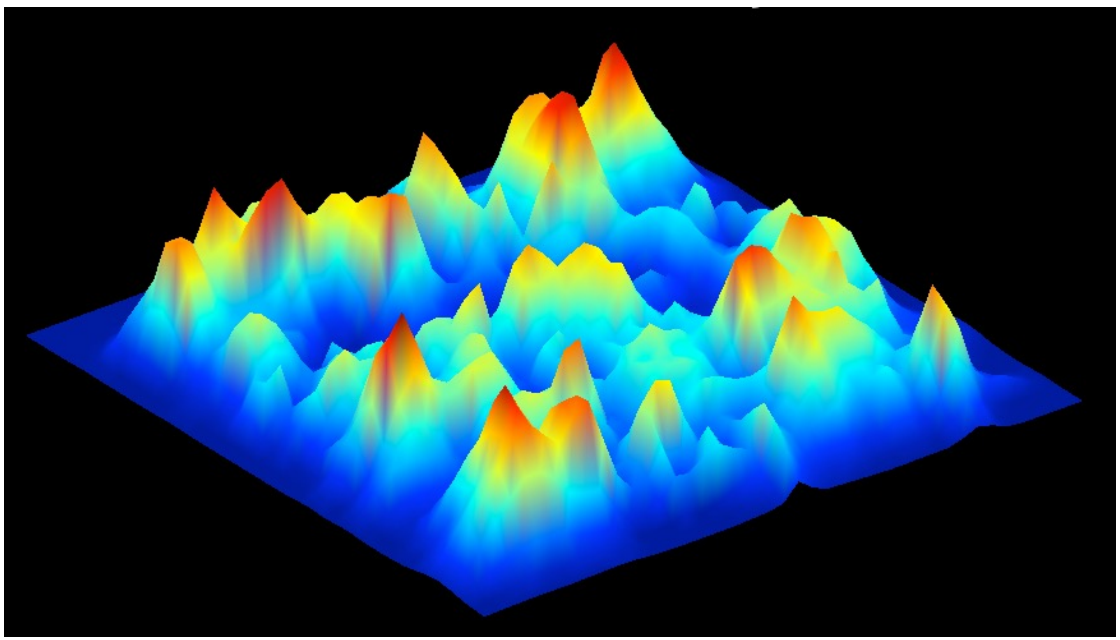

Lessons I take from these abstractions

- That one can not assume that any particular measure will be adequate for all SES
- Similarly, there is probably no *generic* approach for managing SES to keep them from disaster
- Rather, one needs to understand the possibilities inherent in the kind of SES one is dealing with
- And base your your approach on this understanding

Part 2

Lessons from elsewhere

Lessons from robotics: Part I

Robotics in the 70s and 80s tried to (iteratively):

1. build a map of its situation (i.e. a predictive model)
2. use this model to plan its best action
3. then try to do this action
4. check it was doing OK go back to (1)

But this did not work in any realistic situation.

- It was far too slow to react to its world
- to make useable predictions it had to make too many dodgy assumptions about its world

Lessons from robotics: Part II

Rodney Brooks (1991) Intelligence without representation. Artificial Intelligence, 47:139–160

A different approach:

1. Sense the world in rich fast ways
2. React to it quickly
3. Use a variety of levels of reaction
    a. low simple reactive strategies
    b. switched by progressively higher ones

Do not try to predict the world, but react to it quickly

This worked much better.

Lessons from Weather Forecasting

- Taking measurements at a few places and trying to predict what will happen based on simple models based on averages does not work well
- Understanding the weather improved with very detailed simulations fed by rich and comprehensive sensing of the system
- Even then they recognize that there are more than one possibilities concerning the outcomes (using ensembles of specific outcomes)
- If these indicate a *risk* of severe weather they issue a warning so mitigating measures can be taken

Lessons from Radiation Levels

- The human body is a very complex system
- It has long been known that too much radiation can cause severe illness or death in humans
- In the 30s & 40s it was assumed there was a "safe" level of radiation
- However it was later discovered that ***any*** level of radiation carried a risk of illness
- Including naturally occurring levels
- Although an increase in radiation might not seem to affect many people, it did result in more illnesses in some

North Atlantic Cod Fishery Collapse

- In July 1992 Canada's fisheries minister placed a moratorium on all cod fishing off the NE coast of Newfoundland and Labrador. That day 30,000 people lost their jobs and hundreds of years fishing for cod off those coasts ended.
- Models being used predicted healthy stocks up until 1989, and hence *had made the problem worse*.
- Subsequent Harris report: "…scientists, lulled by false data signals and, to some extent, overconfident of the validity of their predictions, failed to recognize the statistical inadequacies in their bulk biomass model and failed to properly acknowledge and recognize the high risk involved with state-of-stock advice based on relatively short and unreliable data series."

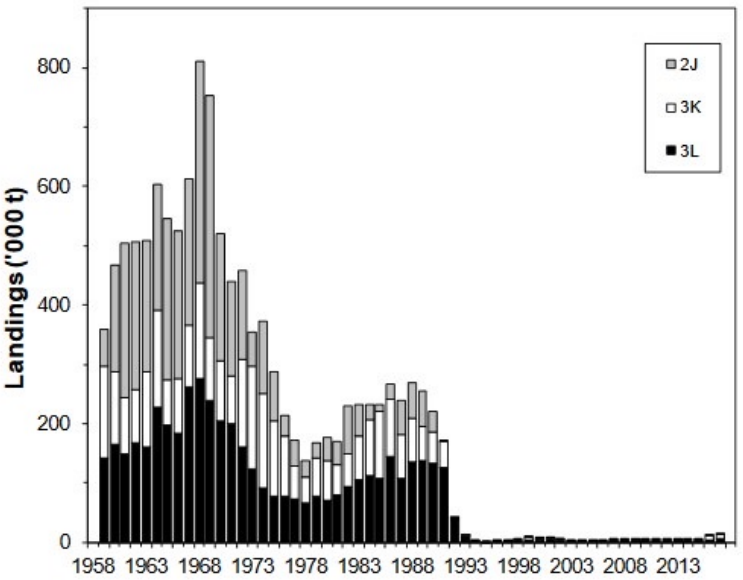

Global Fisheries Collapses

- Not limited to Atlantic Cod
- Complete lack of primary data
- Models do not capture complex inter-species interactions
- Let alone the possible consequences of fishing

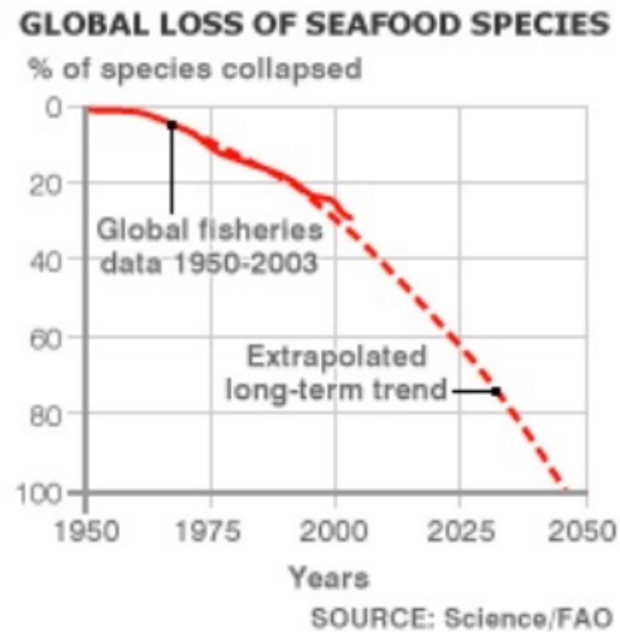

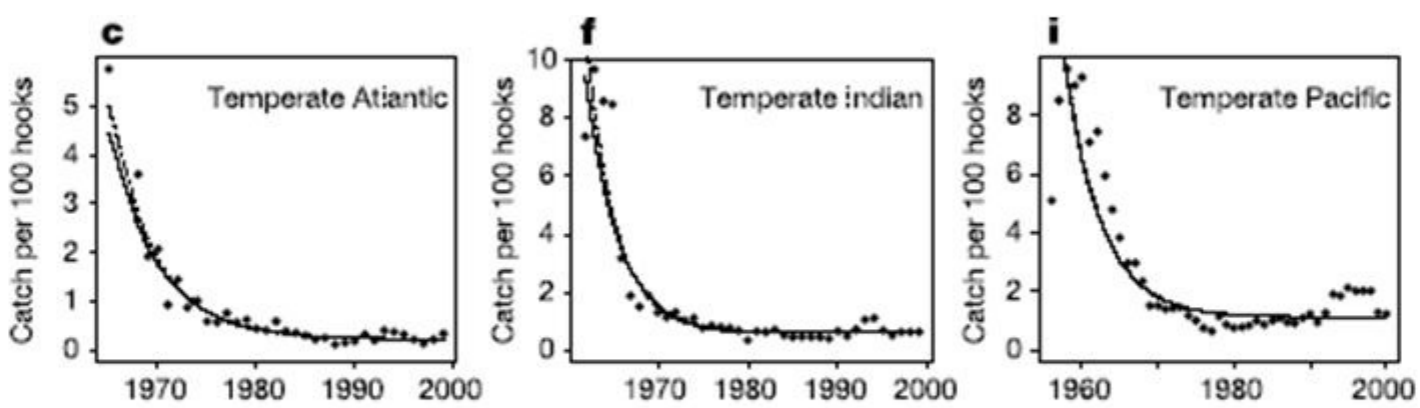

Why simple models won't work

- Simple models are far more convieniant, so many excuses for using them are made but…
- Simpler models do not necessarily get things "roughly" right
- …and they are *not* more general
- They can also be very deceptive – especially with regards to complex ways things can go wrong
- In complex systems the detailed interactions can take outcomes 'far from equilibrium' and far from average behaviour
- Sometimes, with complex systems, a simple model that relies on strong assumptions *can be far worse* than having no models at all

A risk-analysis approach

1. Give up on estimating future impact or “safe” levels of exploitation
2. Make simulation models that include more of the observed complication and complex interactions
3. Run these lots of times with various scenarios to discover some of the ways in which things can go surprisingly wrong (or surprisingly right)
4. Put in place sensors/measures that would give us the earliest possible warning that these might be occurring in real life
5. React quickly if these warning emerge

In this talk I describe…

- …a complex simulation of an ecosystem in which the impact of humans can be included to illustrate such a risk-analysis approach
- The model does not intend to be approximately right or give any indication of what will probably happen
- But rather to reveal some of the real possibilities – things that might happen
- It shows how unpredictable its outcomes can be

Part 3

Exploring an abstract model of complex fishery ecosystems

– general description

Design Criteria for the Model

To exhibit emergent:

- detailed entity-entity interactions
- complex food webs between many species
- co-evolutionary development
- spatial complexity (different niches, diffusion processes, predator waves, etc.)
- all entities embedded within the spatial nutiritional ‘economy’
- possibility of invasive species, extinctions, new species by mutation etc.

This model…

- …is a dynamic, spatial, individual-based ecological model that has some of the complexity, adaptability and fragility of observed ecological systems with emergent outcomes
- It evolves complex, local food webs, endogenous shocks from invasive/emergent species, is adaptive but unpredictable as to the eventual outcomes
- Into this the impact of humans can be imposed or even agents representing humans ‘injected’ into the simulation
- The outcomes can be then analysed at a variety of levels over long time scales, and under different scenarios.
- Full details and code at: http://openabm.org/model/4204

The Model

A well-mixed patch

Each individual represented separately

Slow random rate of migration between patches

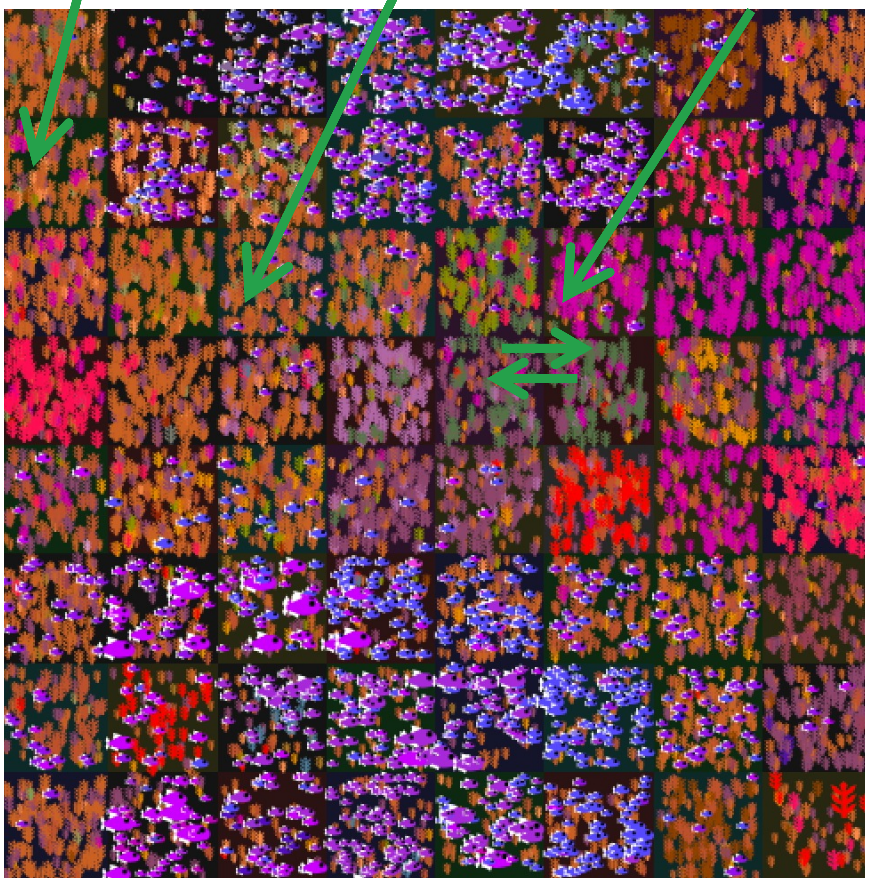

- A wrapped 2D grid of well-mixed patches with:
  - energy (transient)
  - bit string of characteristics
- Organisms (plants and fish) represented individually with its own characteristics, including:
  - bit string of characteristics
  - energy
  - position
  - stats recorders

How Dominance is Decided

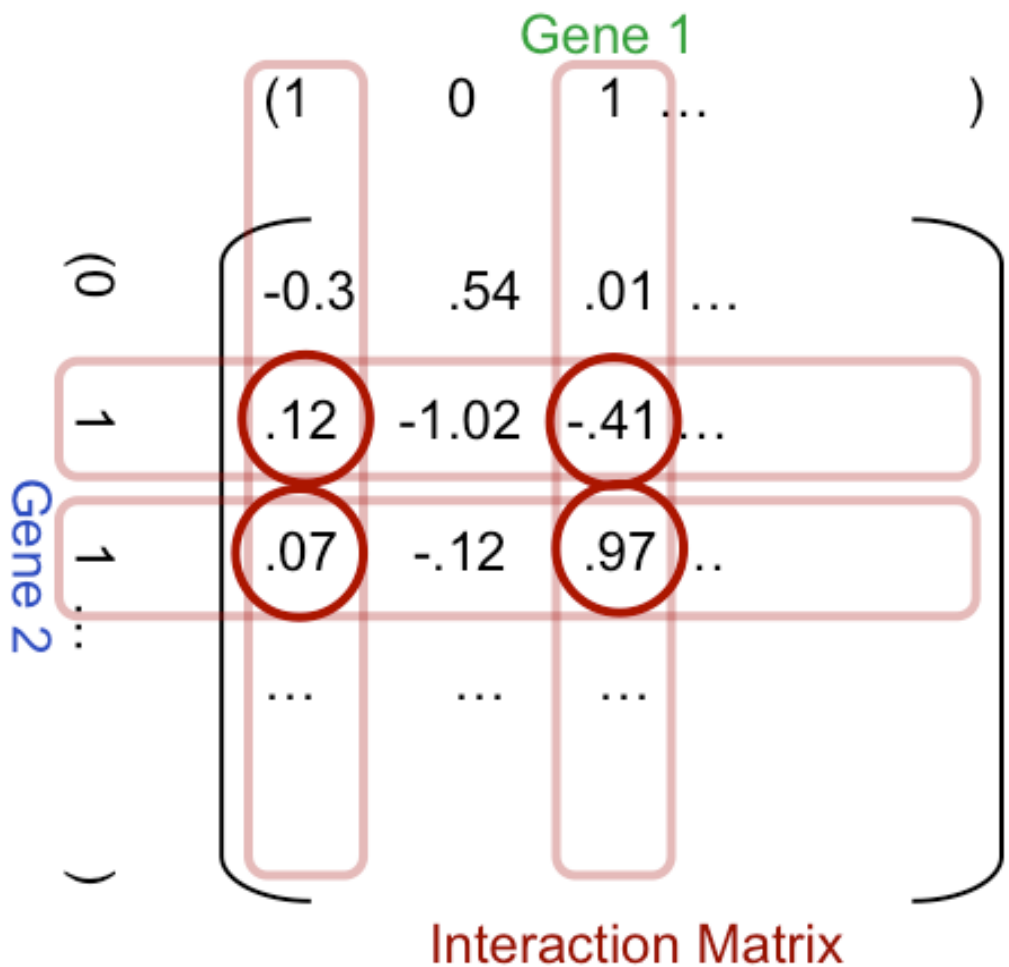

Resulting value:

.12+-.41+.07+.97
= 0.75 (which is > 0)

So an individual with Gene 1 would be able to eat one with Gene 2

(Caldarelli, Higgs, and McKane 1998)

Model sequence each simulation tick

1. **Input energy** equally divided between patches.
2. **Death**. A life tax is subtracted, some die, age incremented
3. **Initial seeding**. until a viable is established, random new individual
4. **Energy extraction from patch**. energy divided among the plants there with positive score when its bit-string is evaluated against patch
5. **Predation**. each non-plant individual is randomly paired with a number of others on the patch, if dominate them, get a % of their energy, other removed
6. **Maximum Store**. energy above a maximum level is discarded.
7. **Birth**. Those with energy > “reproduce-level” gives birth to a new entity with the same bit-string as itself, with a probability of mutation, Child has an energy of 1, taken from the parent.
8. **Migration**. randomly individuals move to one of 4 neighbours
9. **Statistics**. Various statistics are calculated.

Demonstration of the model

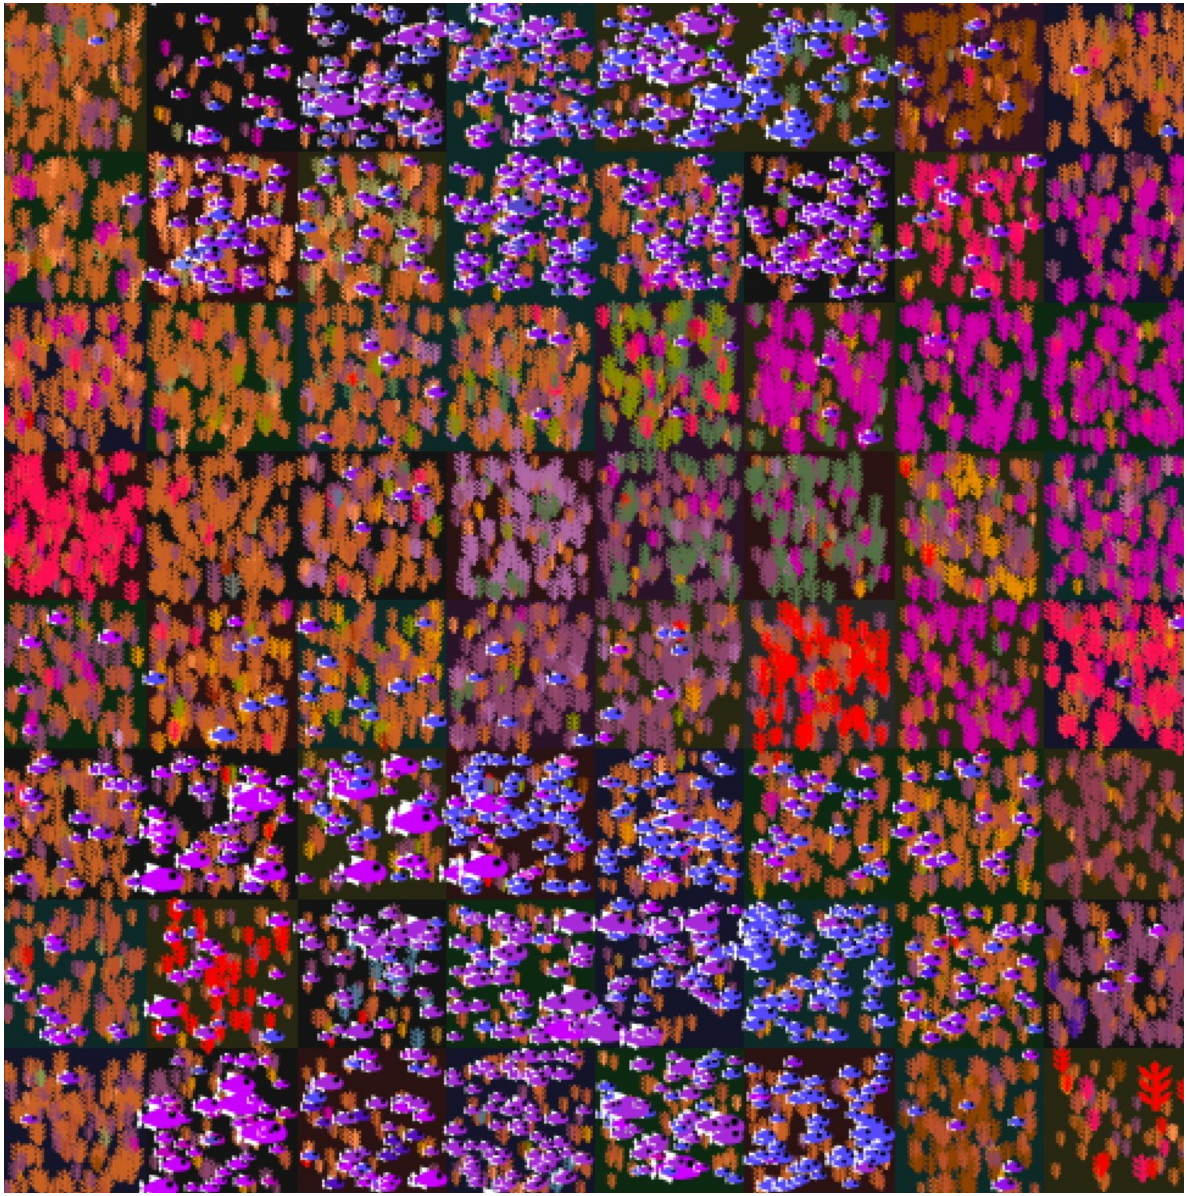

Evolving a moderately complex ecosystem starting point

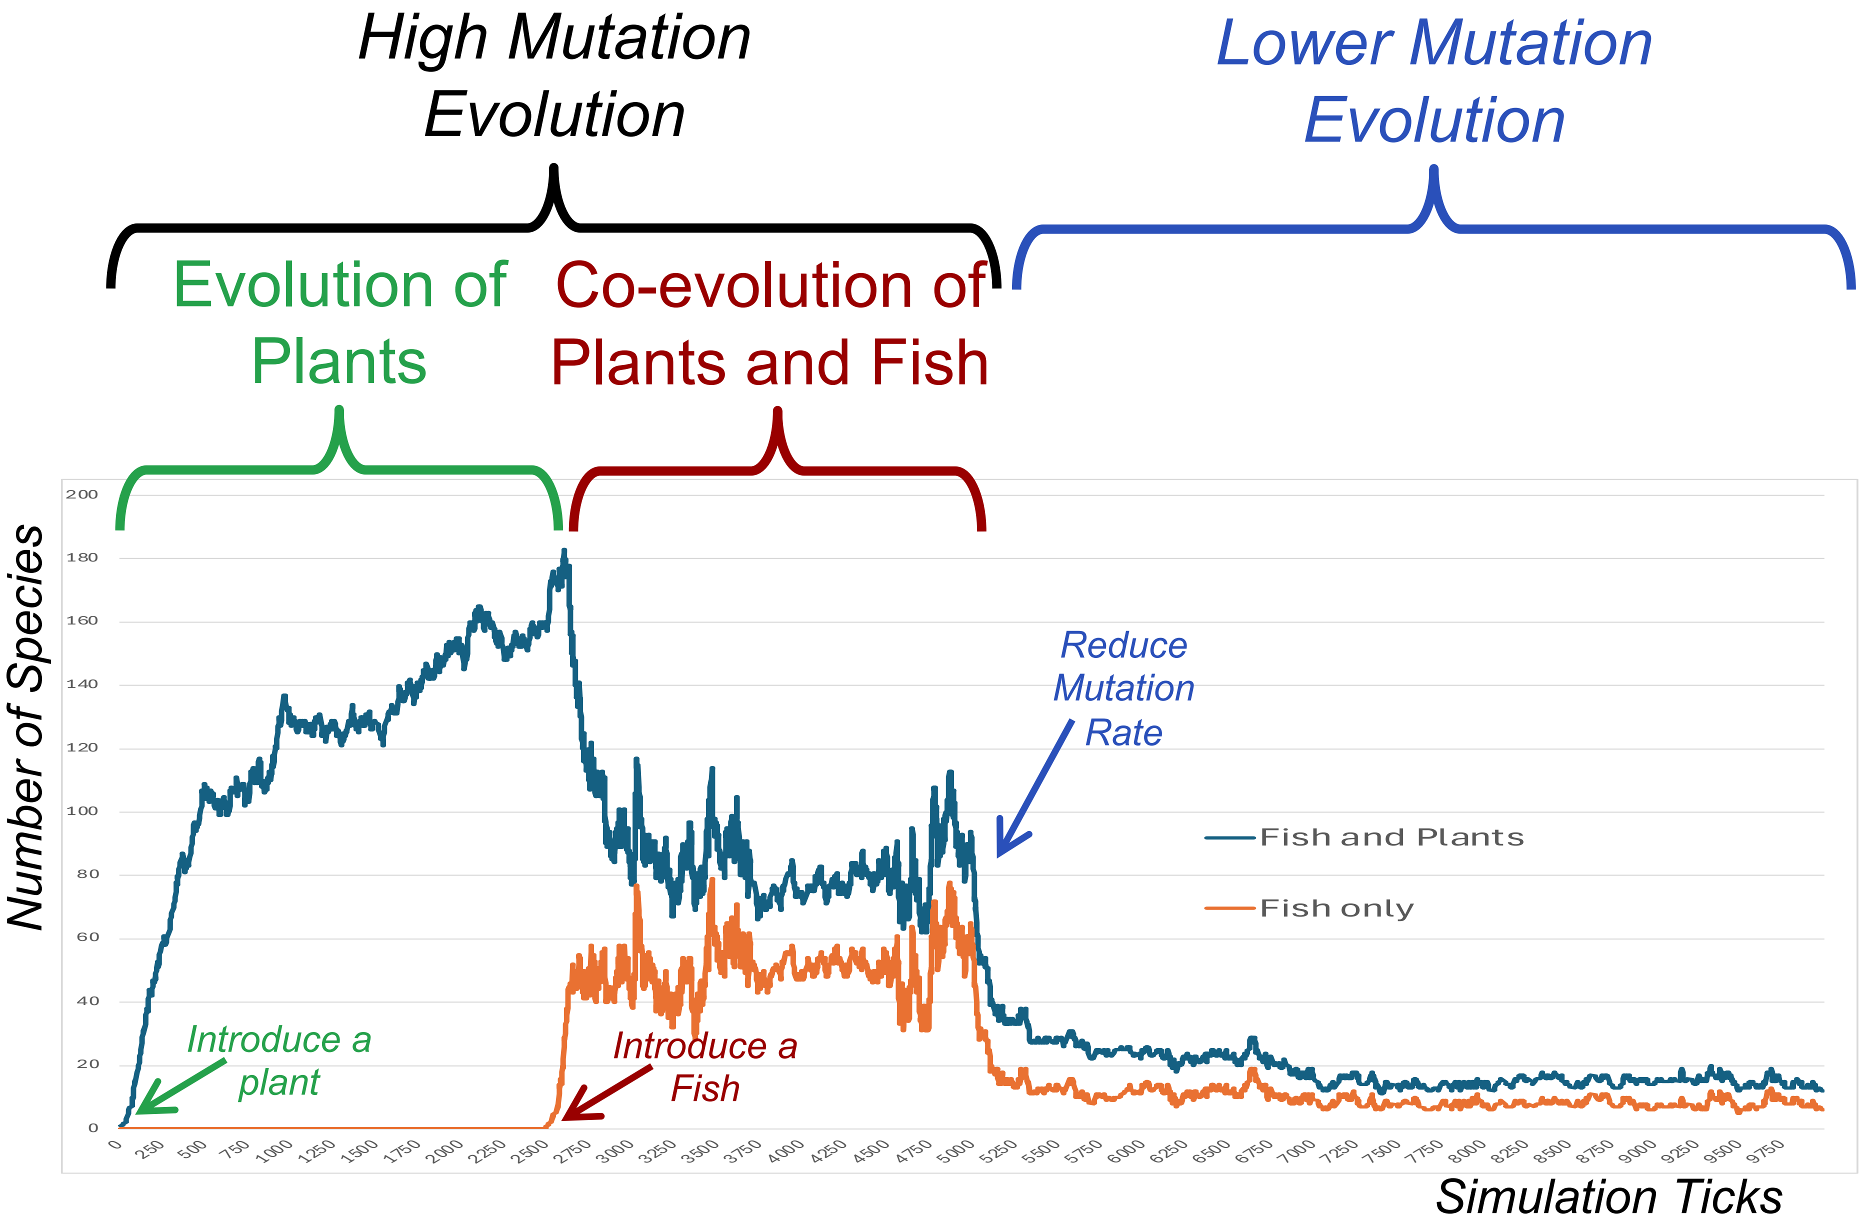

Then explore the system starting from there

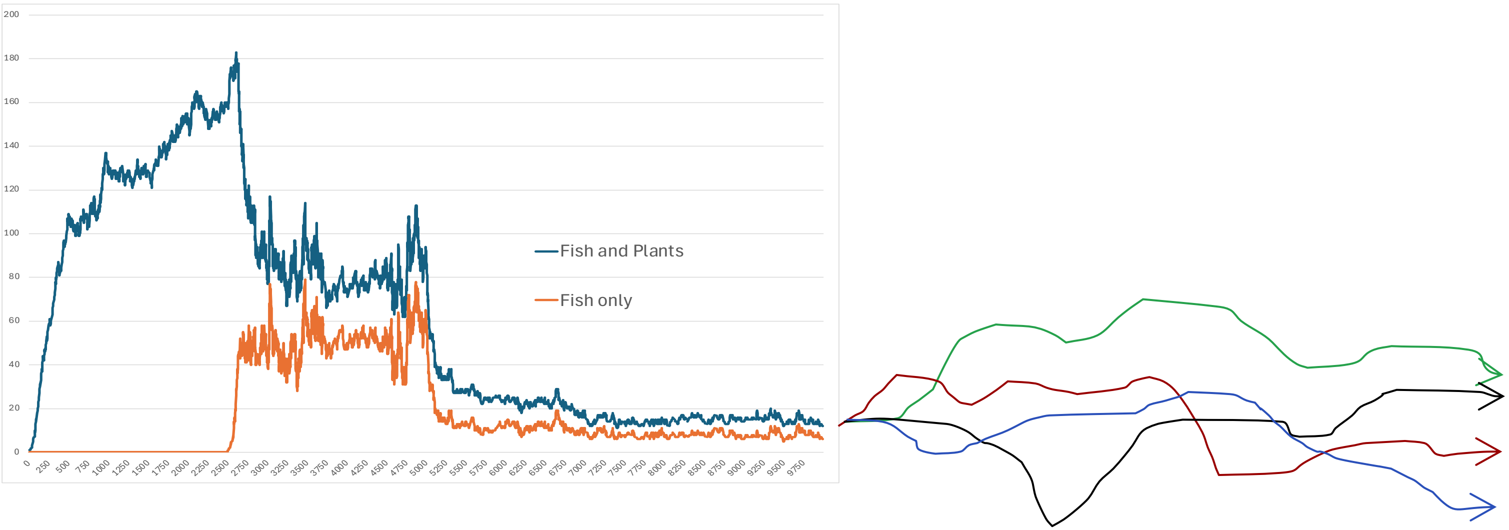

Evolve a complex ecology and save this state

Do multiple runs of the simulation starting from there for each condition to test

After, collect statistics or visualisations about what happened in the runs to understand the possible paths

Part 4

Exploring an abstract model of complex fishery ecosystems

– exploring the outcomes

Experiment 1 – a fishing 'shock'

- Starting from the same ecology…
- Remove a certain amount of fish for the first 100 simulation ticks
- Fish are removed with a uniform probability (a random fish from a random patch repeatedly)
- And run the simulation for another 900 ticks to see the impact of this
- Measure lots of stuff about each run

Fish numbs (within run averages) – no fishing

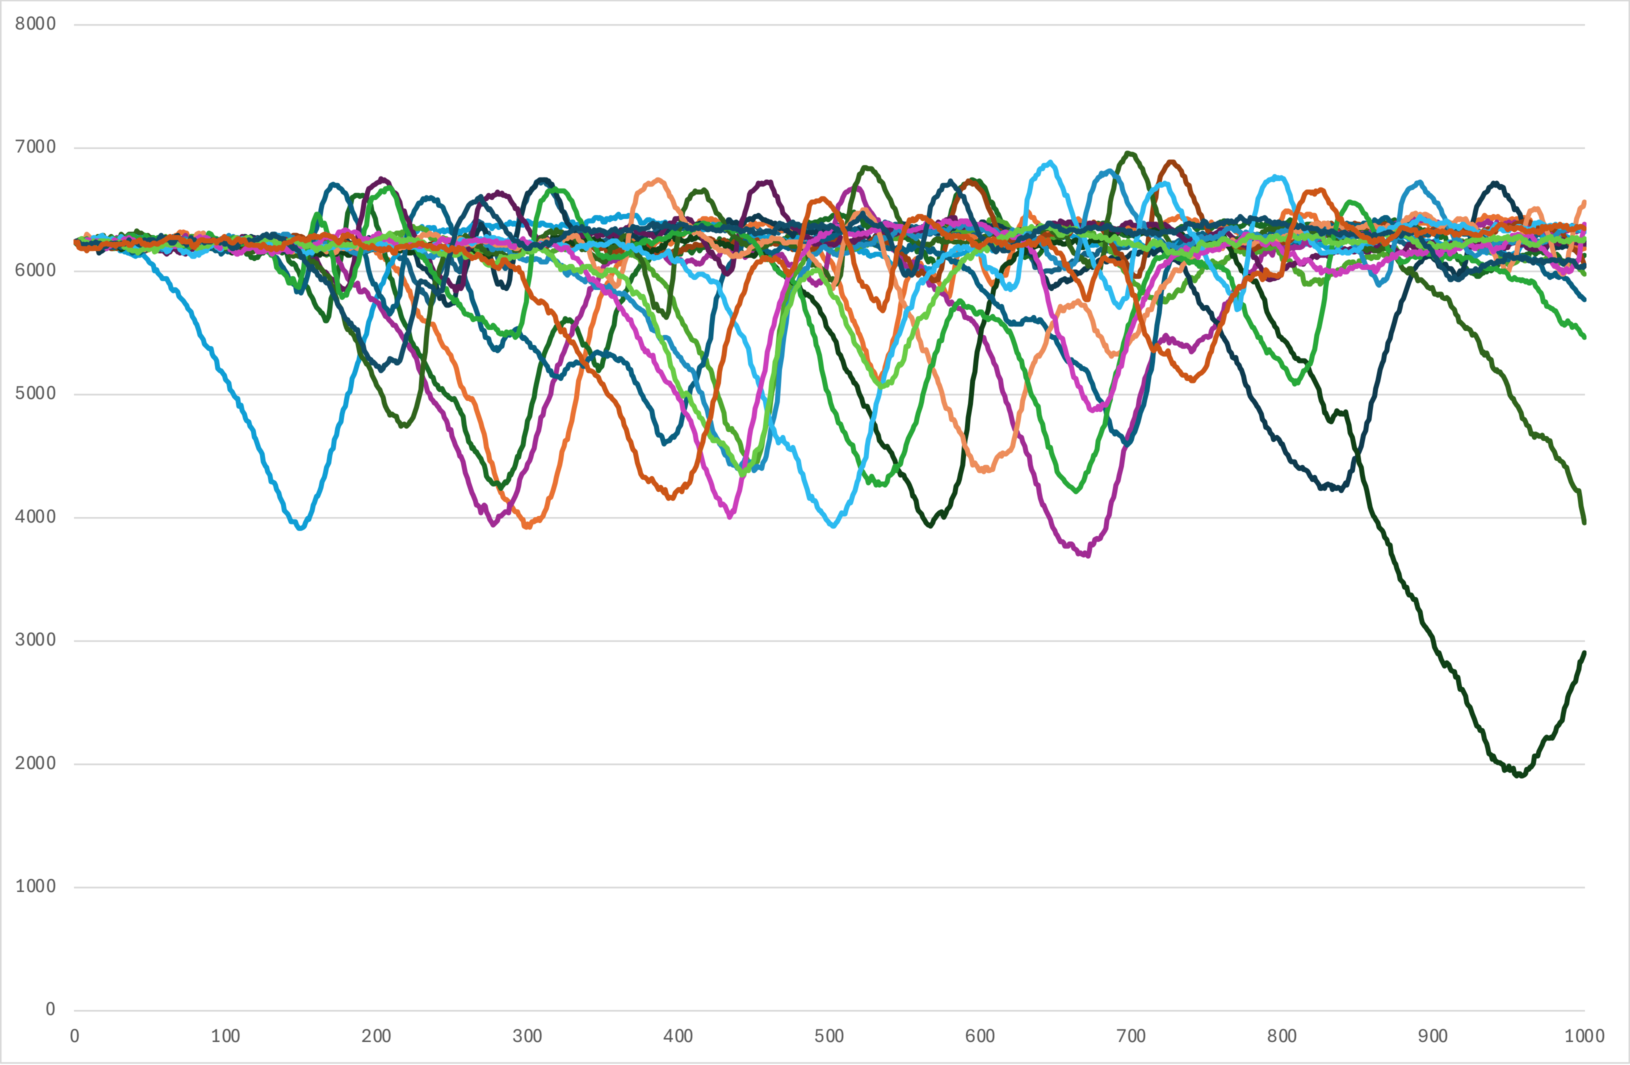

Fish numbs (within run averages) – catch level 70/tick

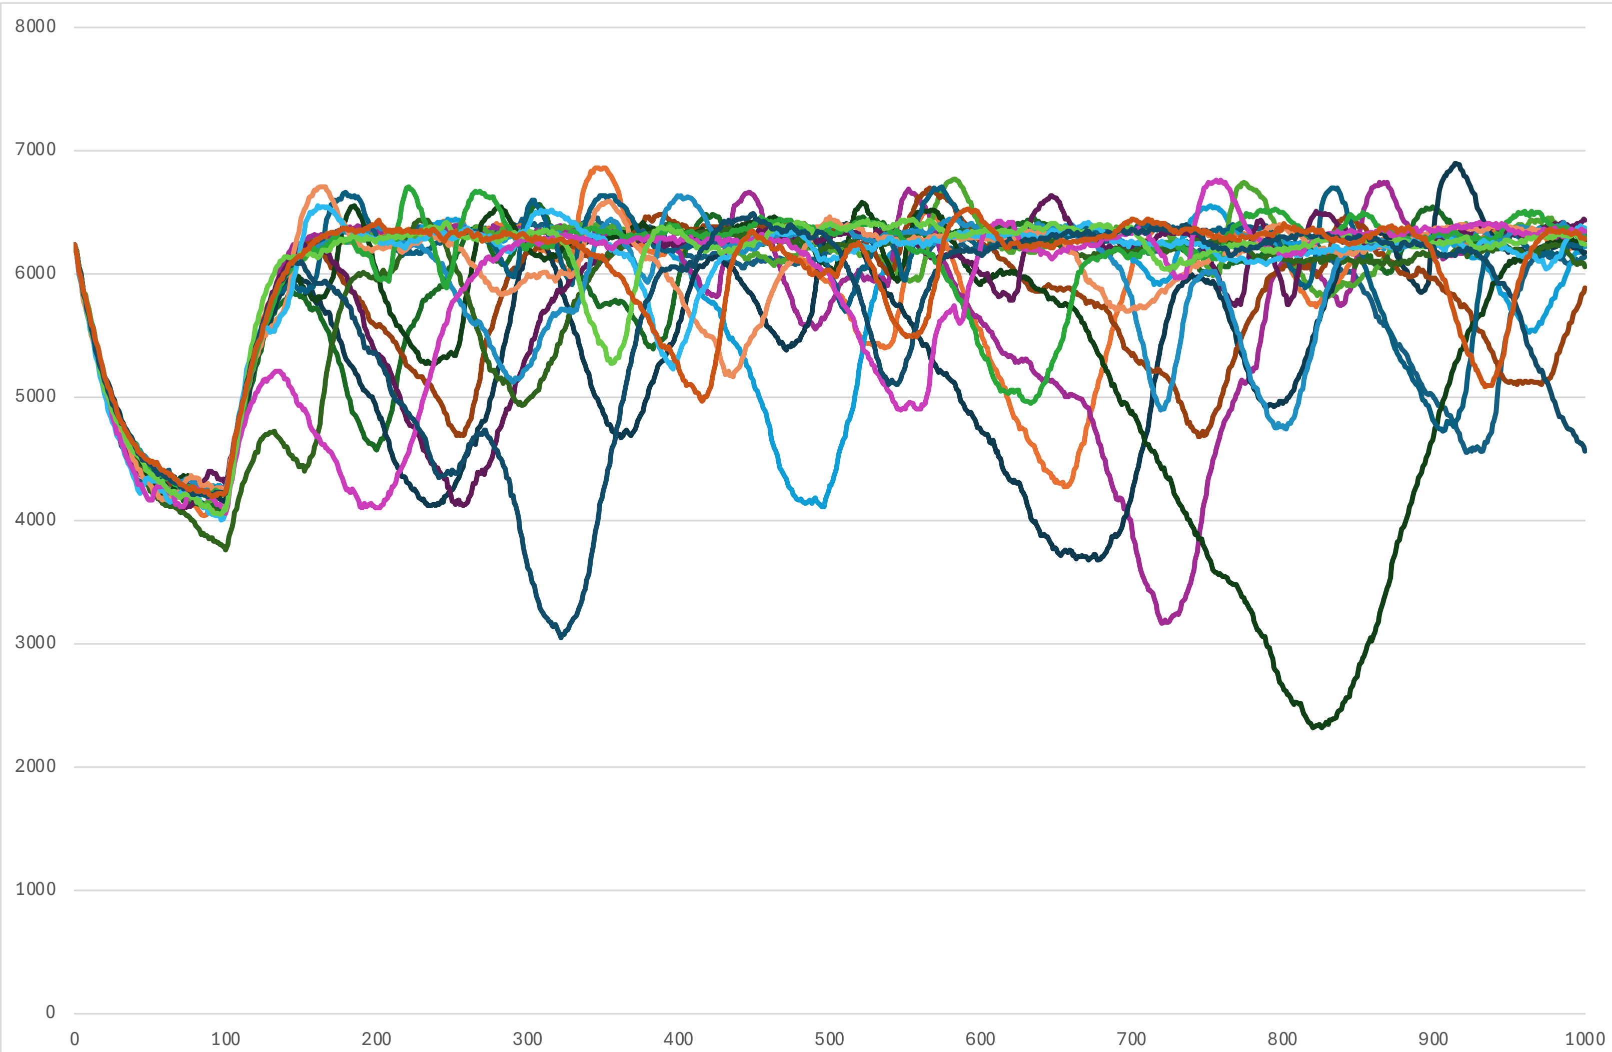

Fish numbs (within run averages) – catch level 140/tick

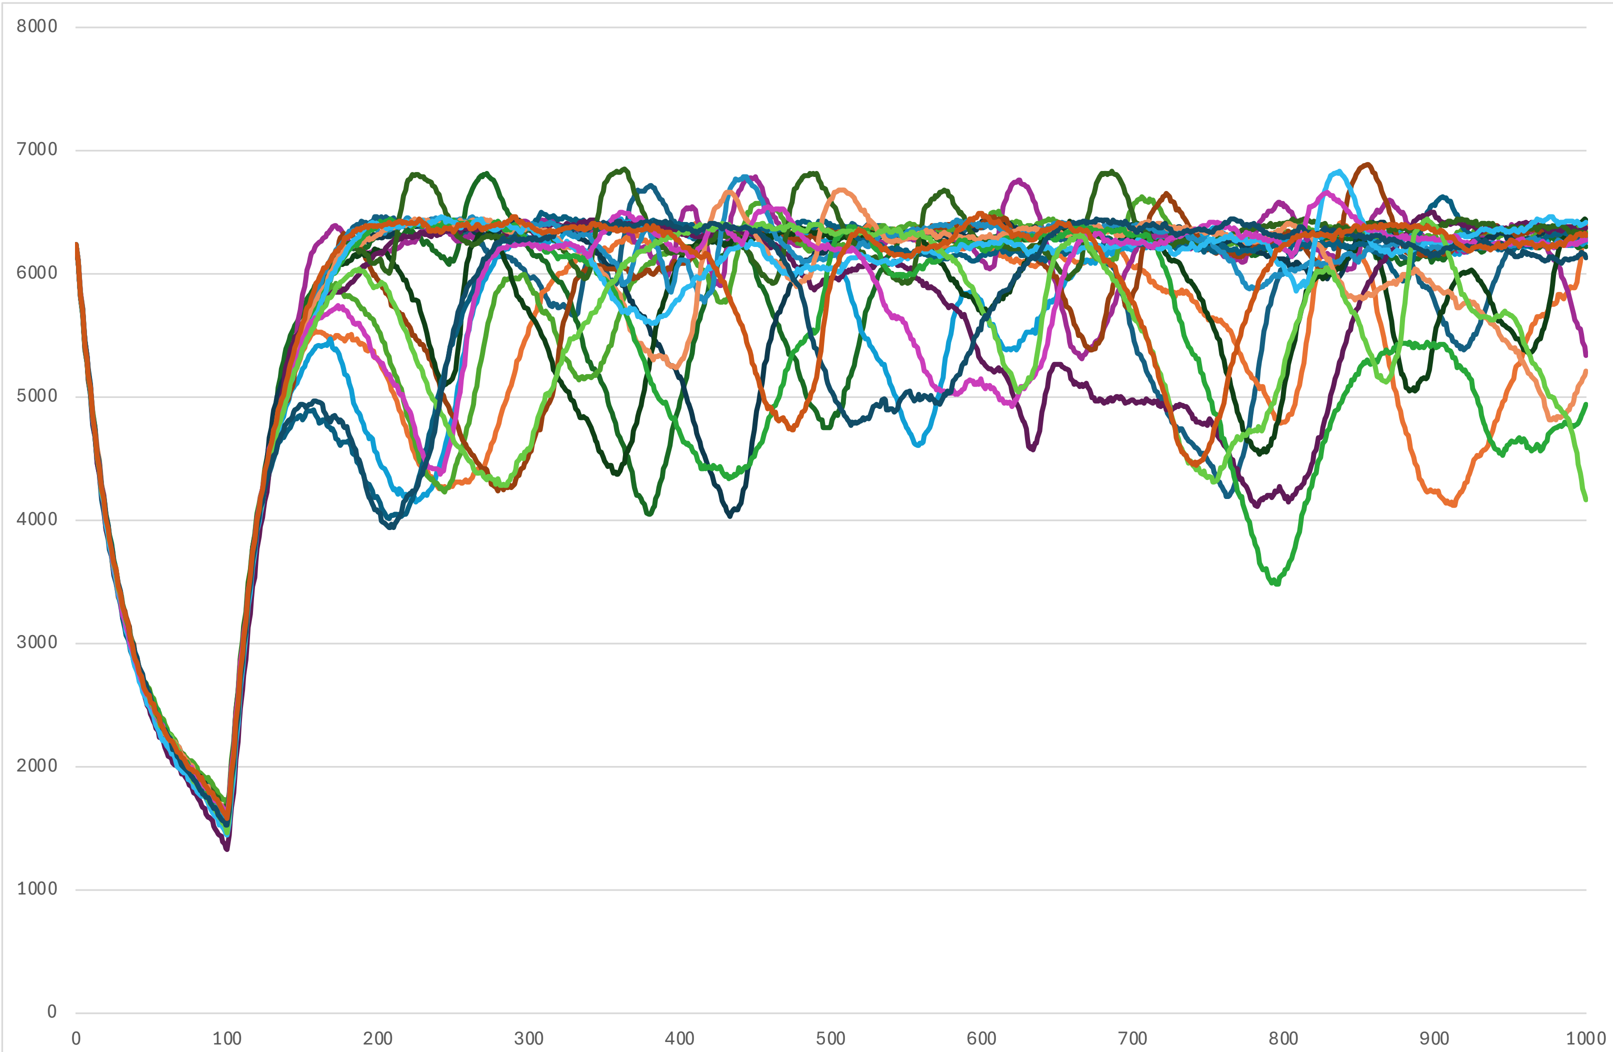

Fish numbs (within run averages) – catch level 150/tick

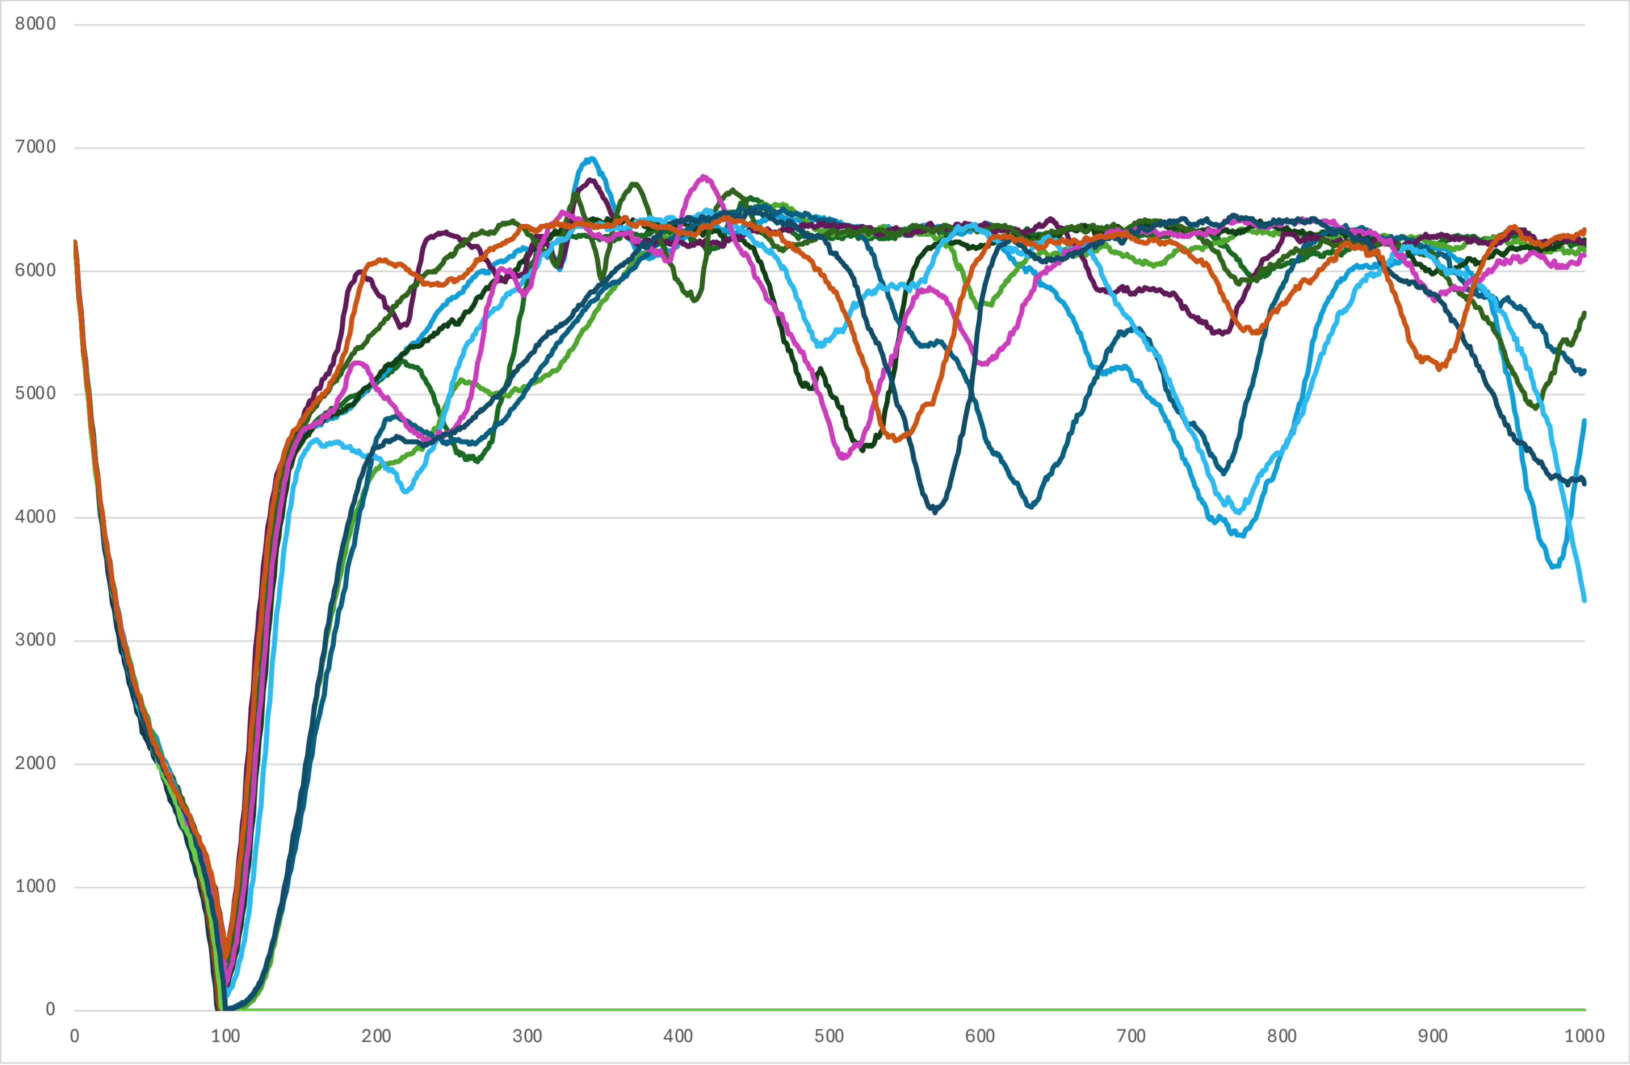

Fish numbs (within run averages) – catch level 160/tick

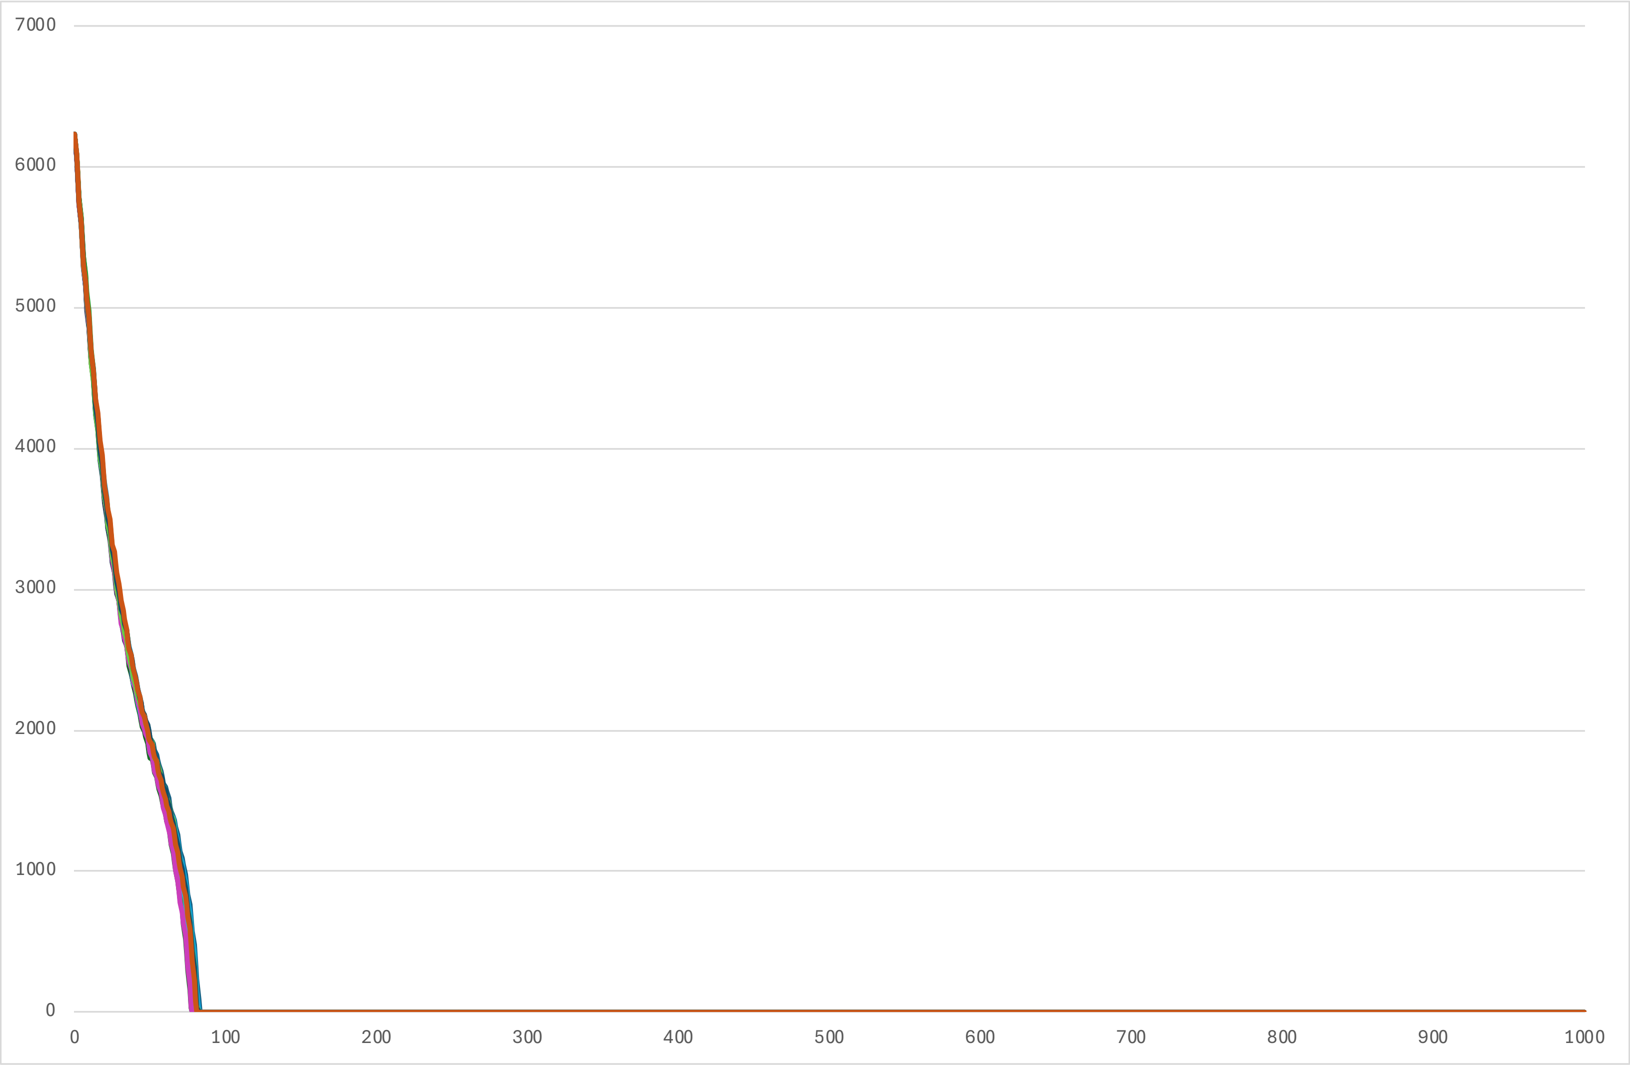

Fish numbers after shock (averaged over the average from each of the 20 runs) different levels of catch

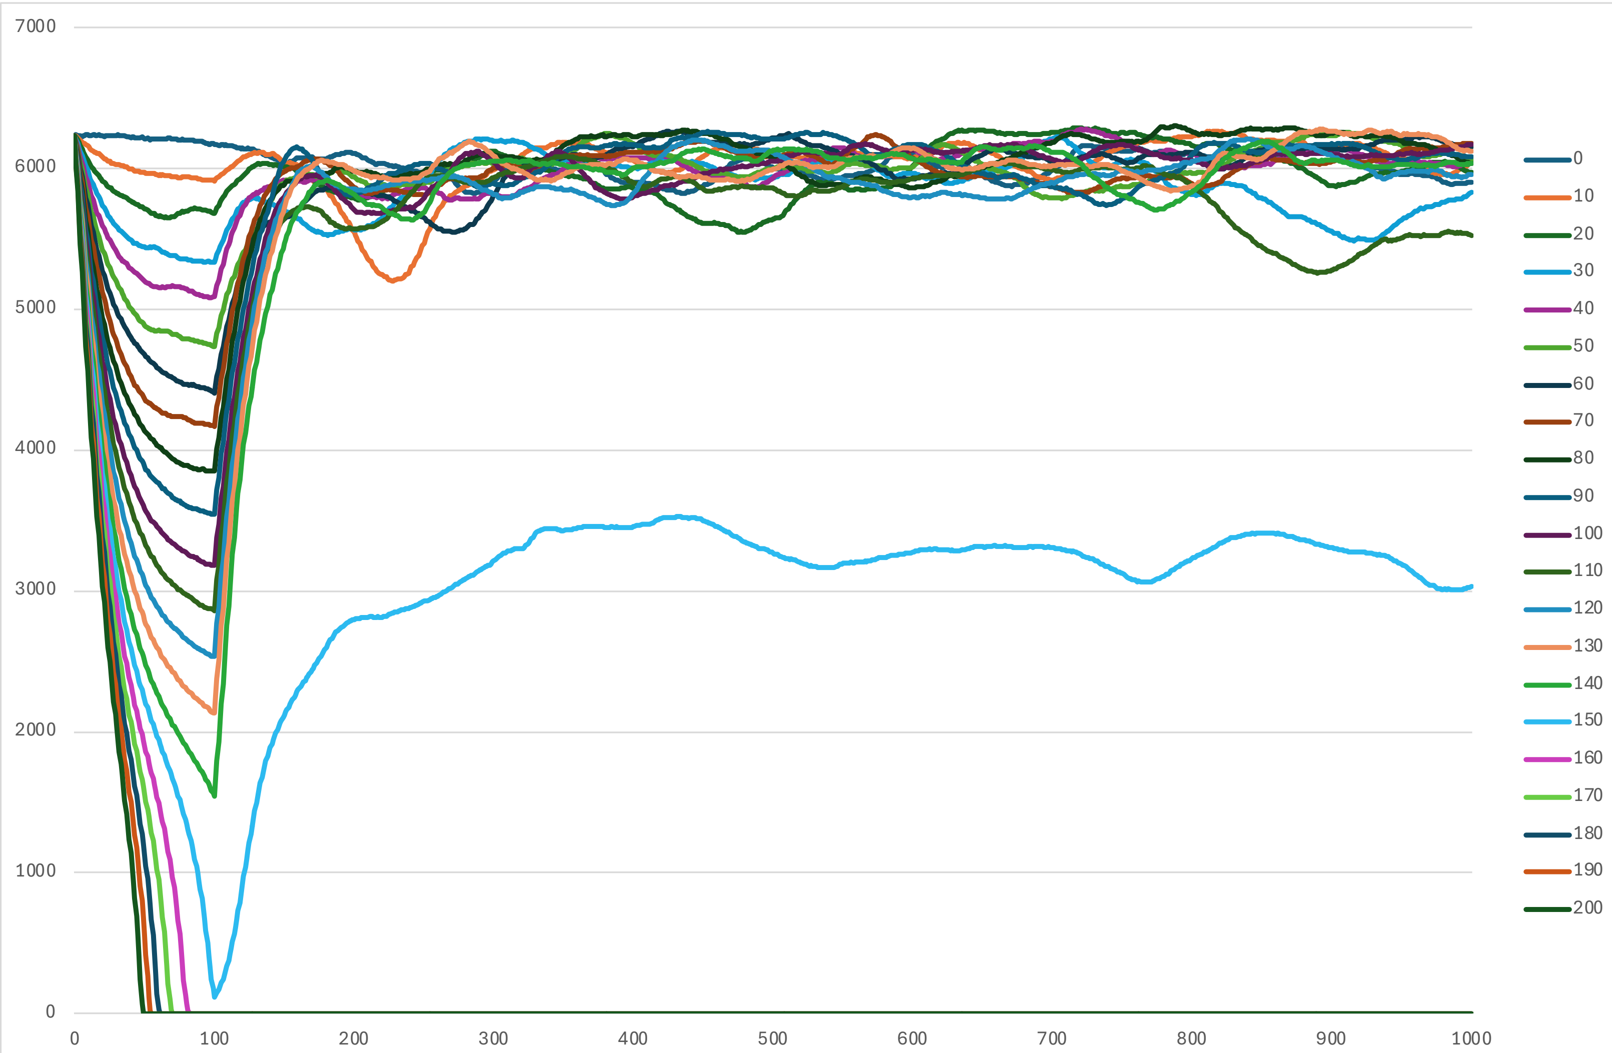

If we tried this (catch level 140/tick for 100 ticks**) what would we measure? What would we conclude?**

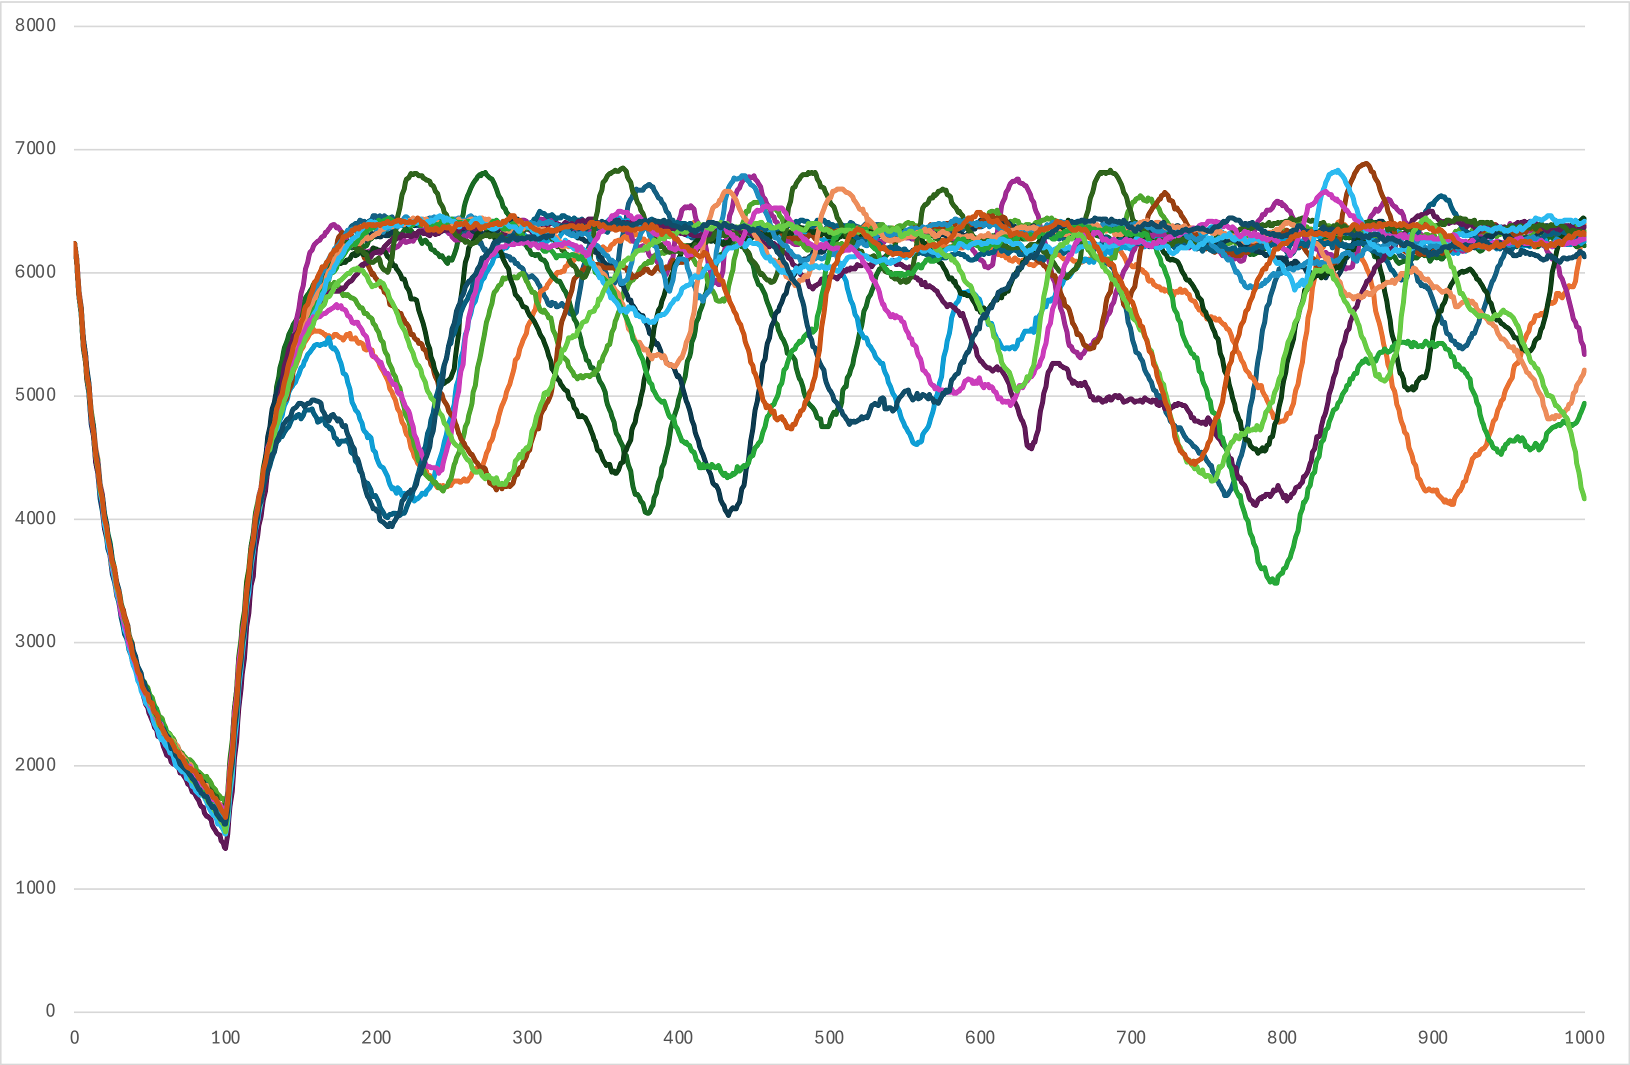

Total Fish Harvested vs Extinction Risk of shock trial

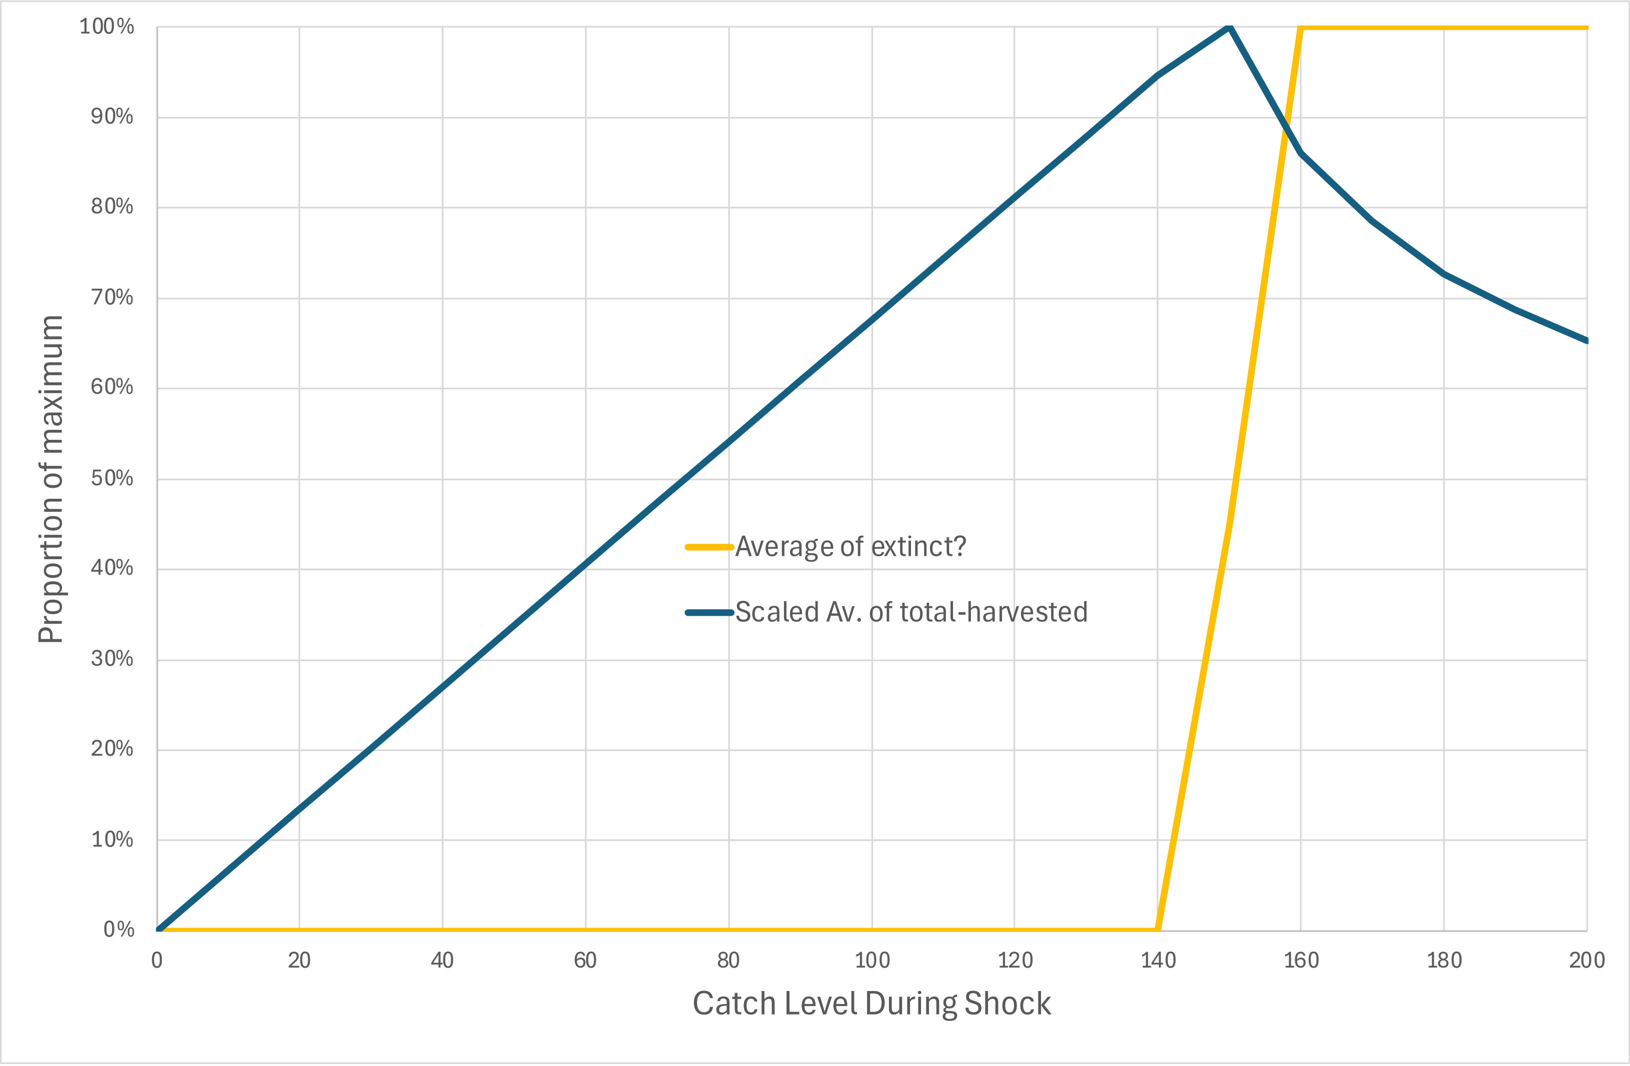

Experiment 2 – long-term fishing levels

- Starting from the same ecology…
- Remove a certain amount of fish every simulation tick until end of run
- Fish are removed with a uniform probability (a random fish from a random patch repeatedly)
- Measure lots of stuff about each run

Total Fish Harvested from long-term fishing (scaled so max=1)

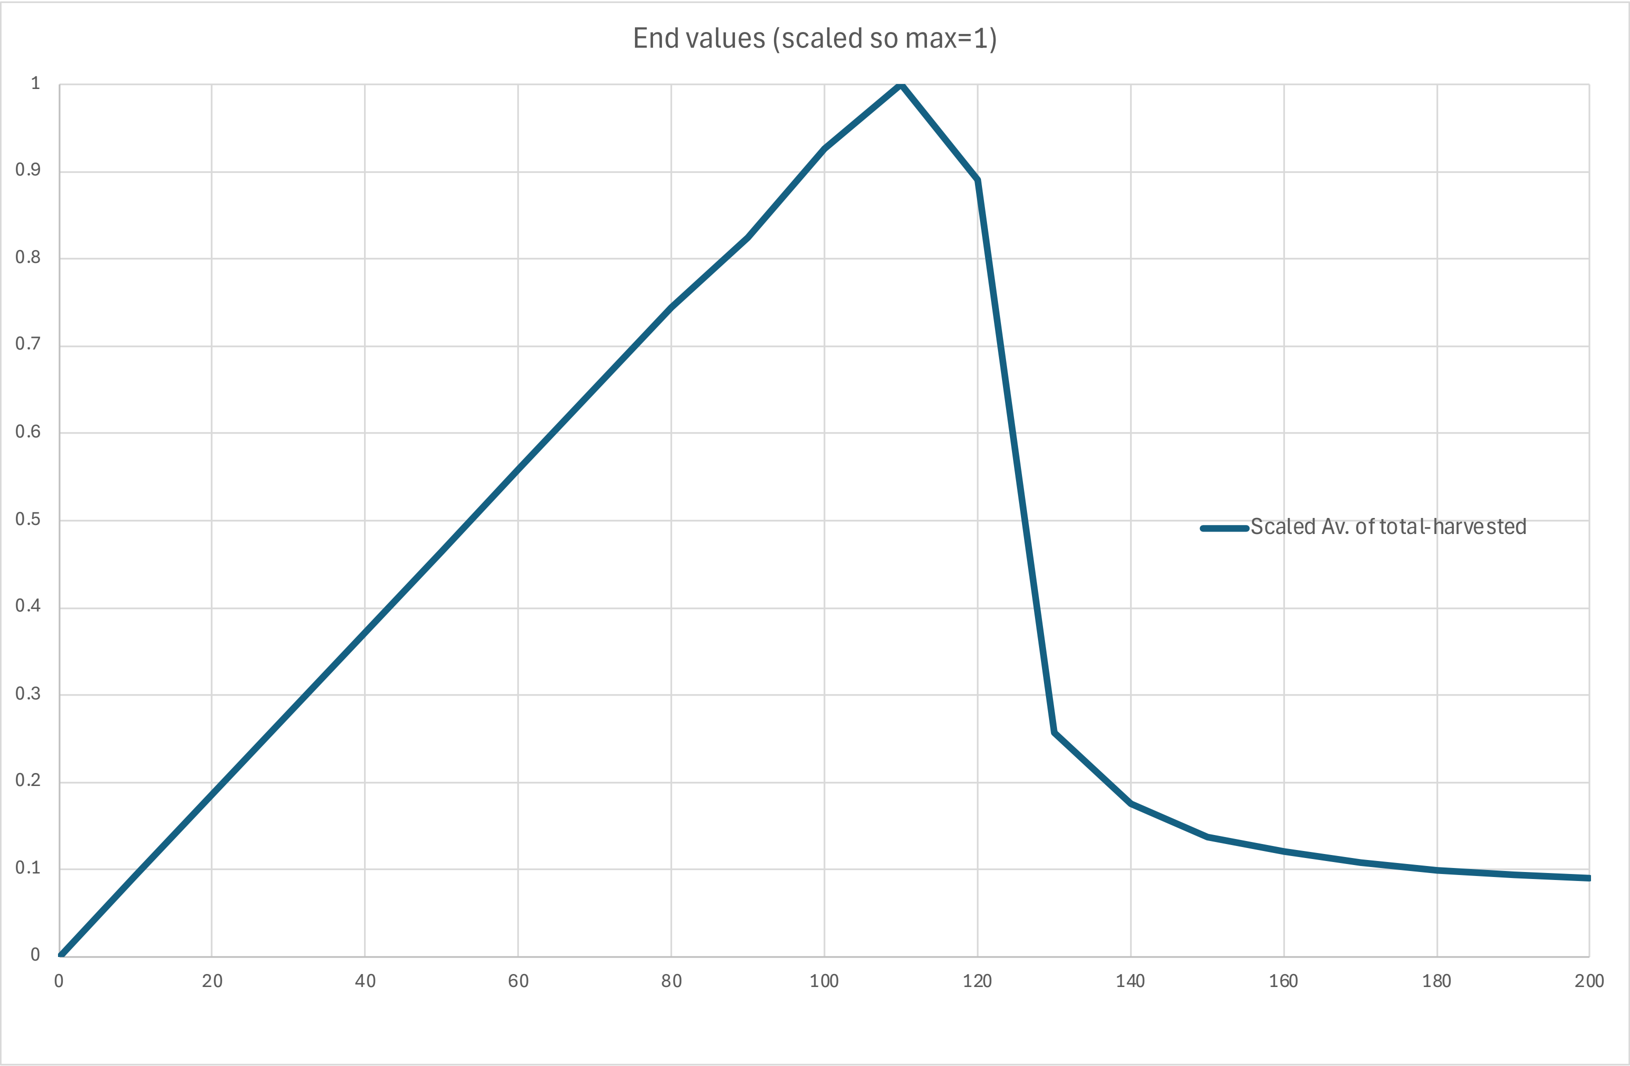

Fish numbers (av. over each run) Catch level = 0

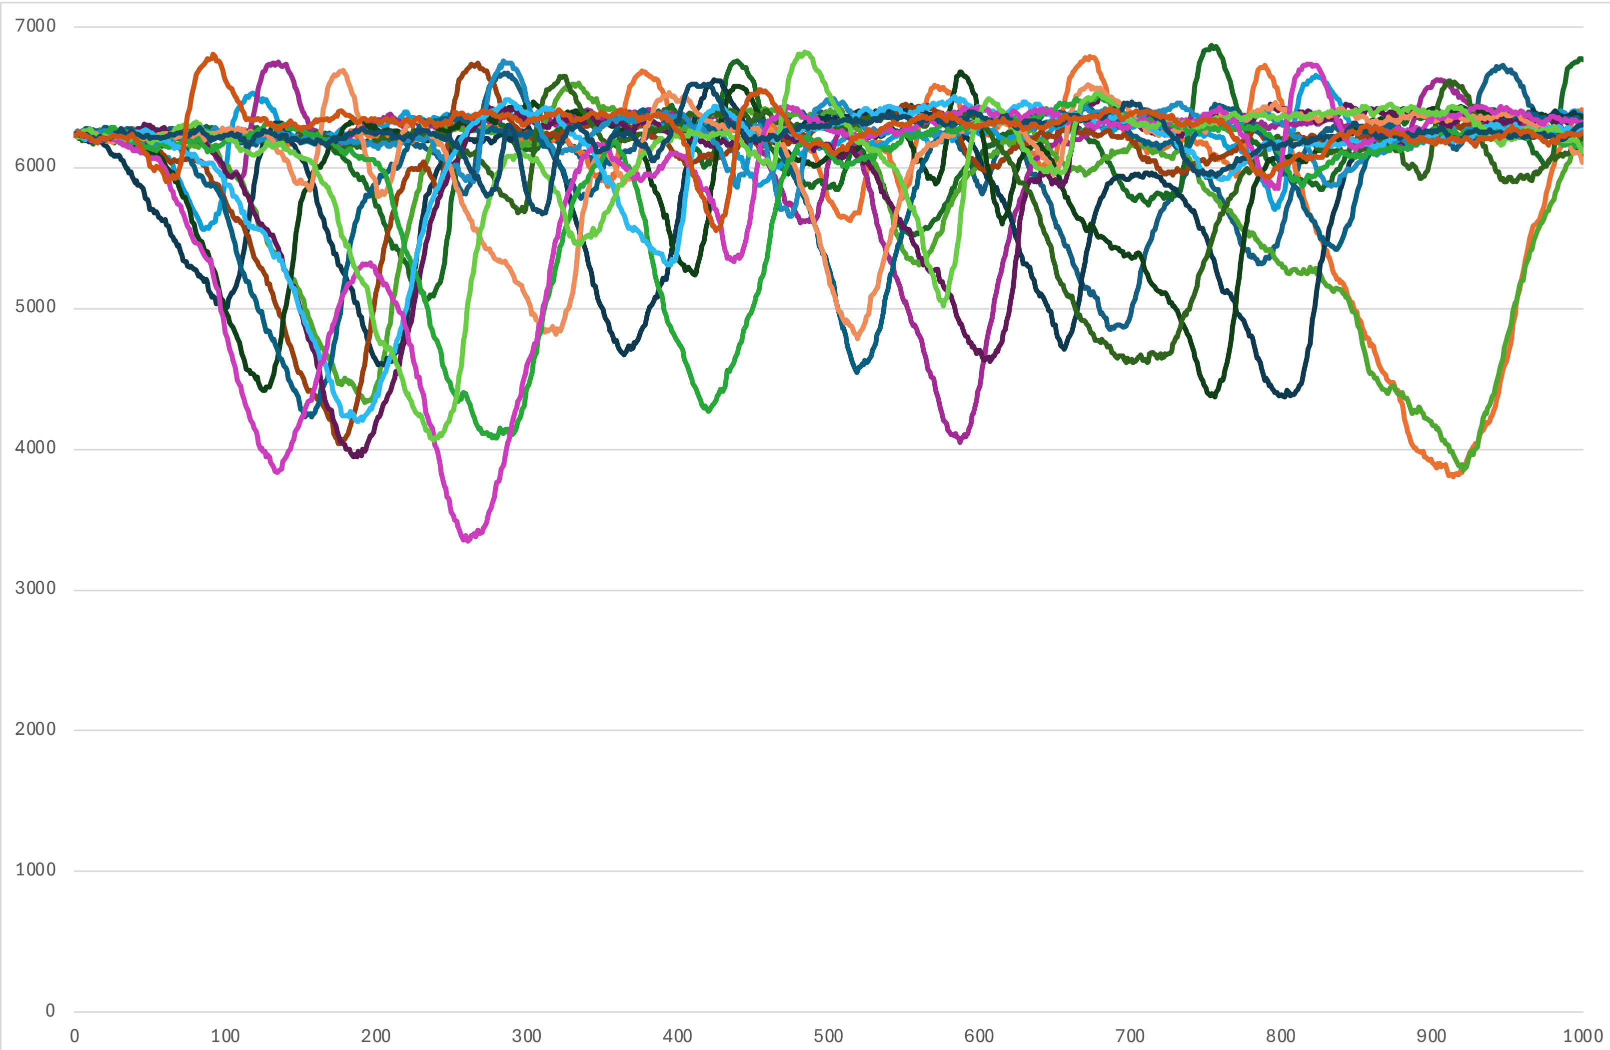

Fish numbers (av. over each run) Catch level = 20

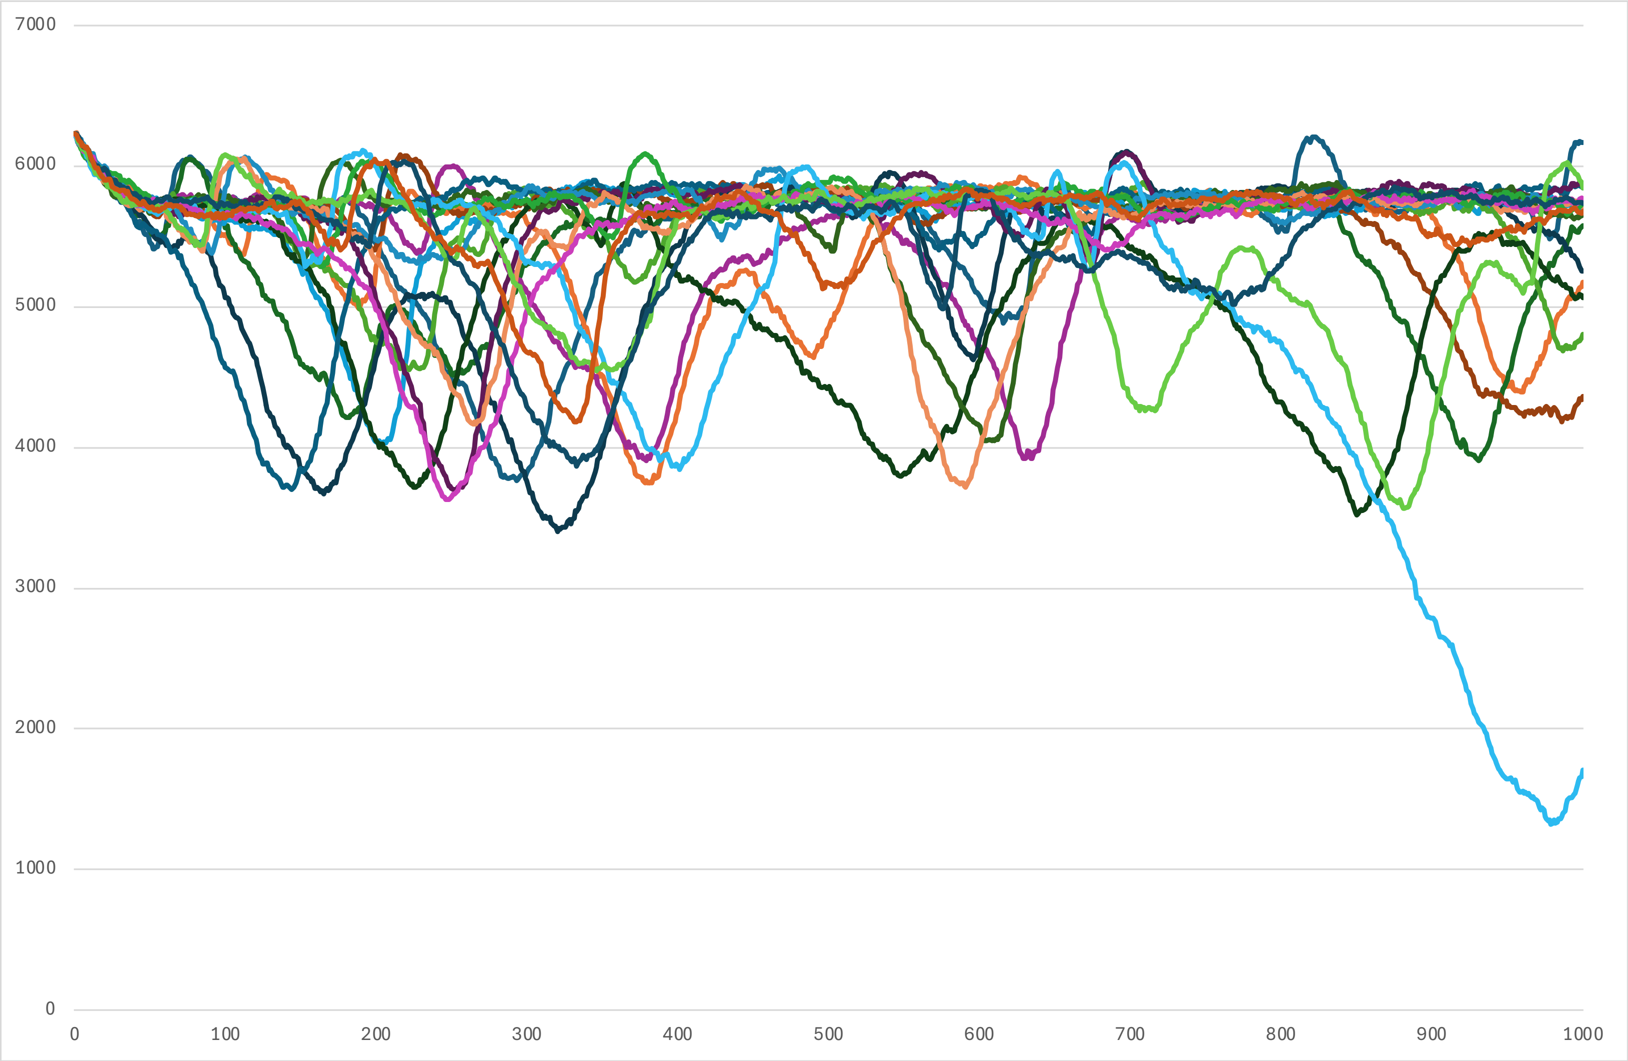

Fish numbers (av. over each run) Catch level = 40

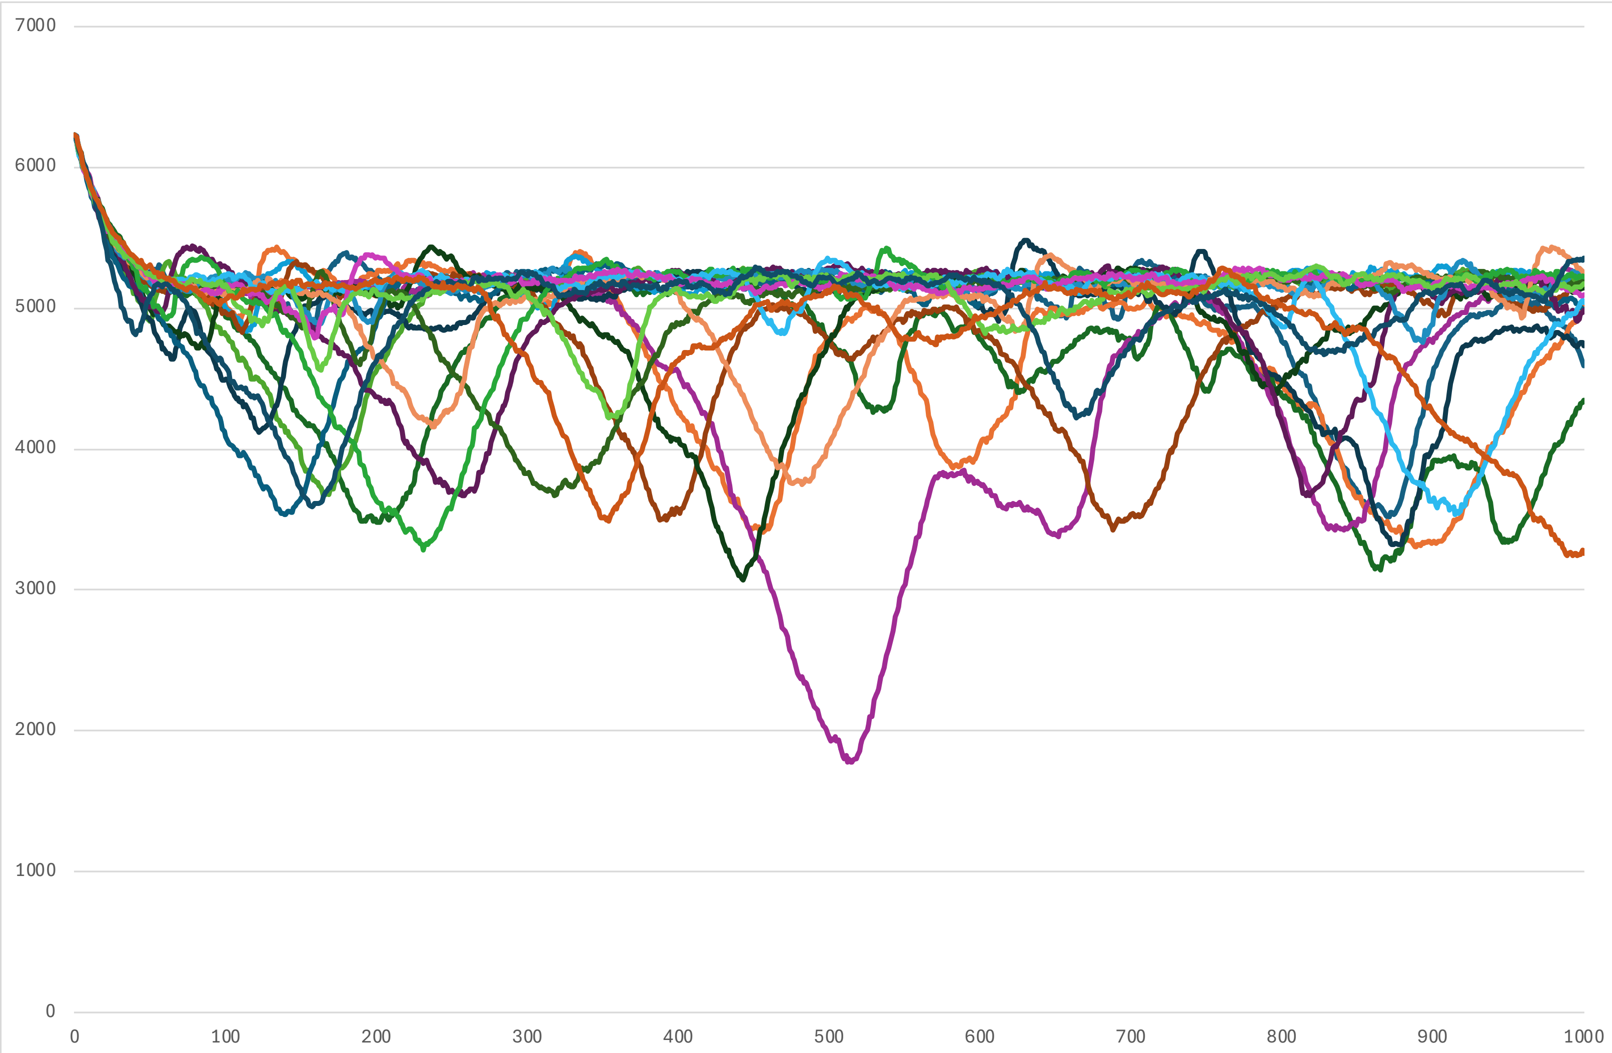

Fish numbers (av. over each run) Catch level = 60

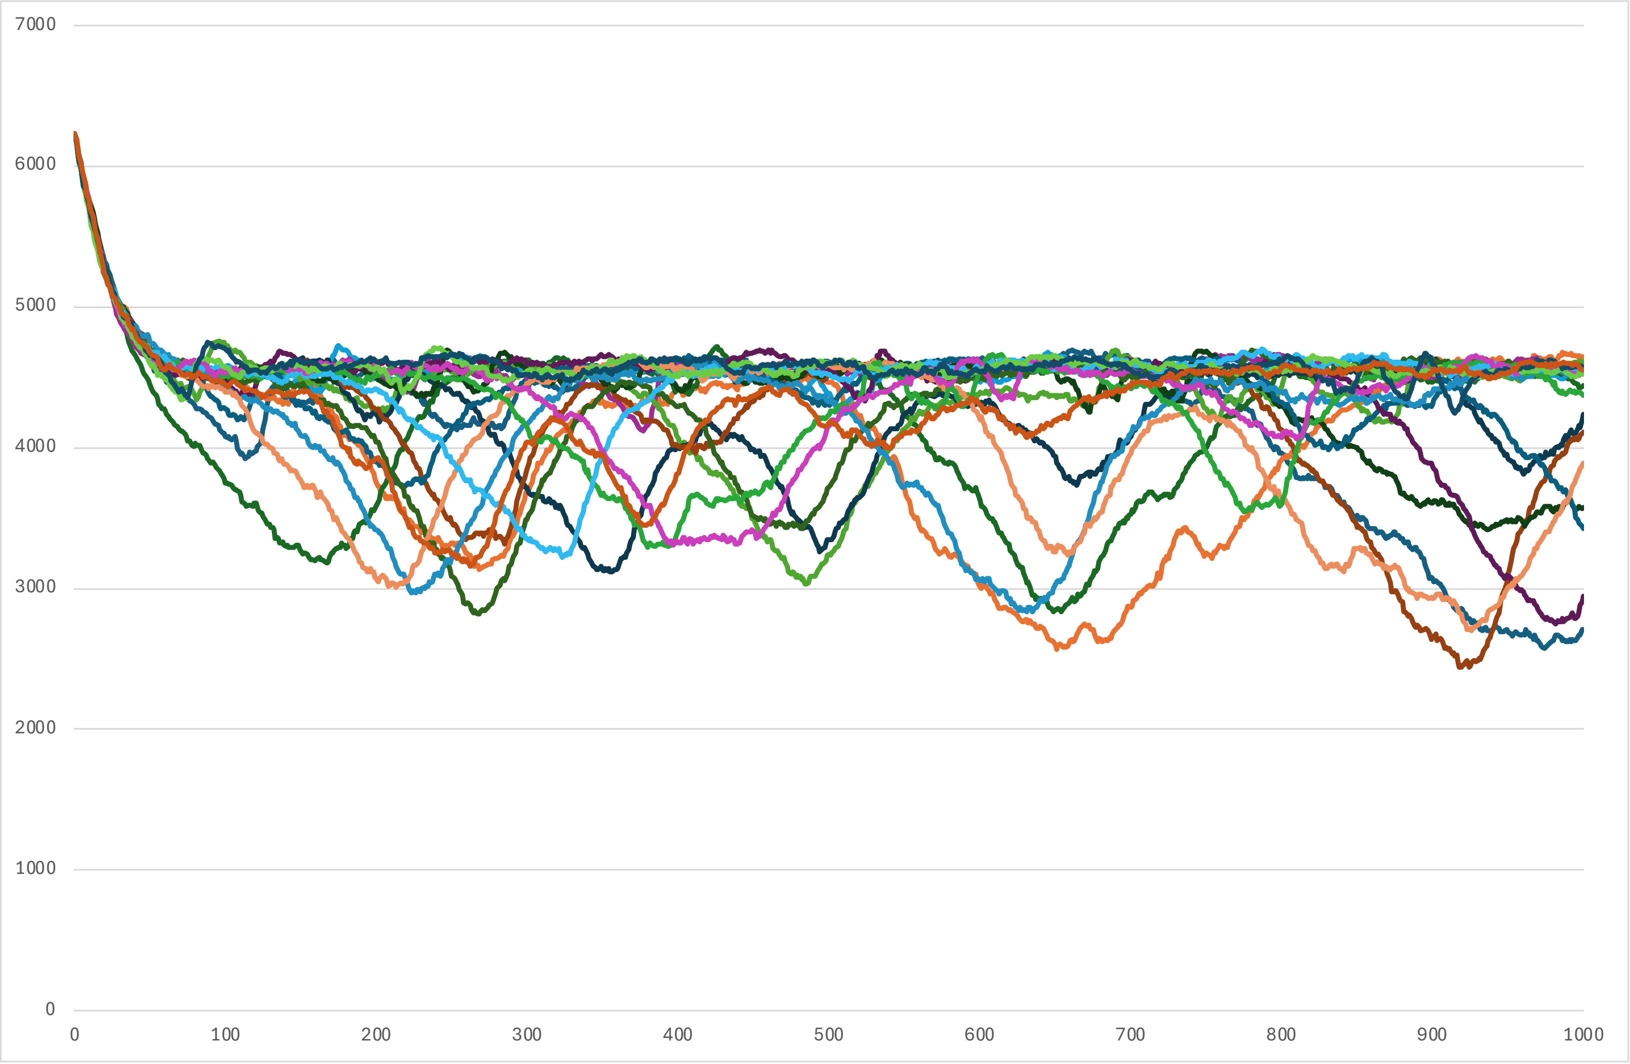

Fish numbers (av. over each run) Catch level = 80

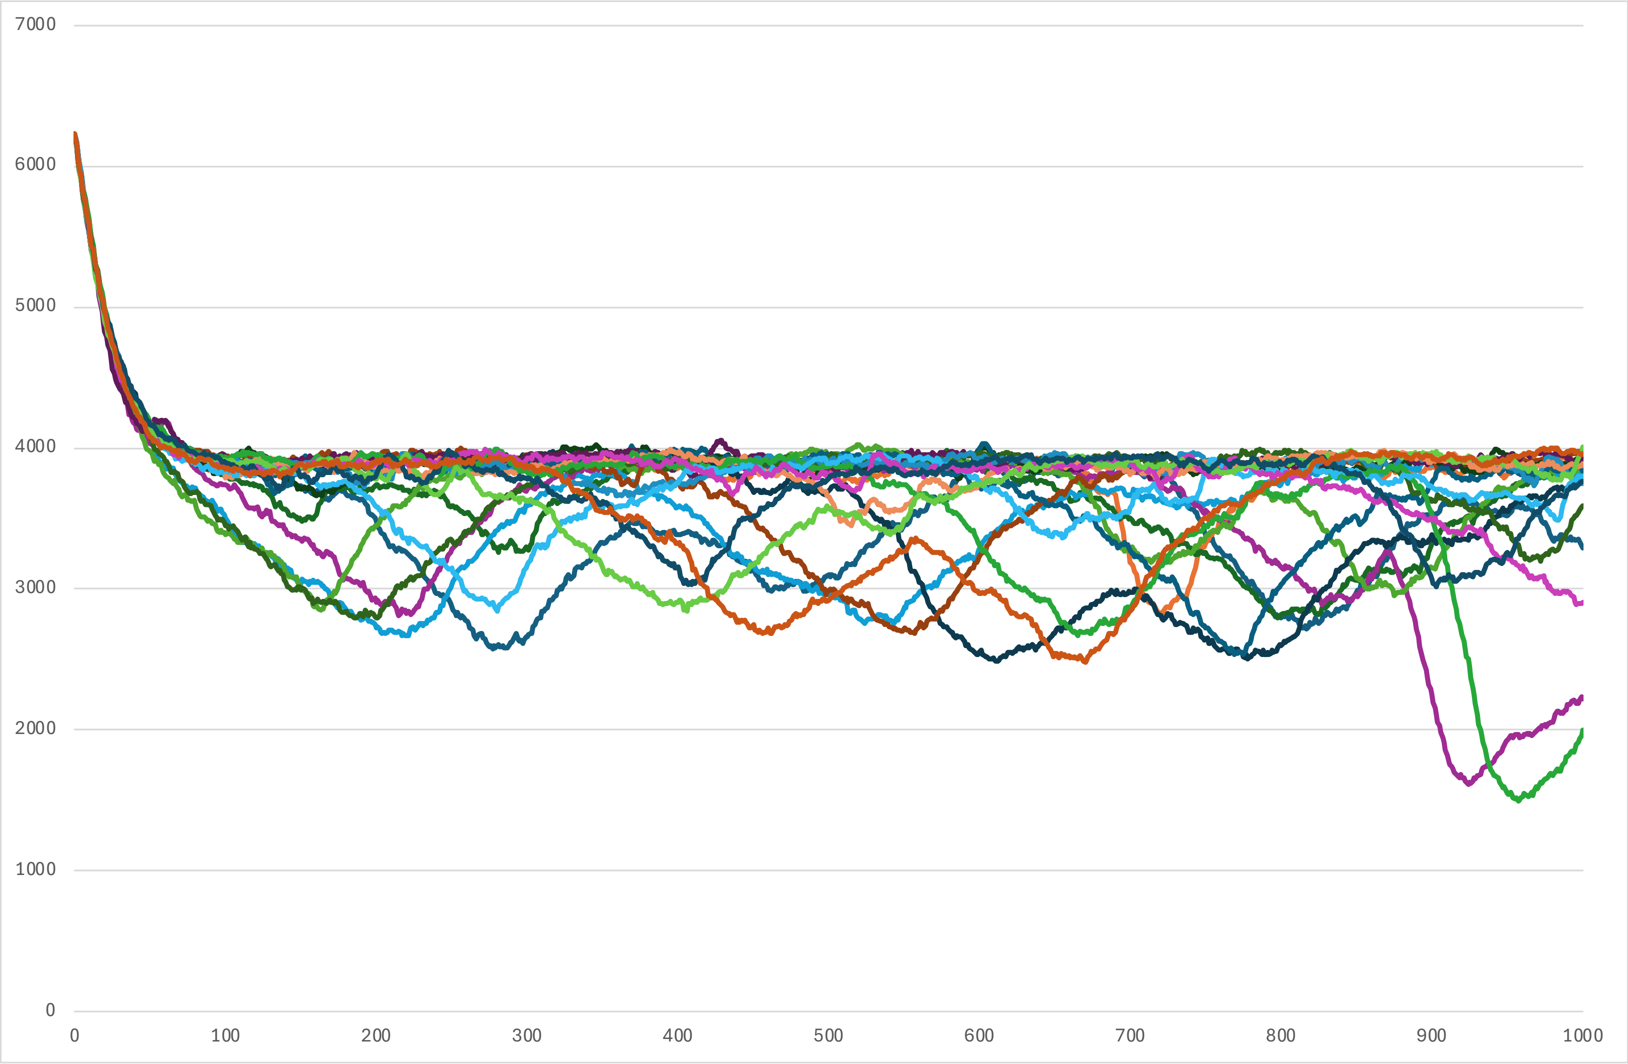

Fish numbers (av. over each run) Catch level = 90

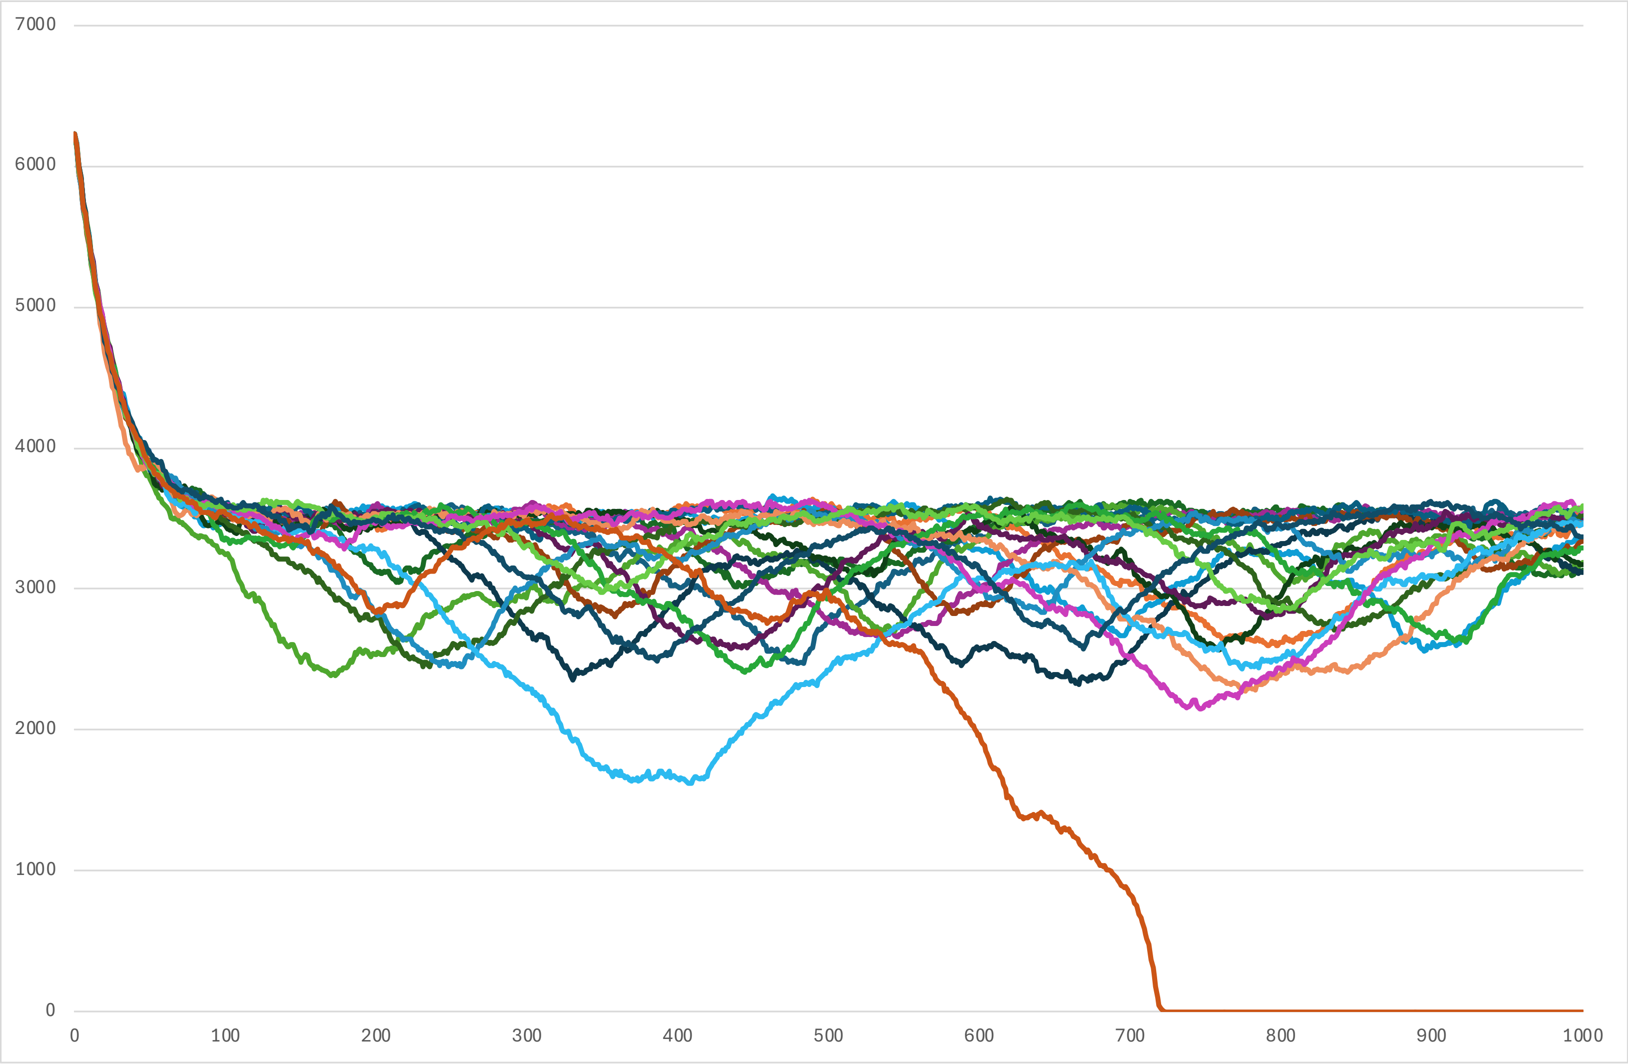

Fish numbers (av. over each run) Catch level = 100

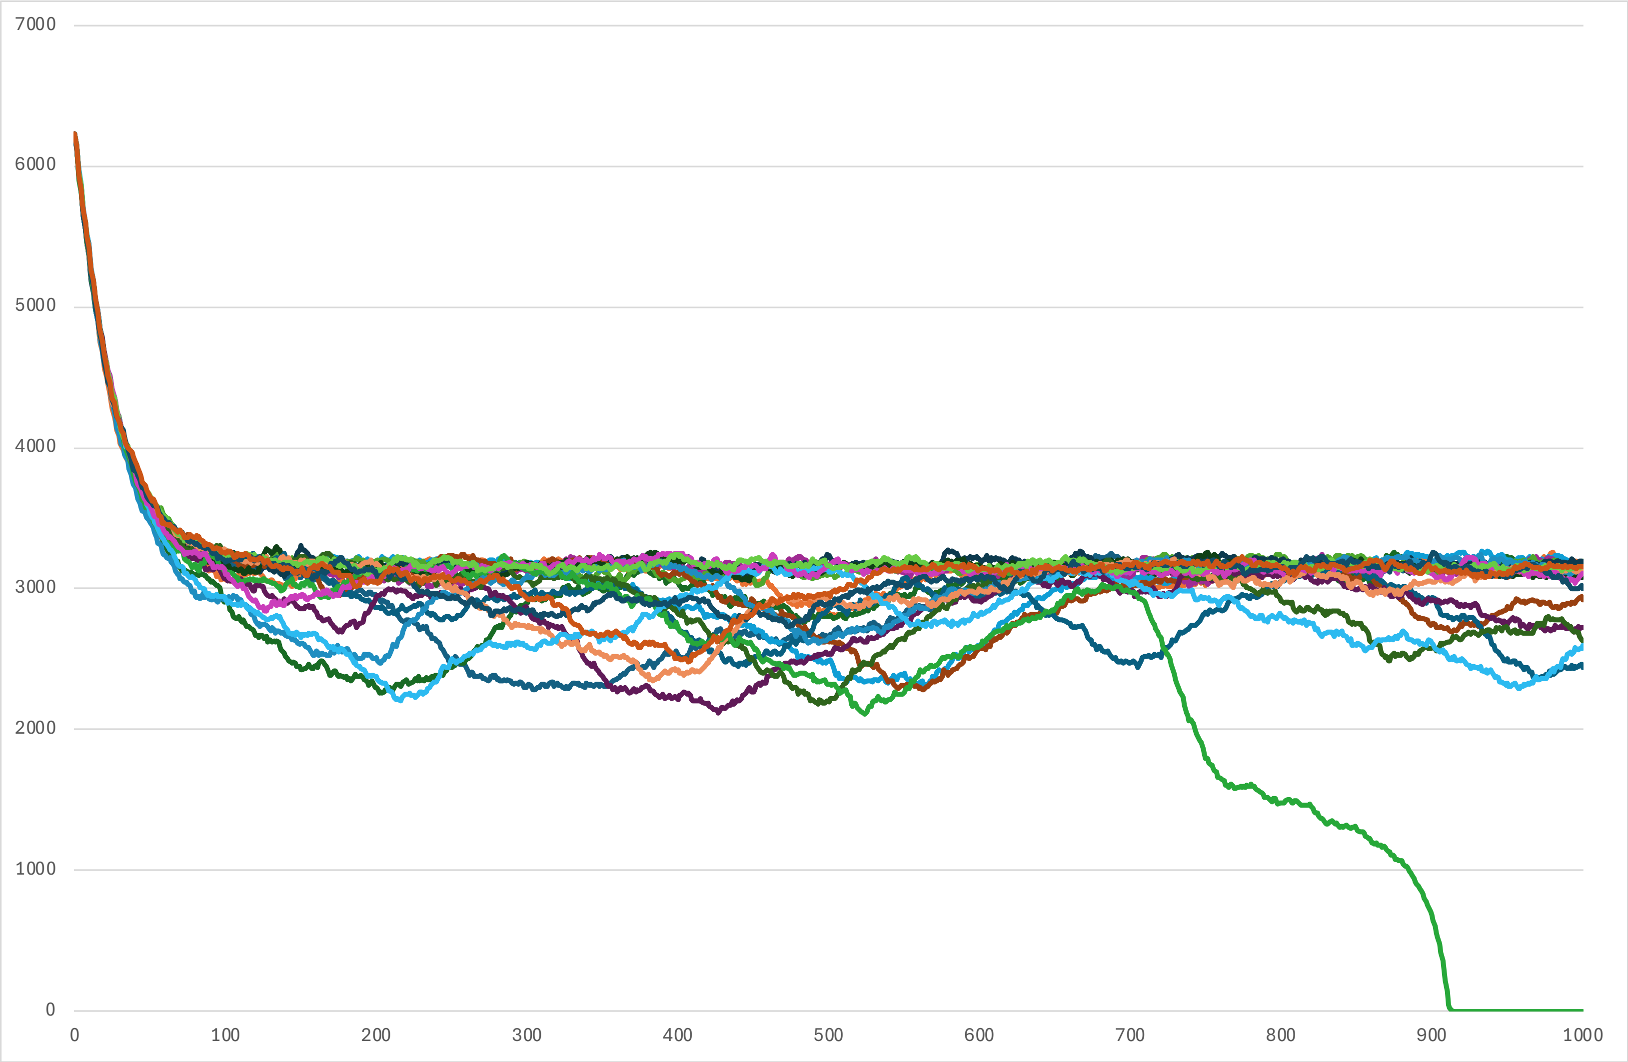

Fish numbers (av. over each run) Catch level = 110

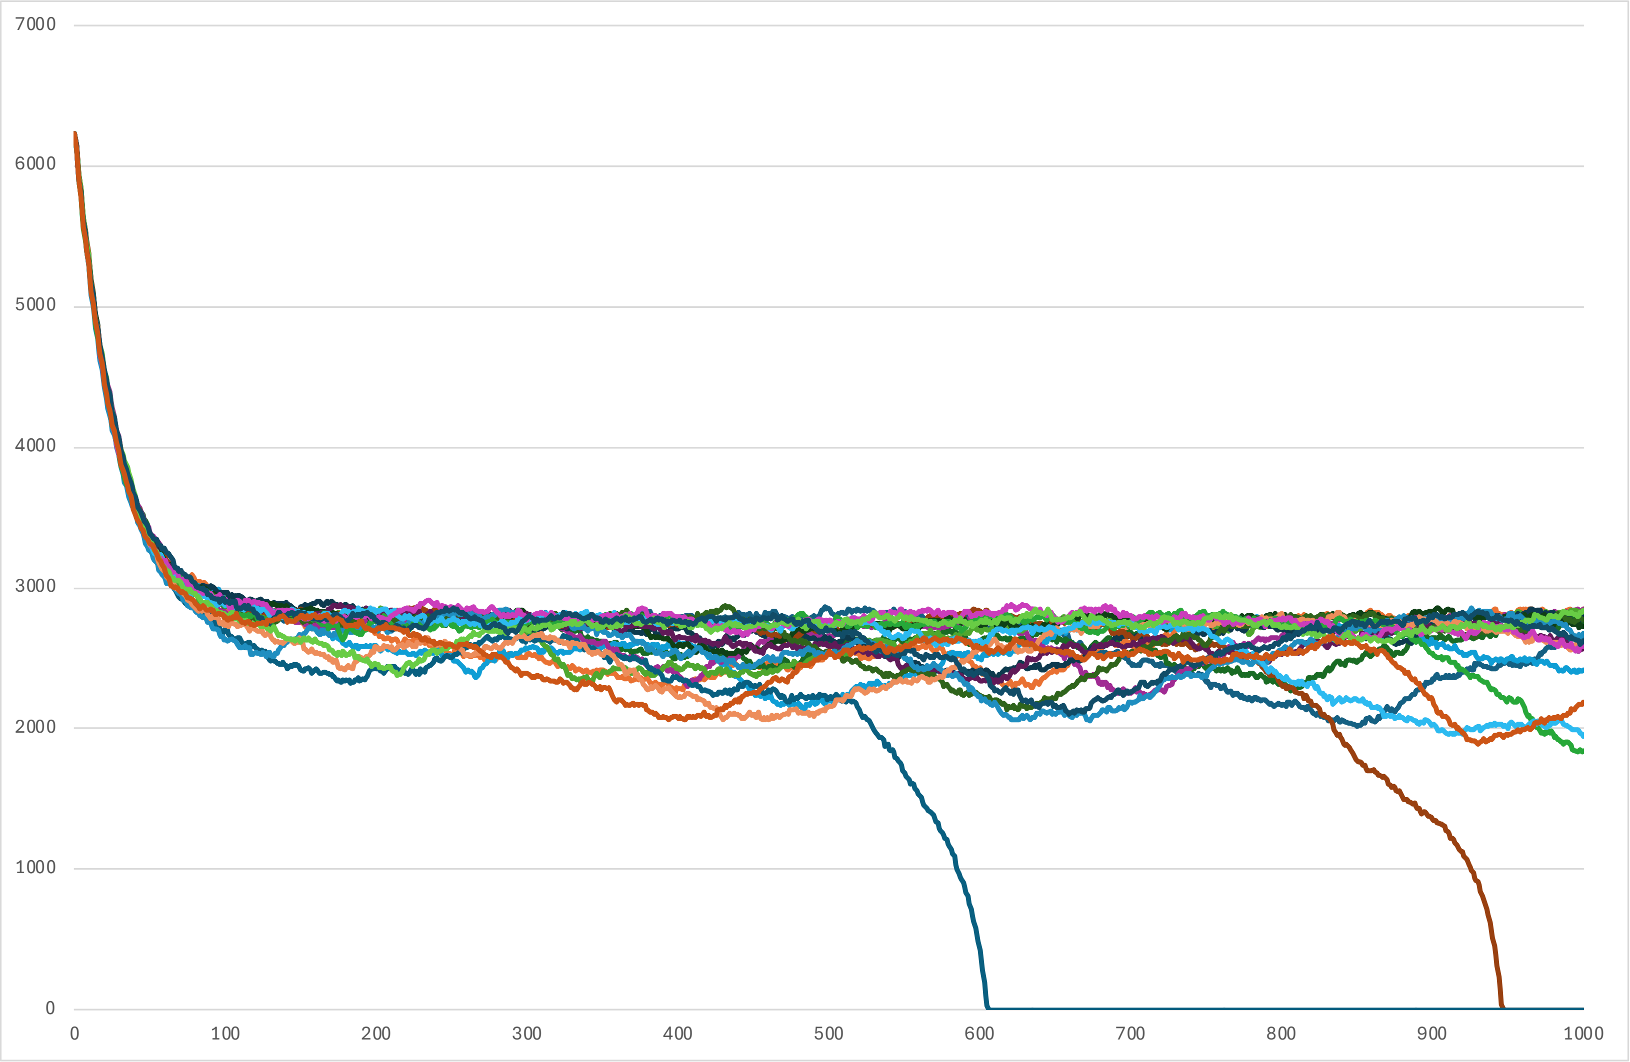

Fish numbers (av. over each run) Catch level = 120

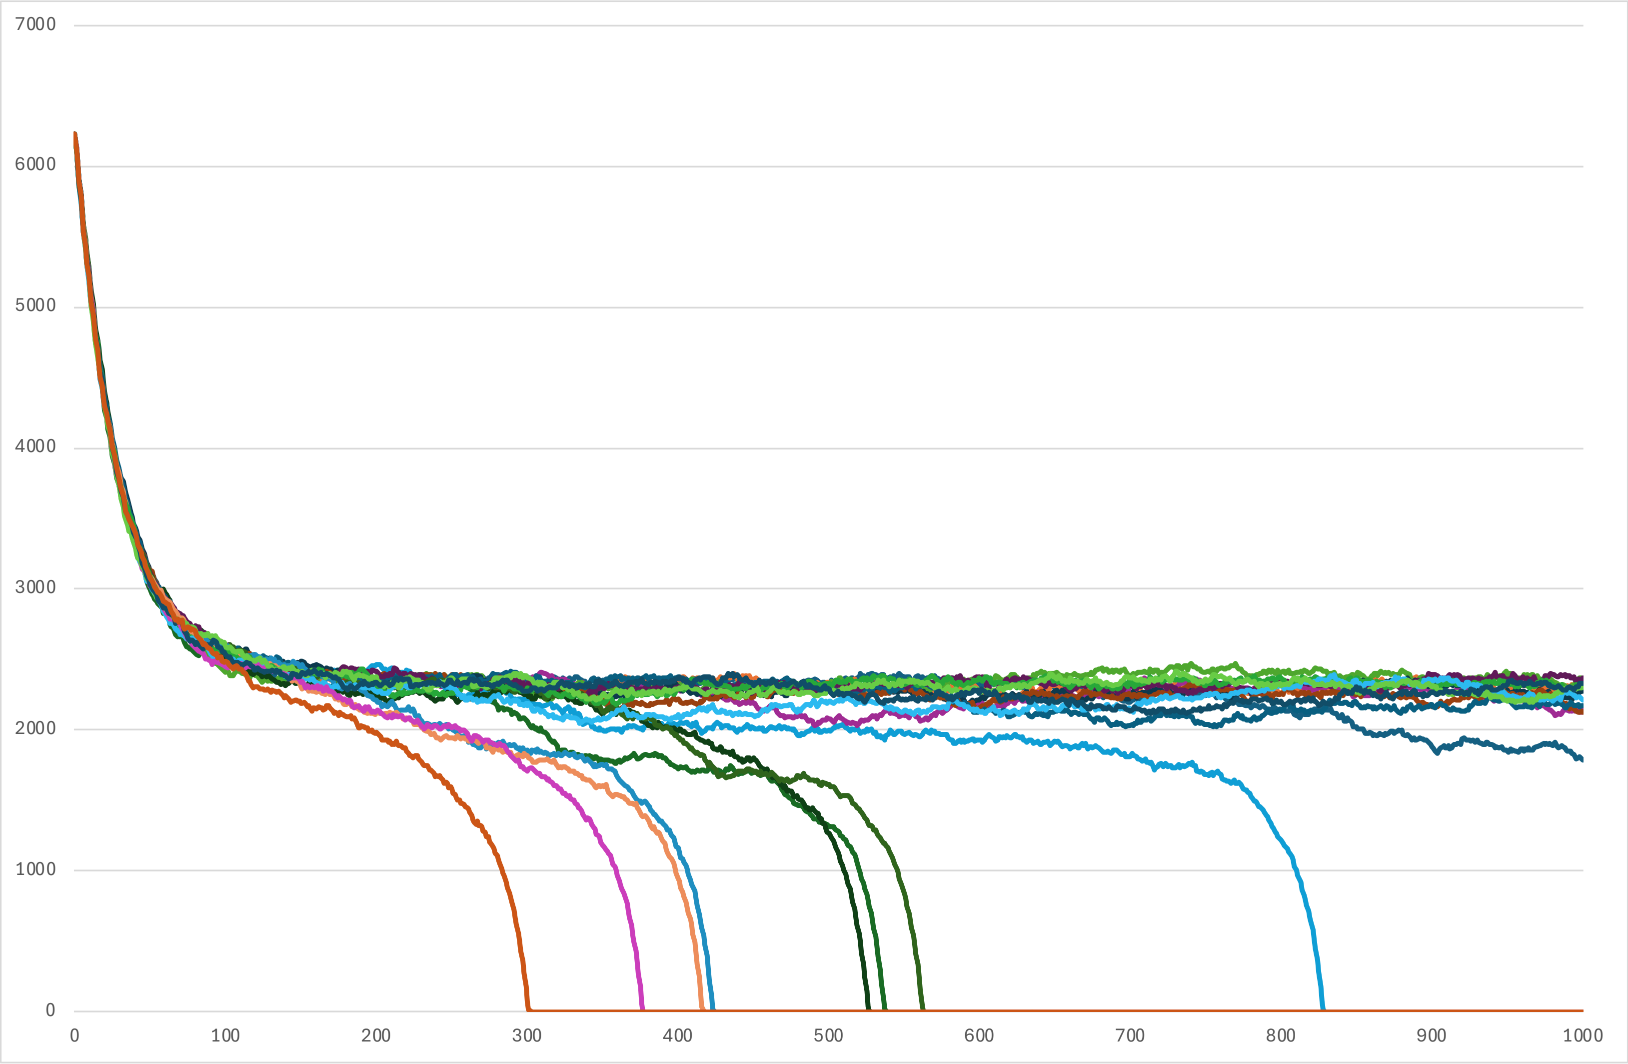

Fish numbers (av. over each run) Catch level = 130

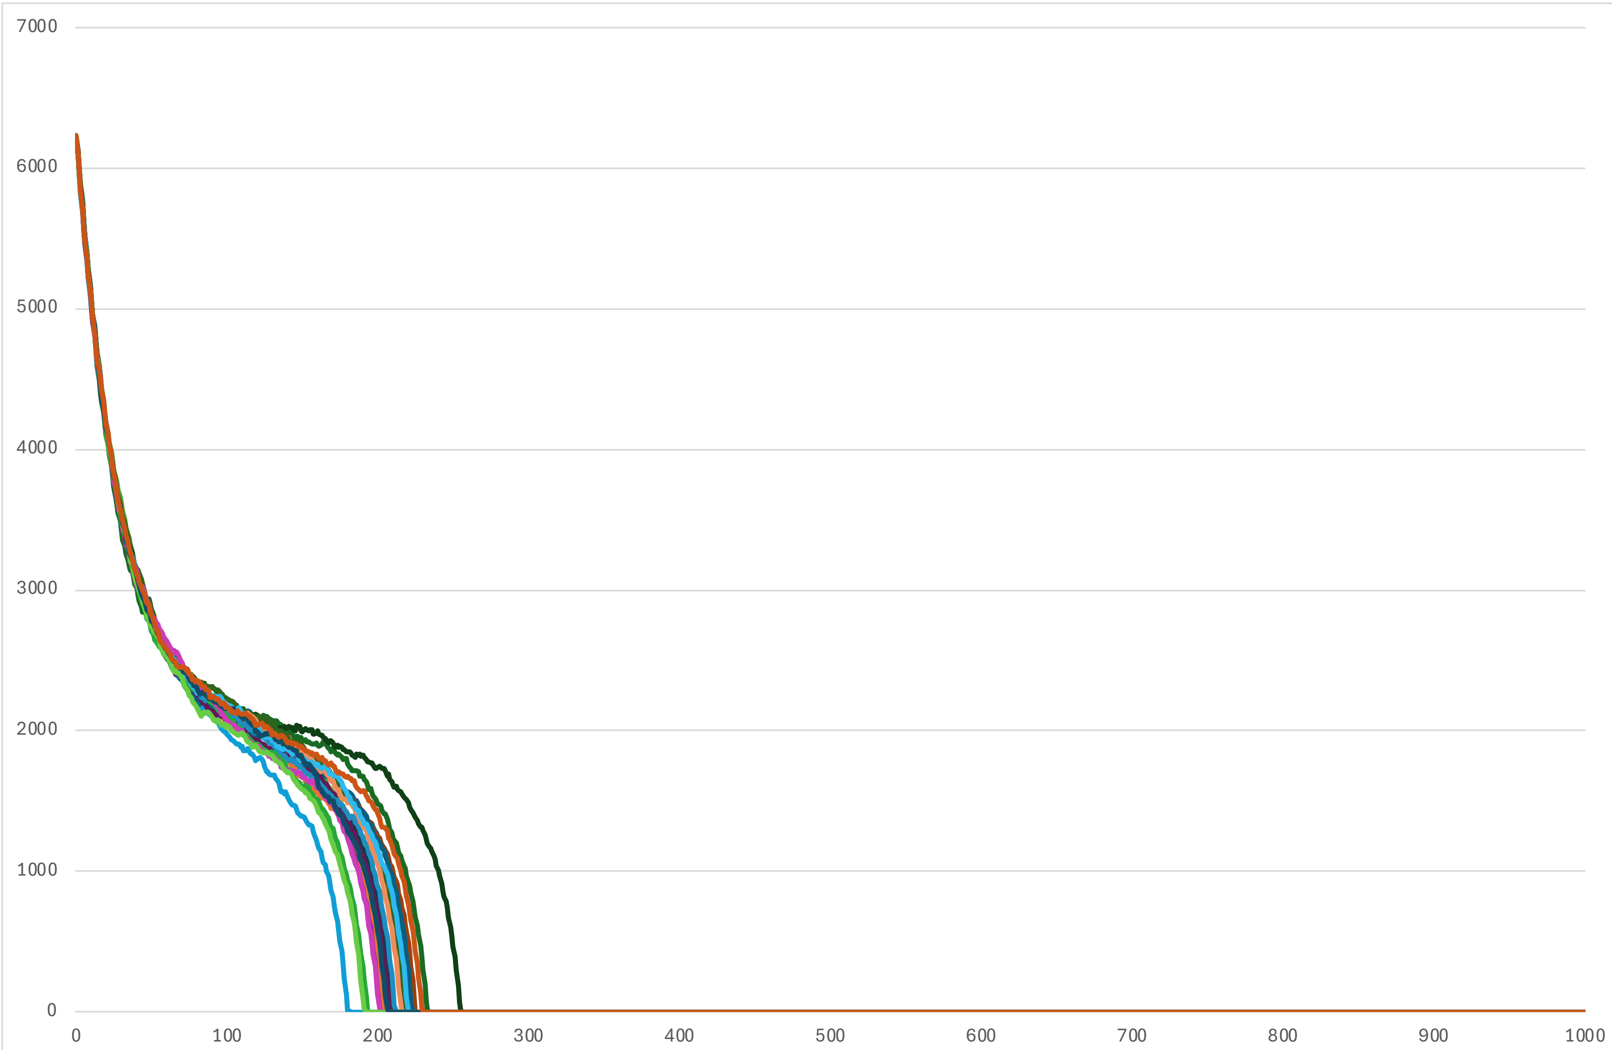

Total Fish Harvested vs Extinction Risk of long-term fishin (scaled so max=1)

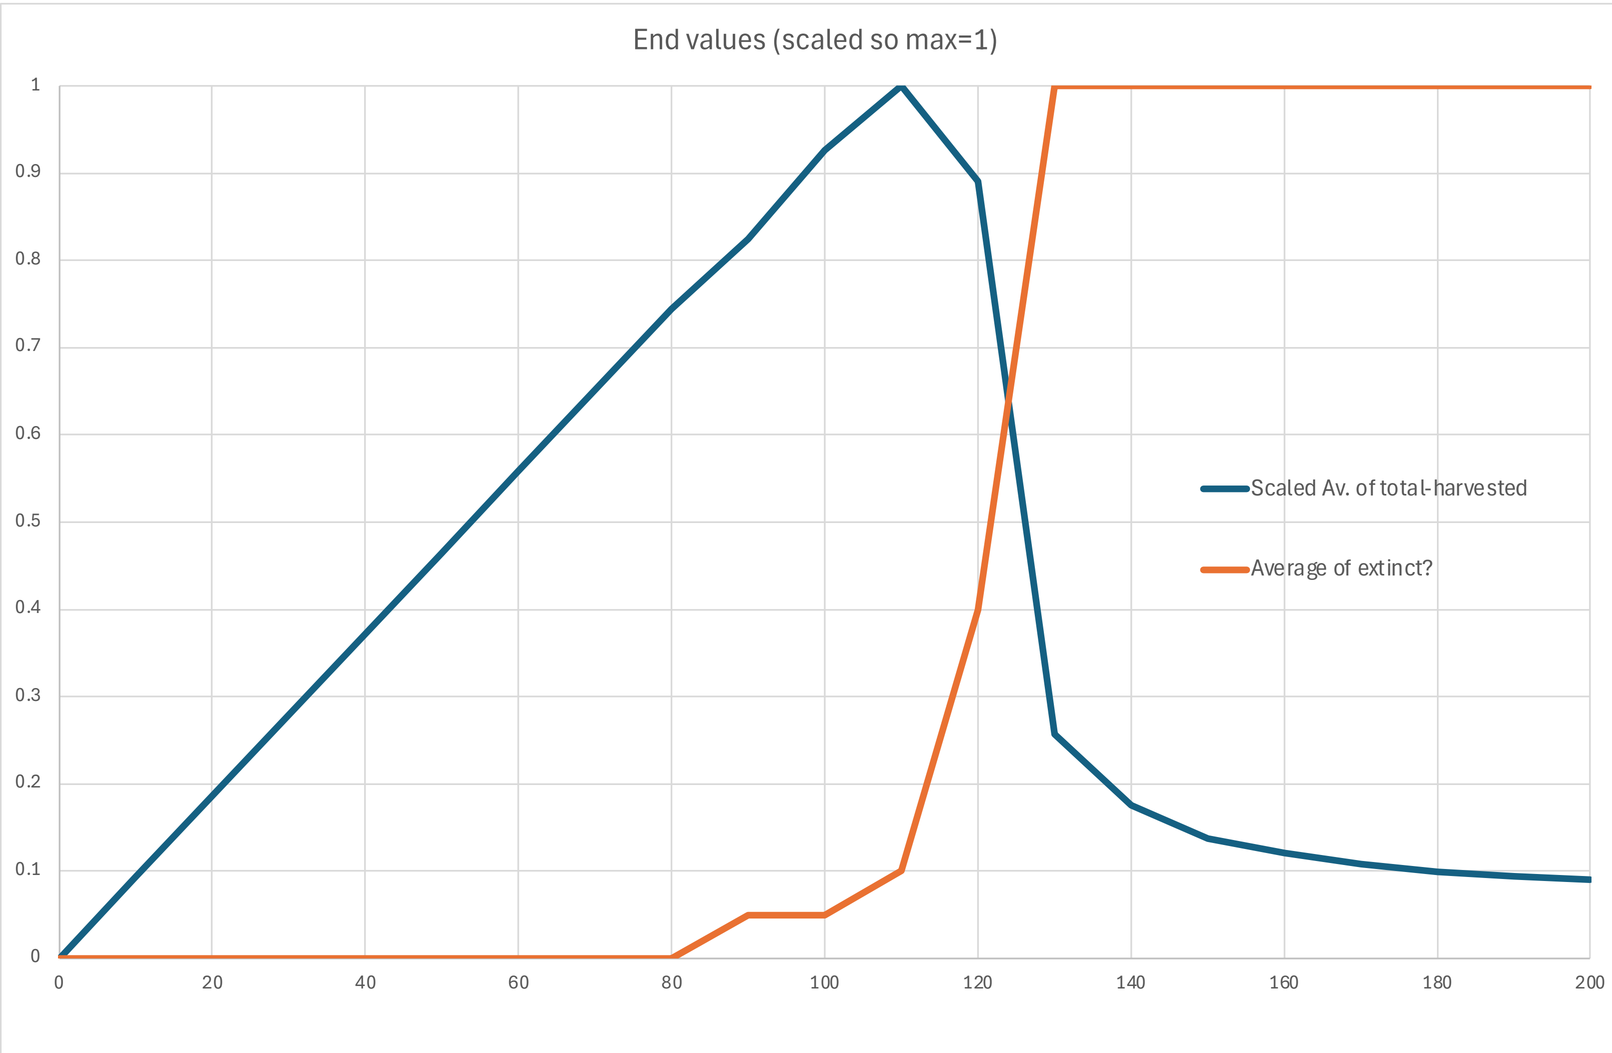

10000 – cont fishing harvested

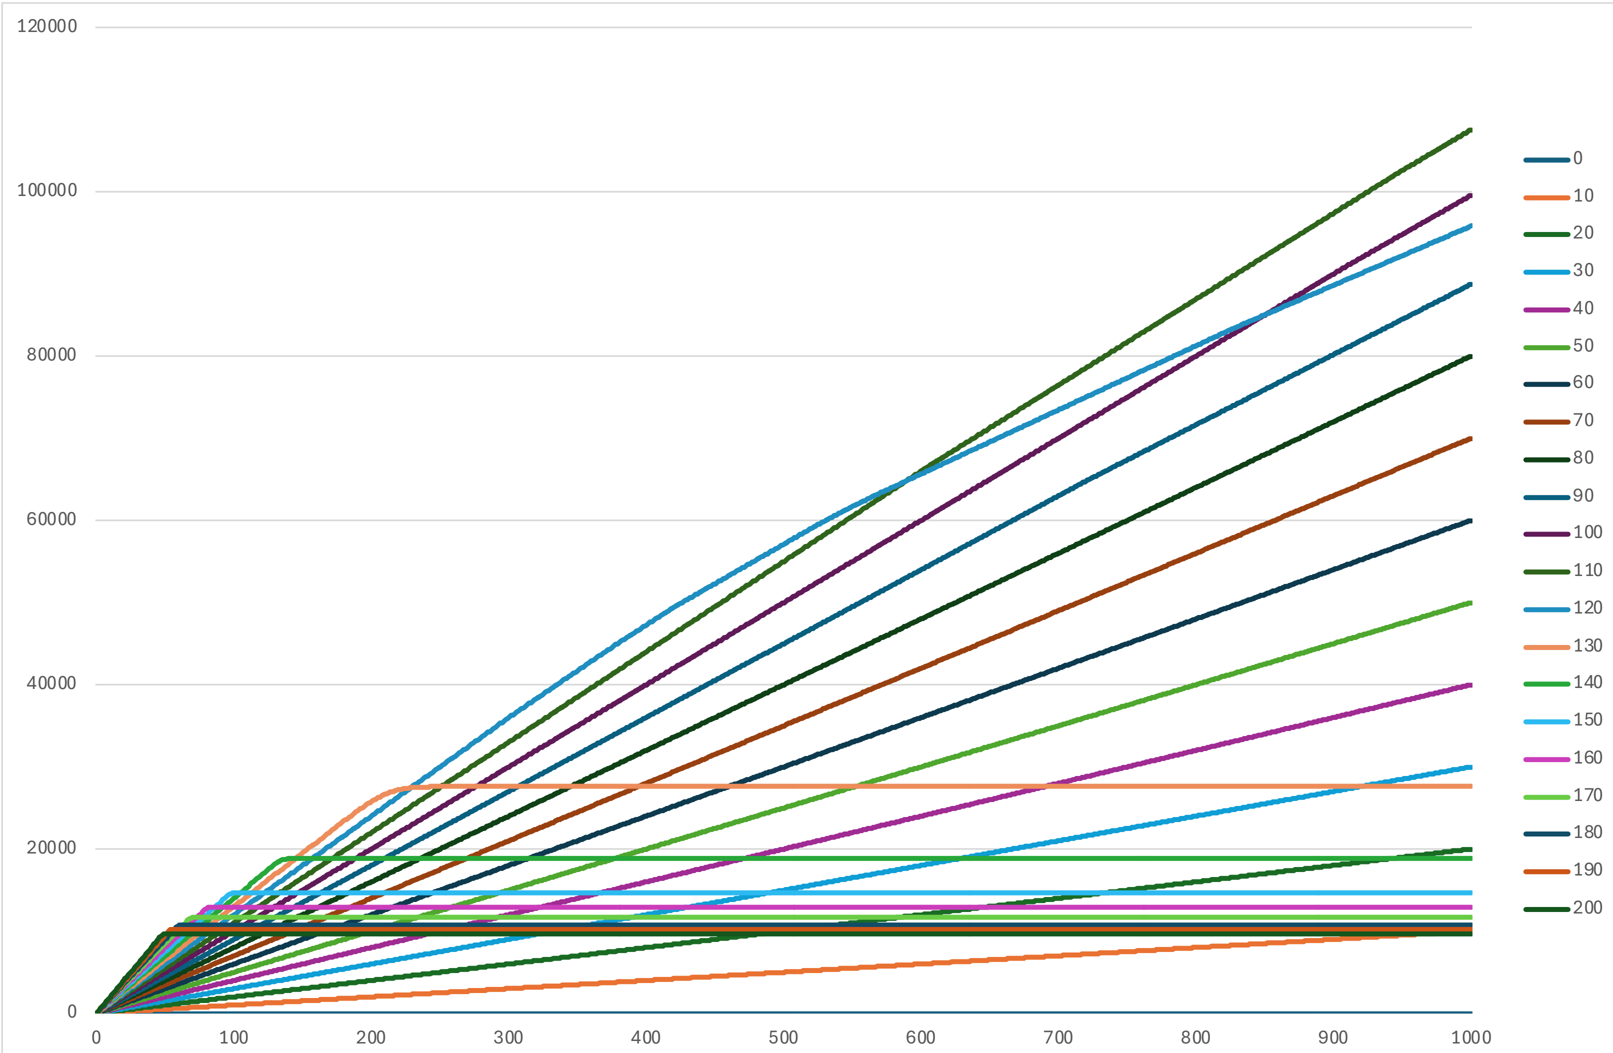

Fish numbers over time for different catch levels

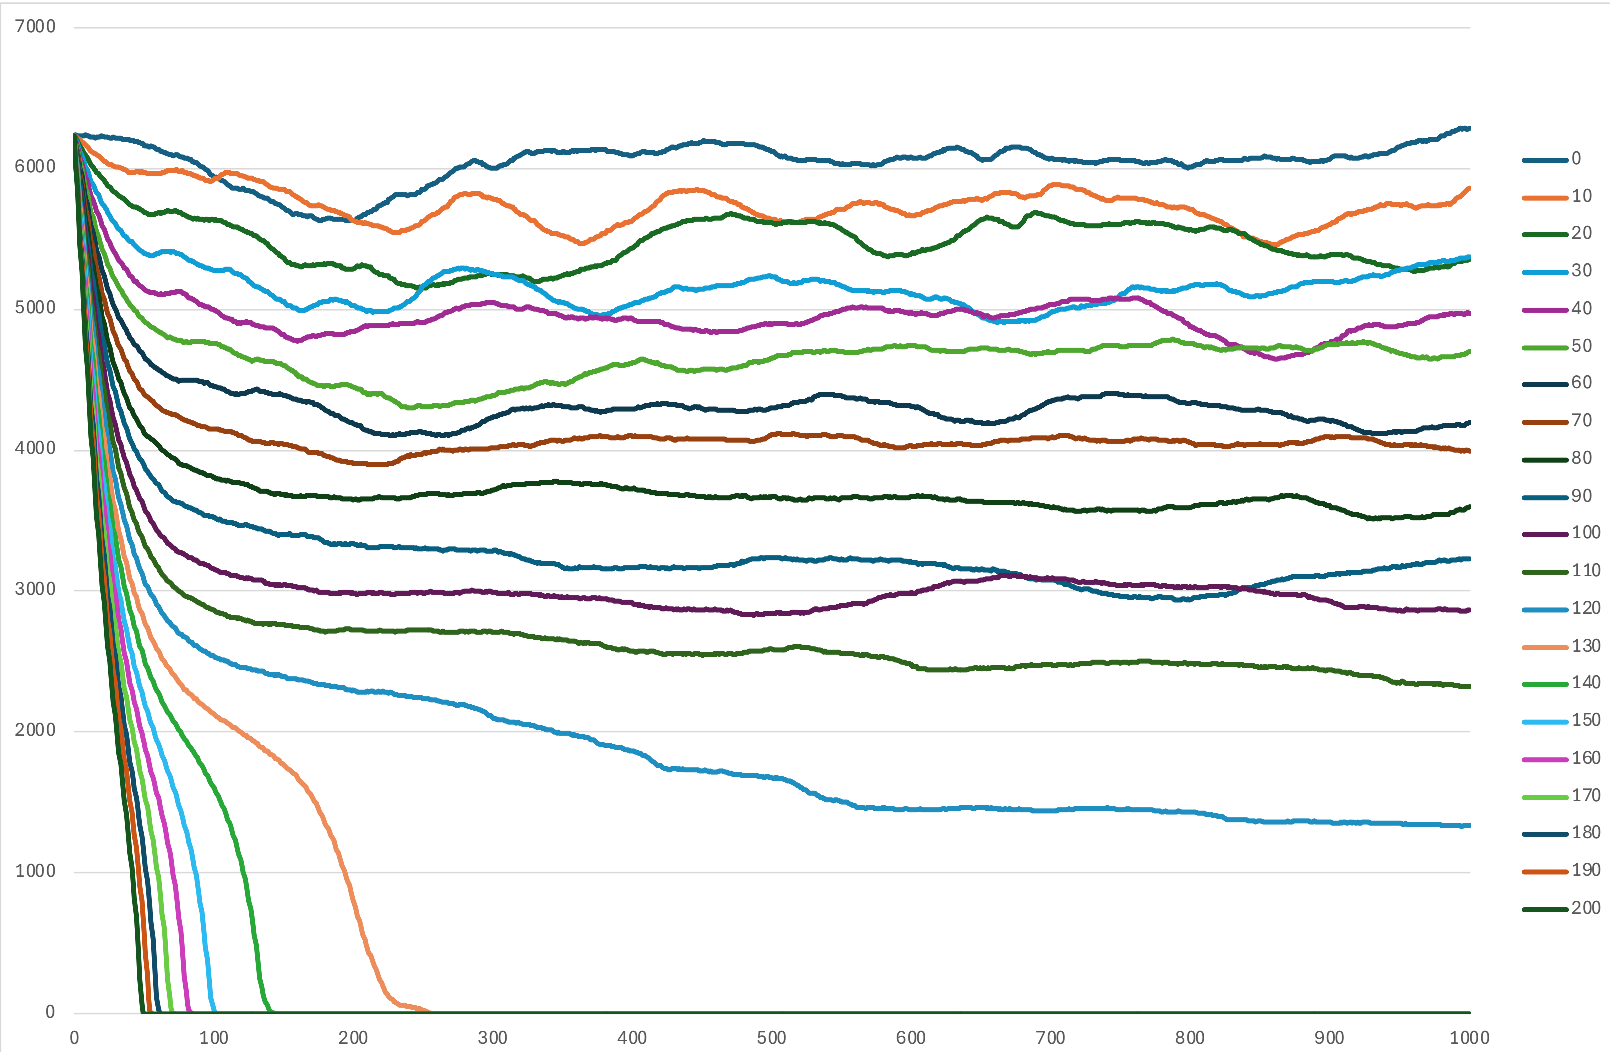

Number of species over time for different catch levels (averaged of 20 runs)

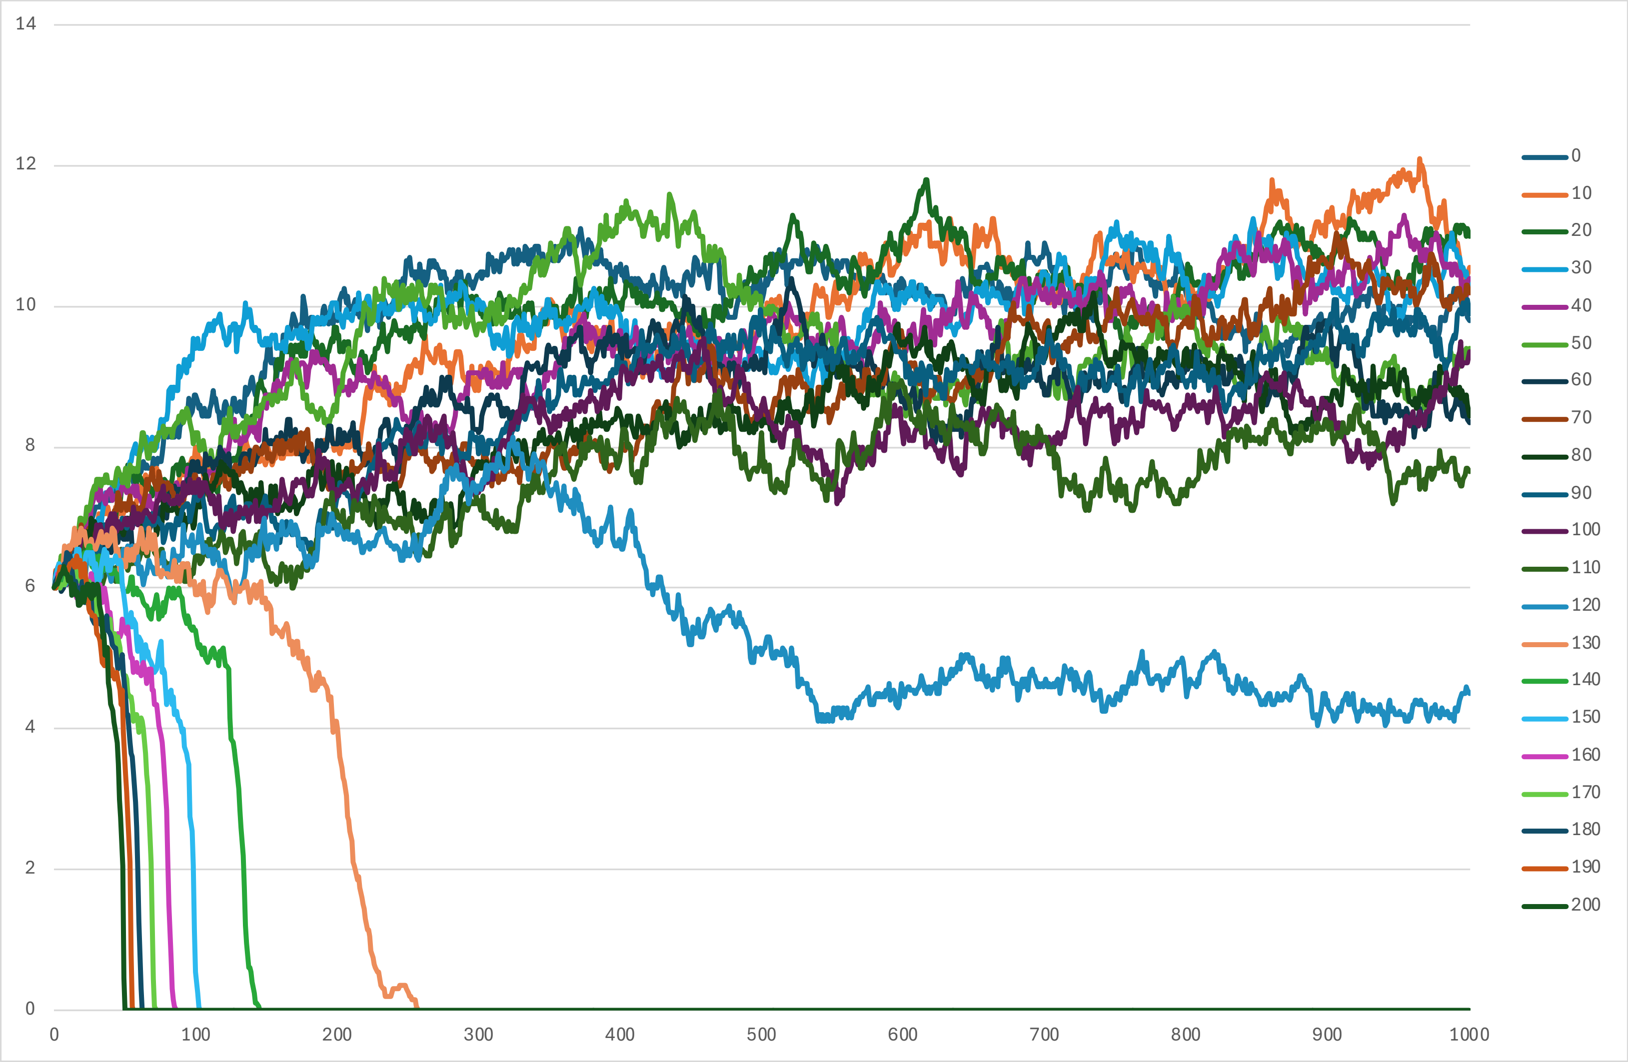

Average number of fish species at end of run, various levels of continuous catch (averaged over 20 runs)

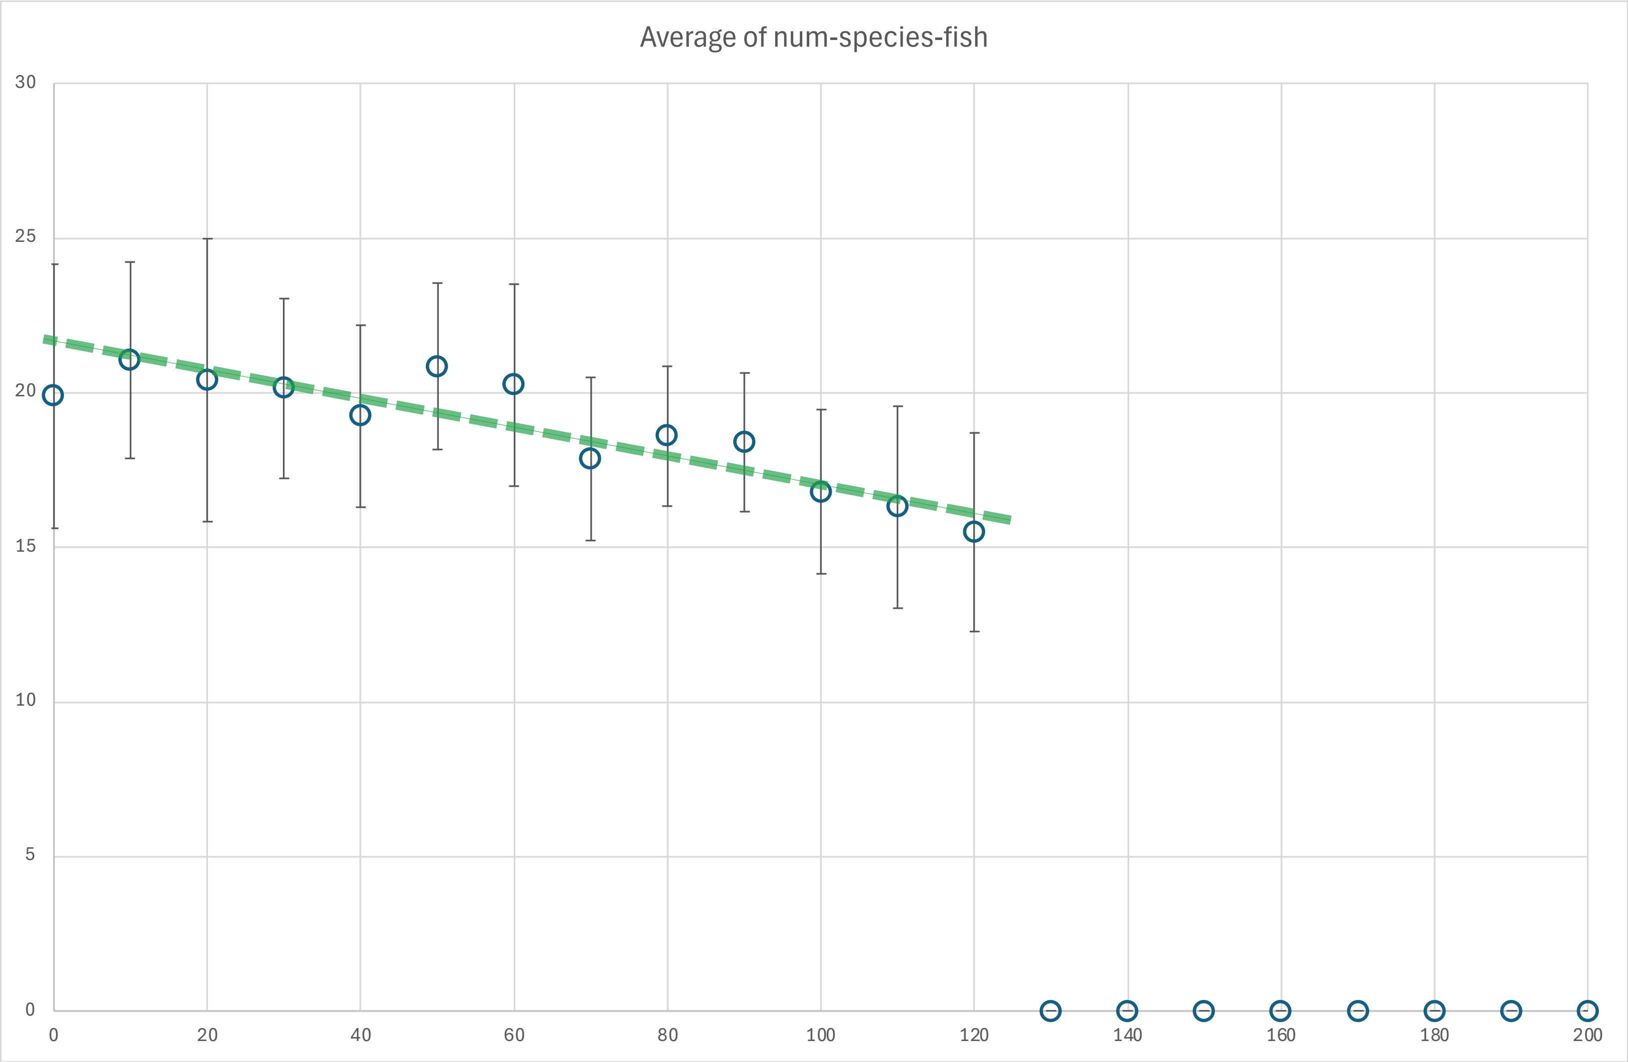

Total harvested against catch level

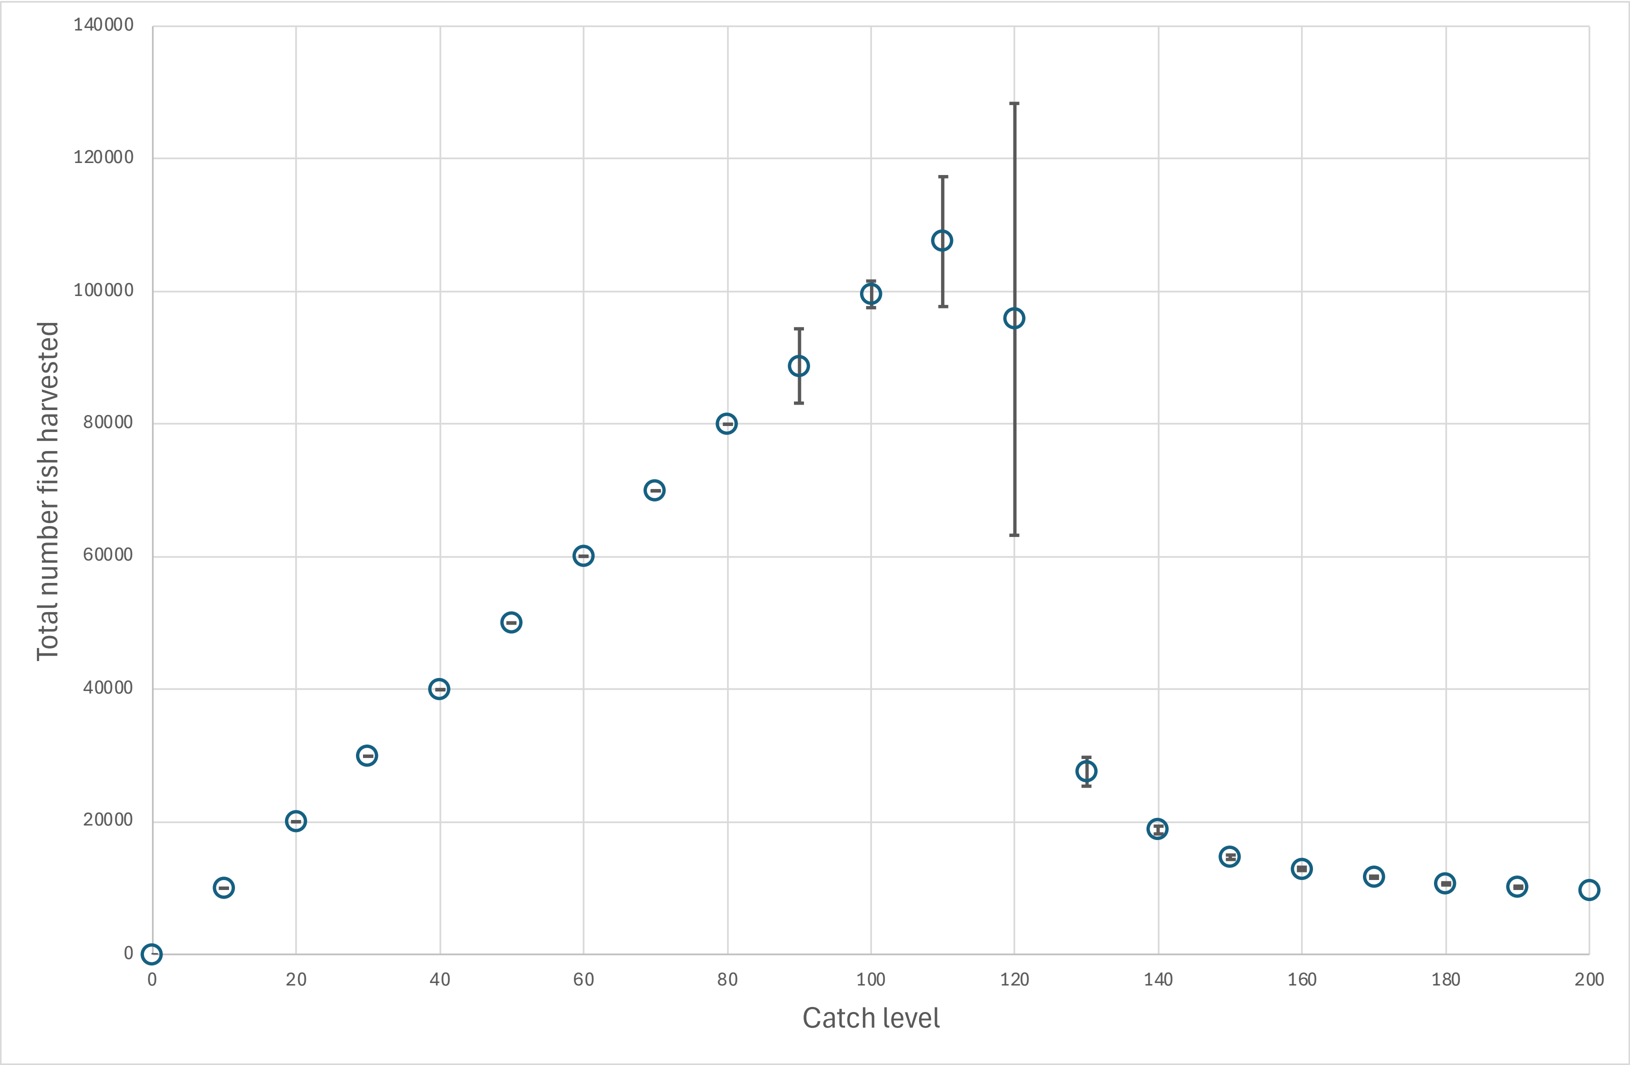

Look in more detail at catch level = 0 cases

- Rerun of catch level=0 cases (20 runs)
- Including measuring ecosystem diversity:
  - The average genetic ‘distance’ between all entities in the system
  - Distance is the number of locations in which the gene of species differ
  - Approximated using a sample of 1000 pairs
  - This is NOT practical in real life, given what data we have and how we fish!
  - Quite time-consuming to calculate, so only do so every 10 ticks

Rerun, catch=0 fish numbers (averaged in each run)

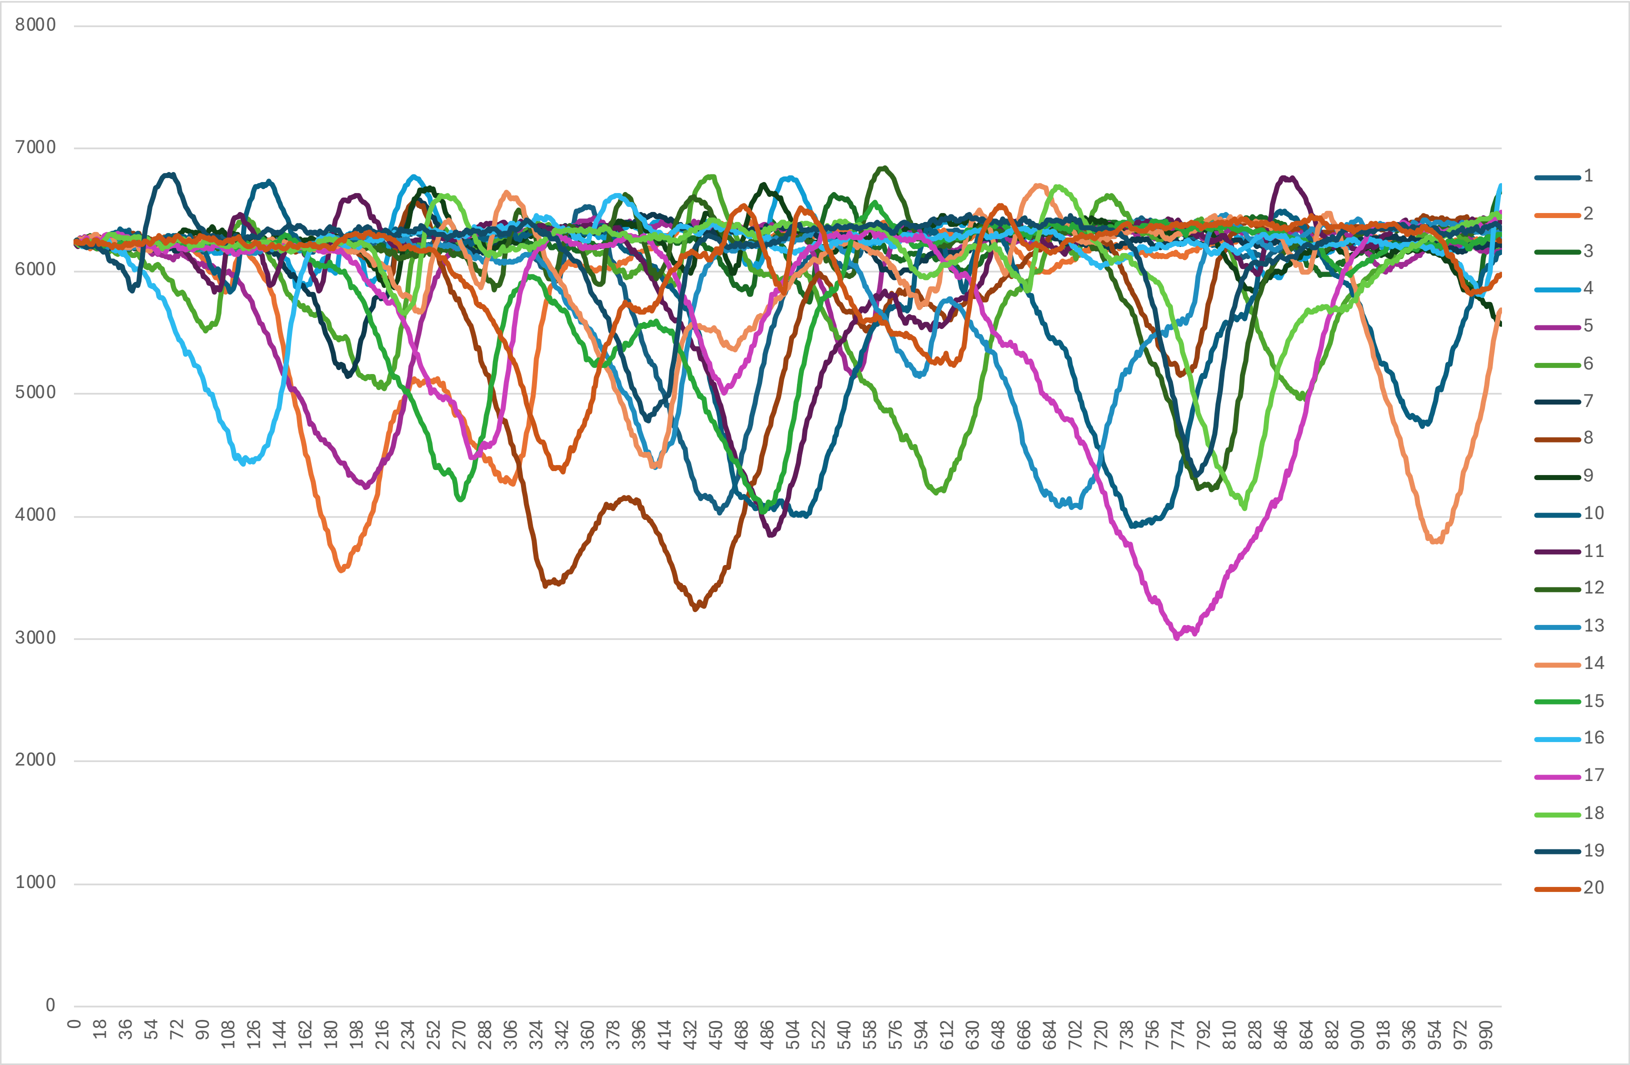

Rerun, catch=0 number fish species (averaged in each run)

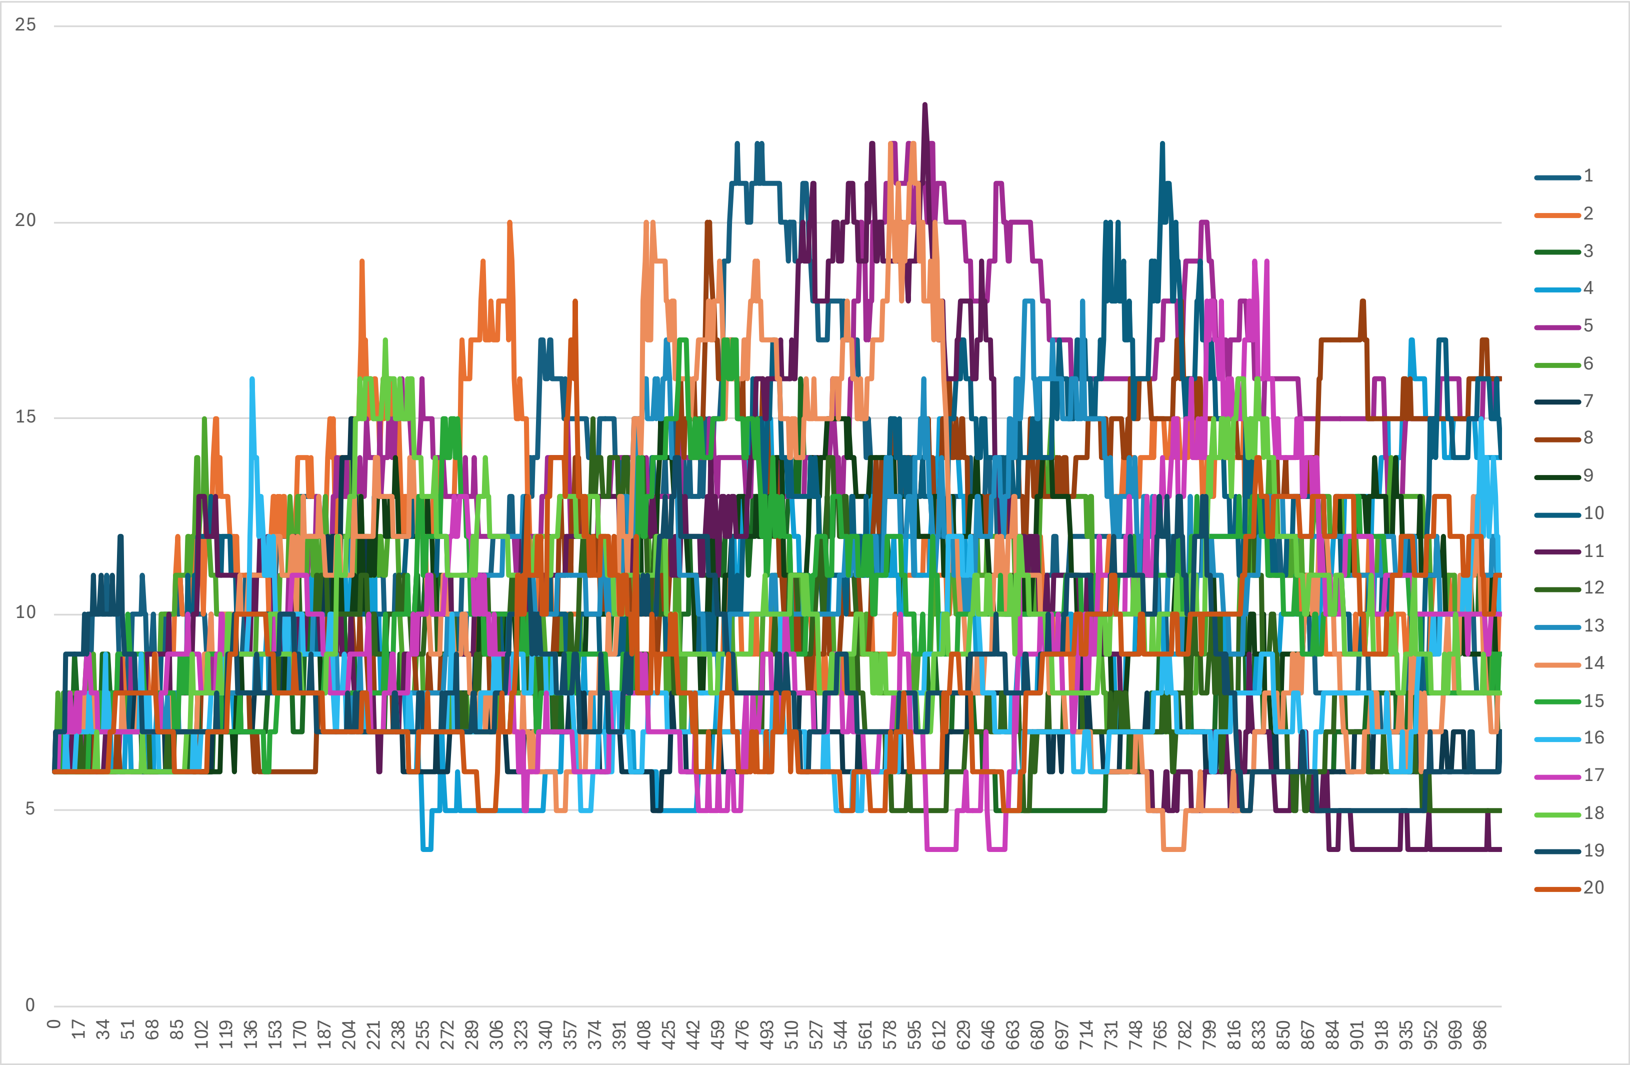

Rerun, catch=0 ecosystem diversity (averaged in each run)

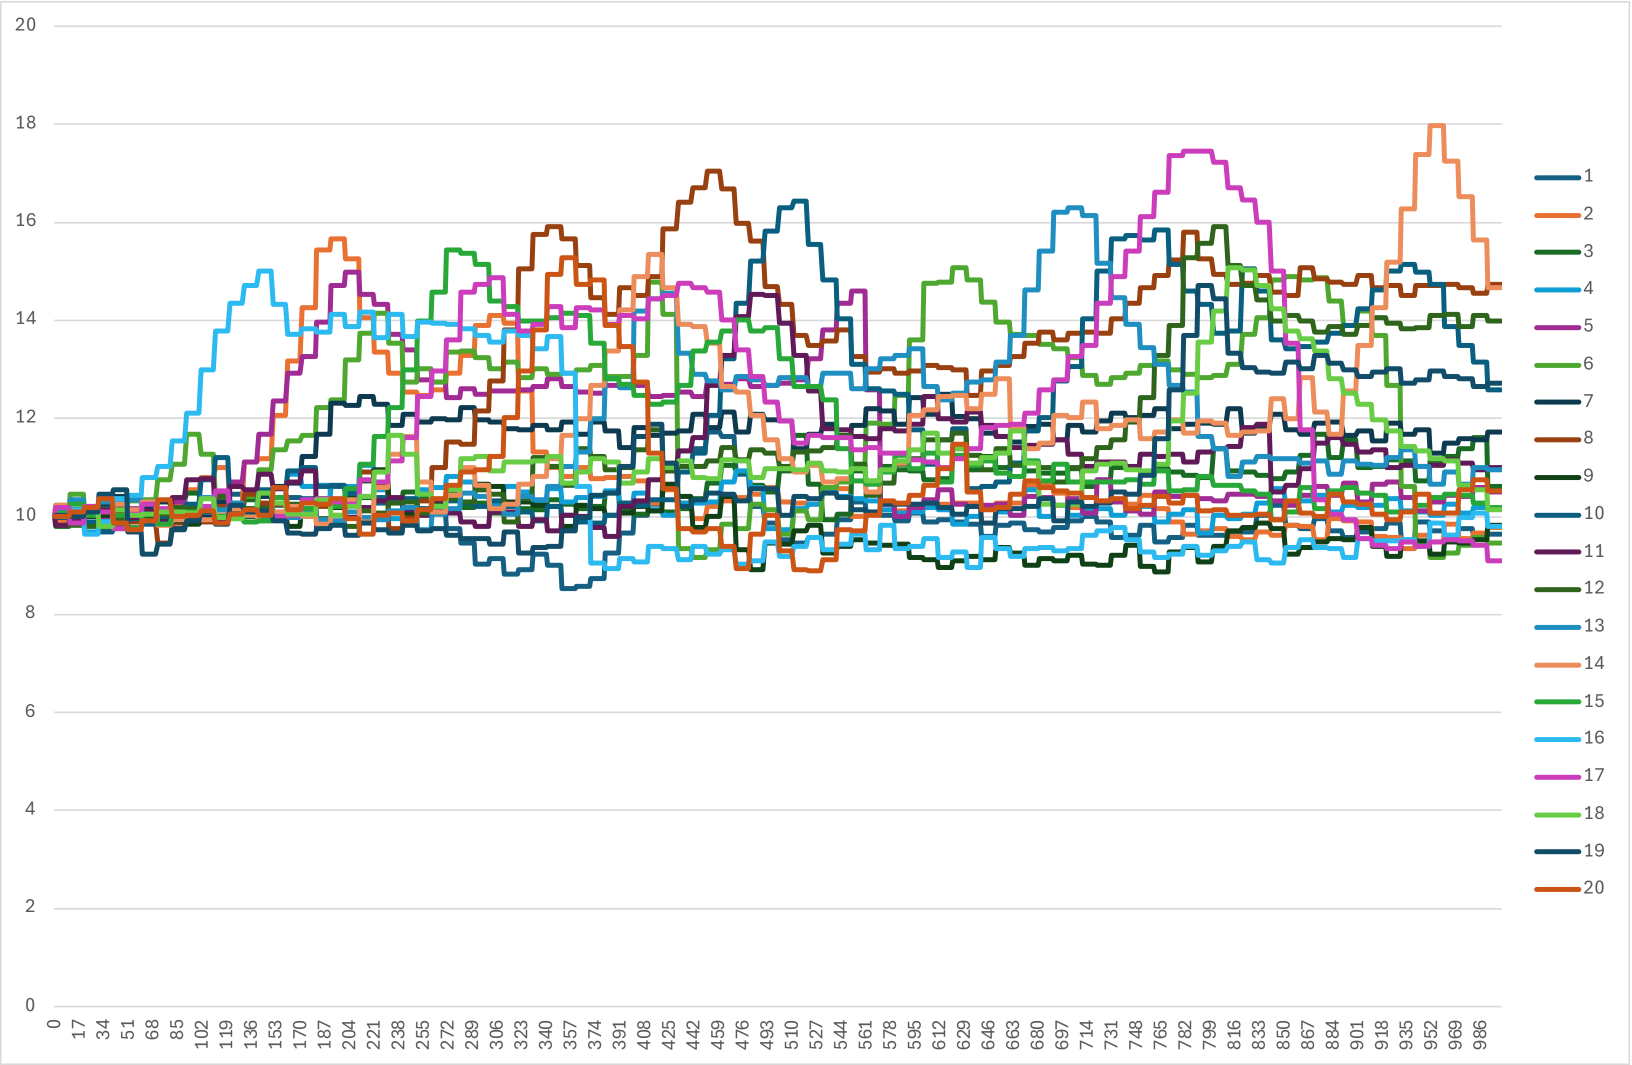

Rerun, catch=0 run 17 diversity & fish numbs (both scaled so max=1)

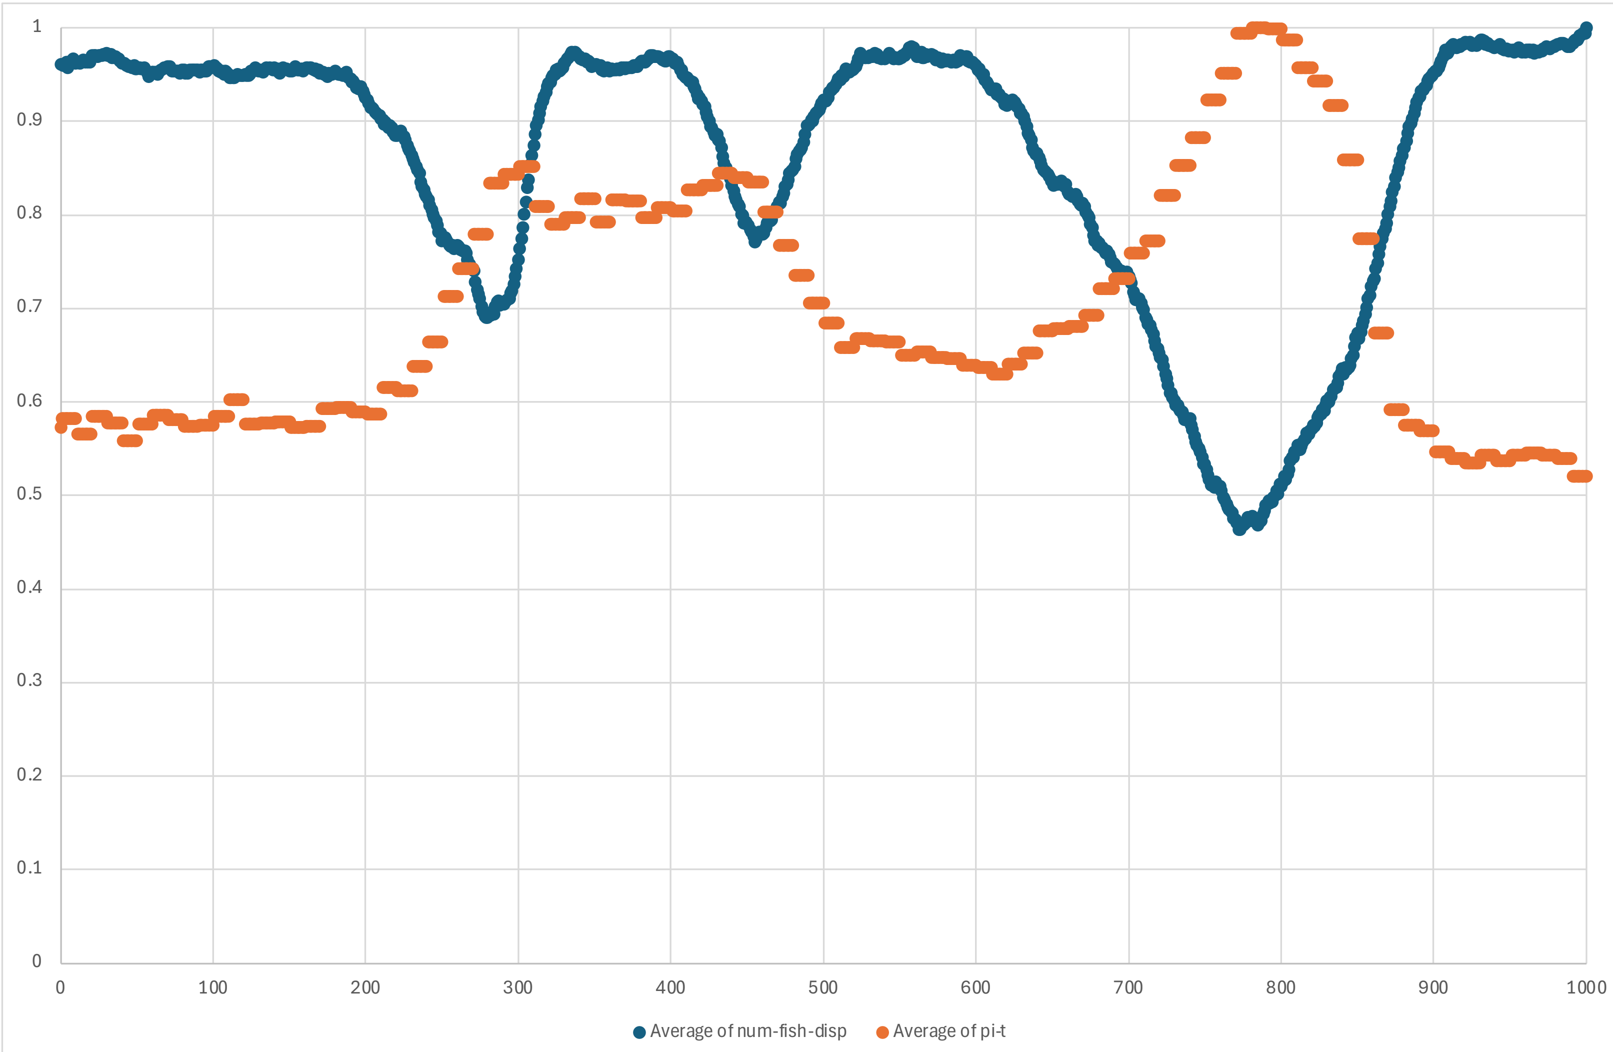

Rerun, catch=0 run 17, numbers of each fish species

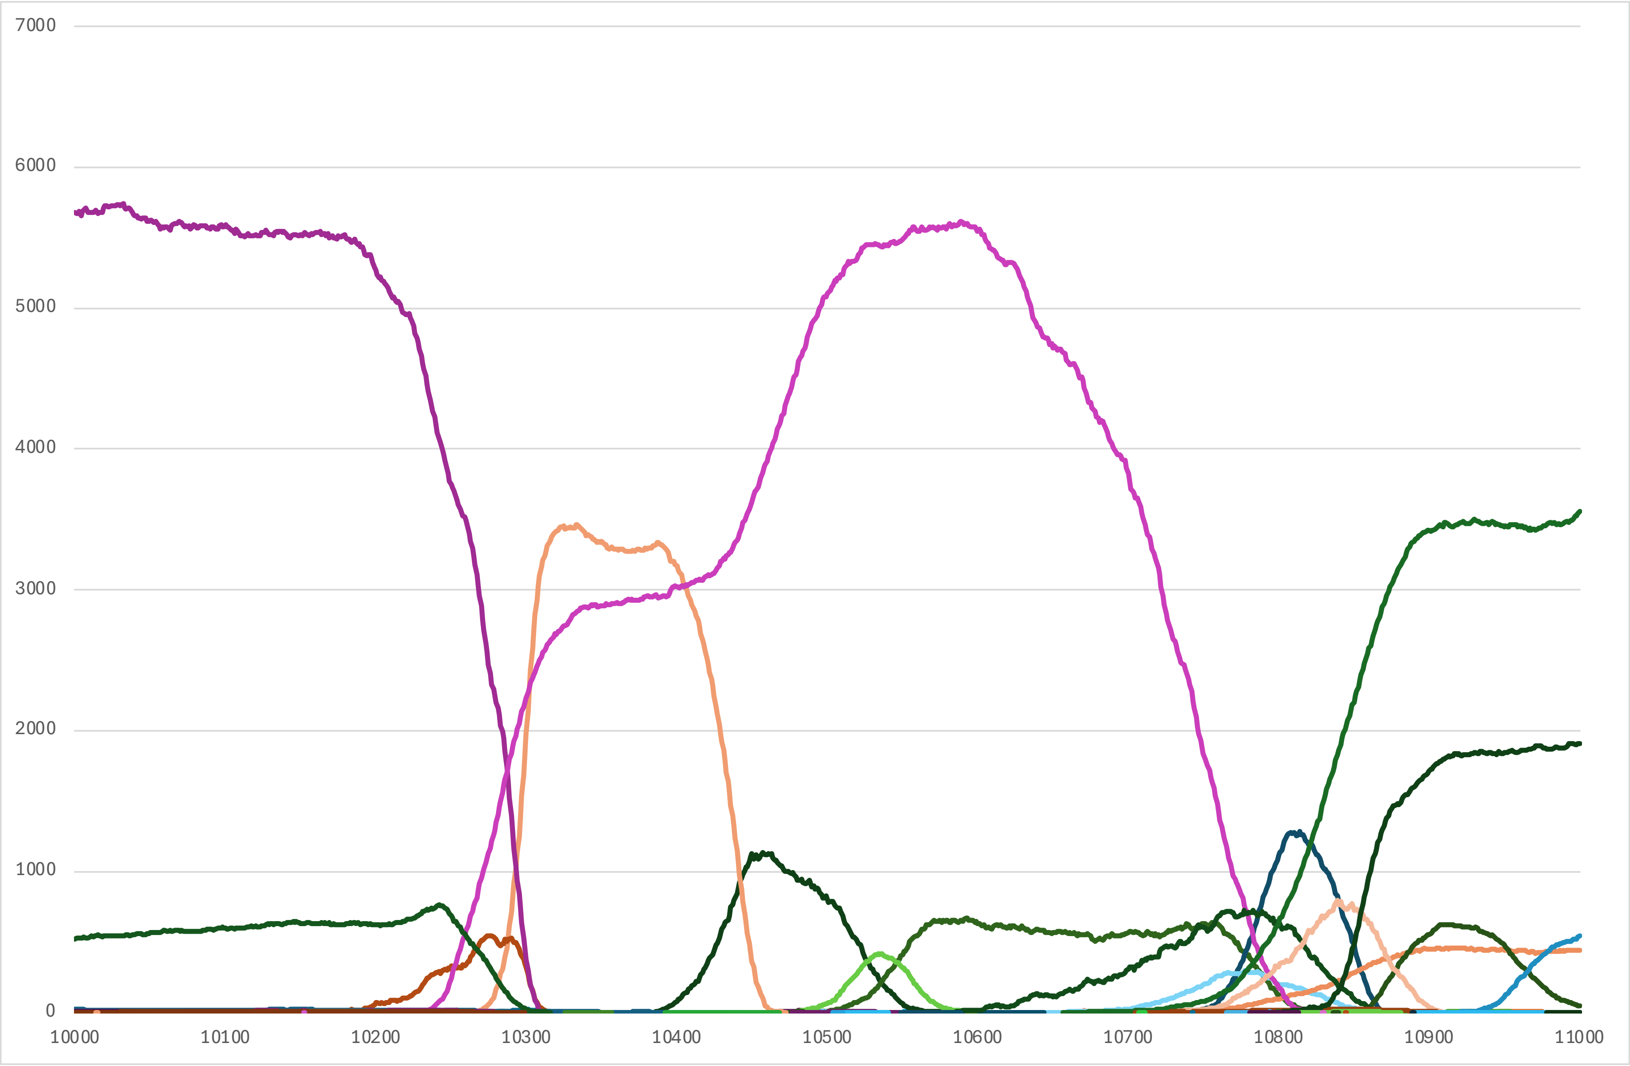

Look in more detail at catch level = 110 cases

- Rerun of catch level=110 cases (20 runs)
- Including measuring ecosystem diversity

10000 – cont fishing, 110 – fish numbers

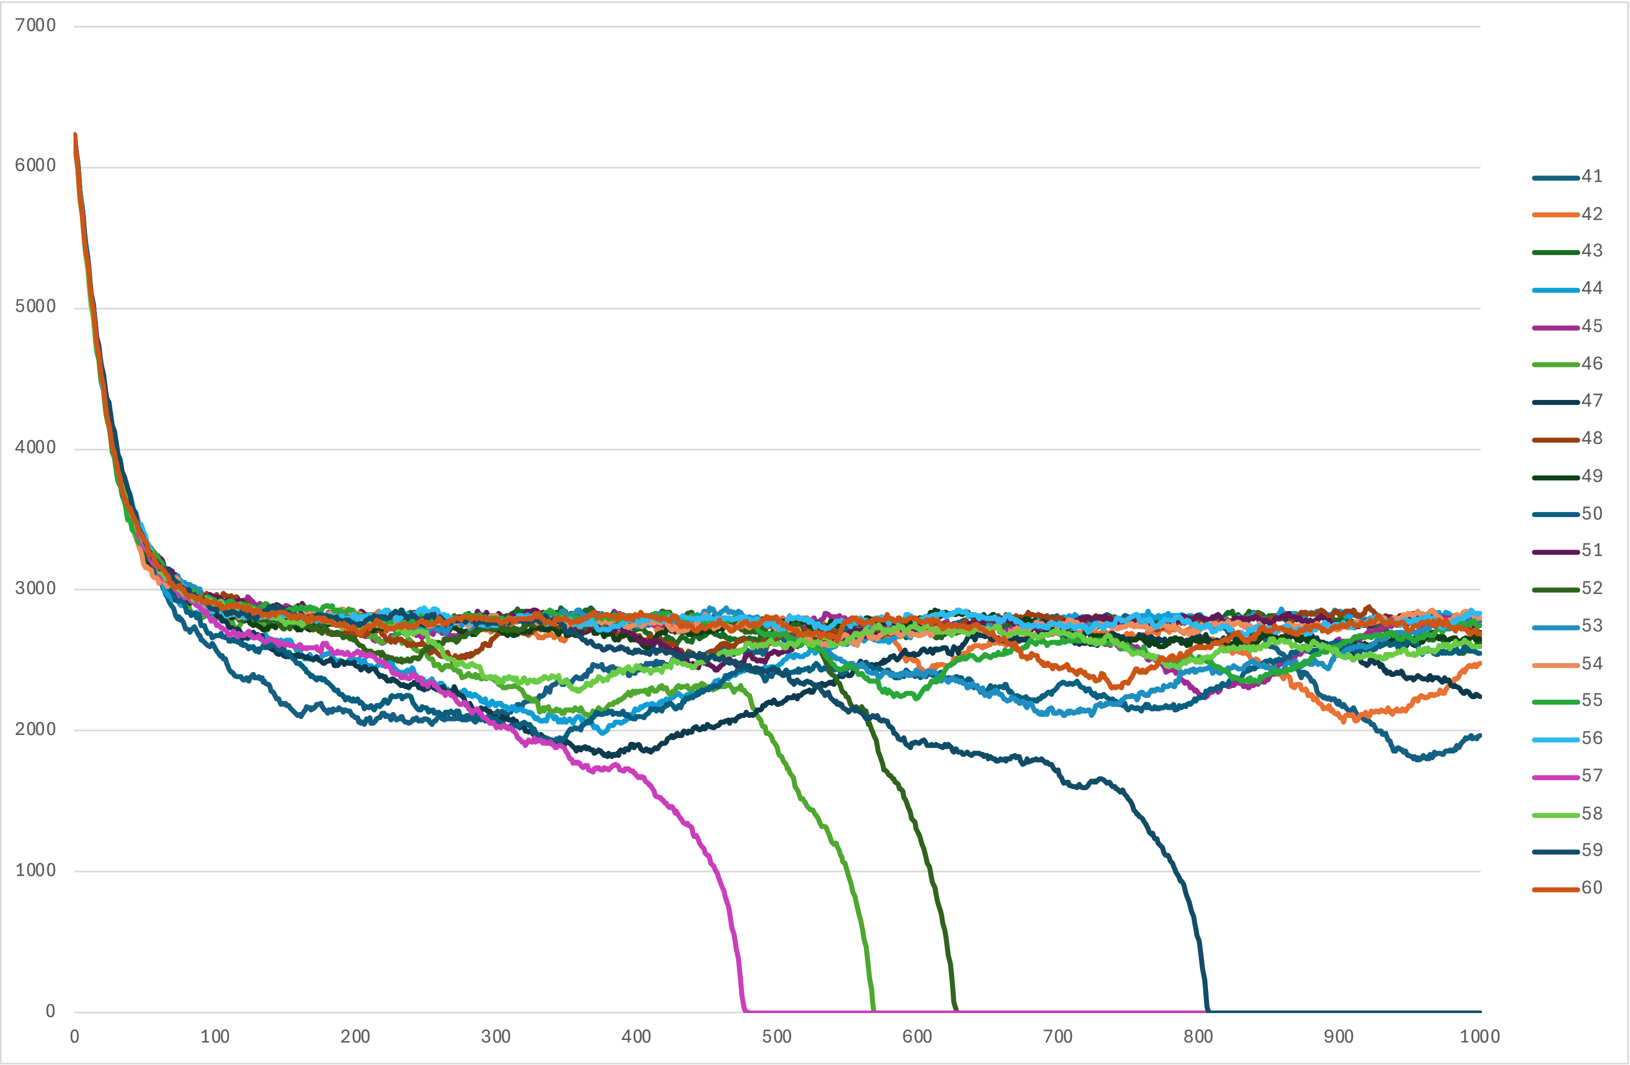

10000 – cont fishing, 110 – diversity

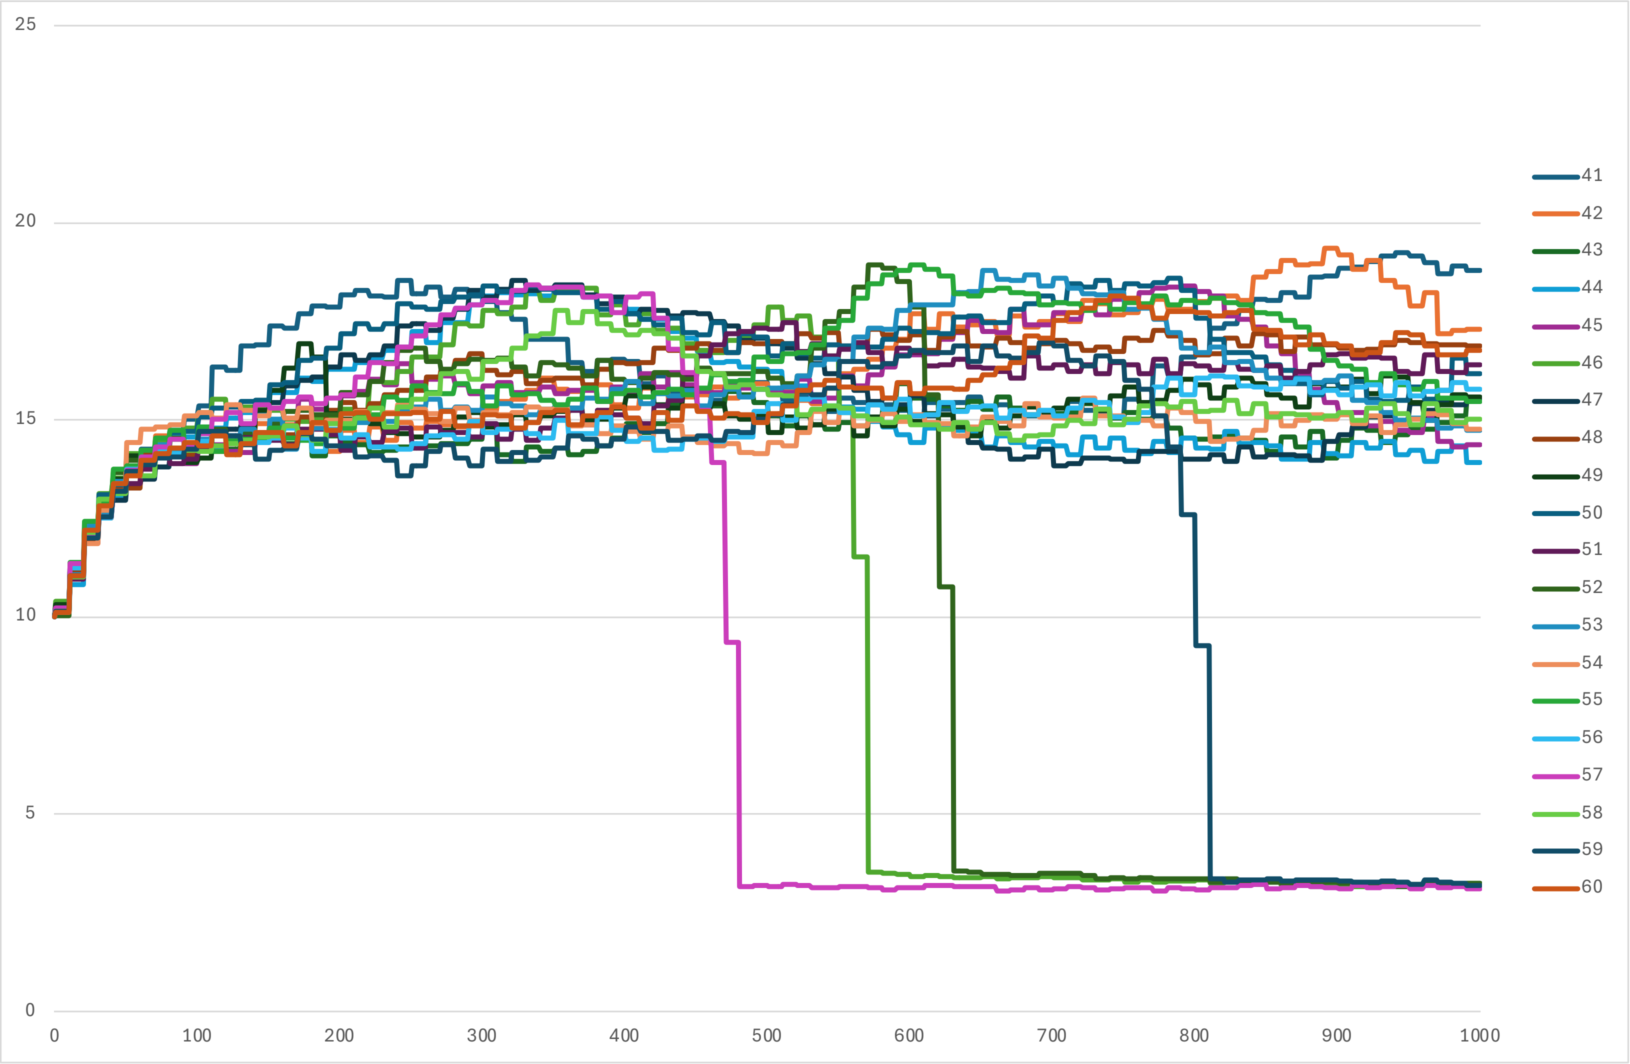

10000 – cont fishing, 110 – diversity (extinct only)

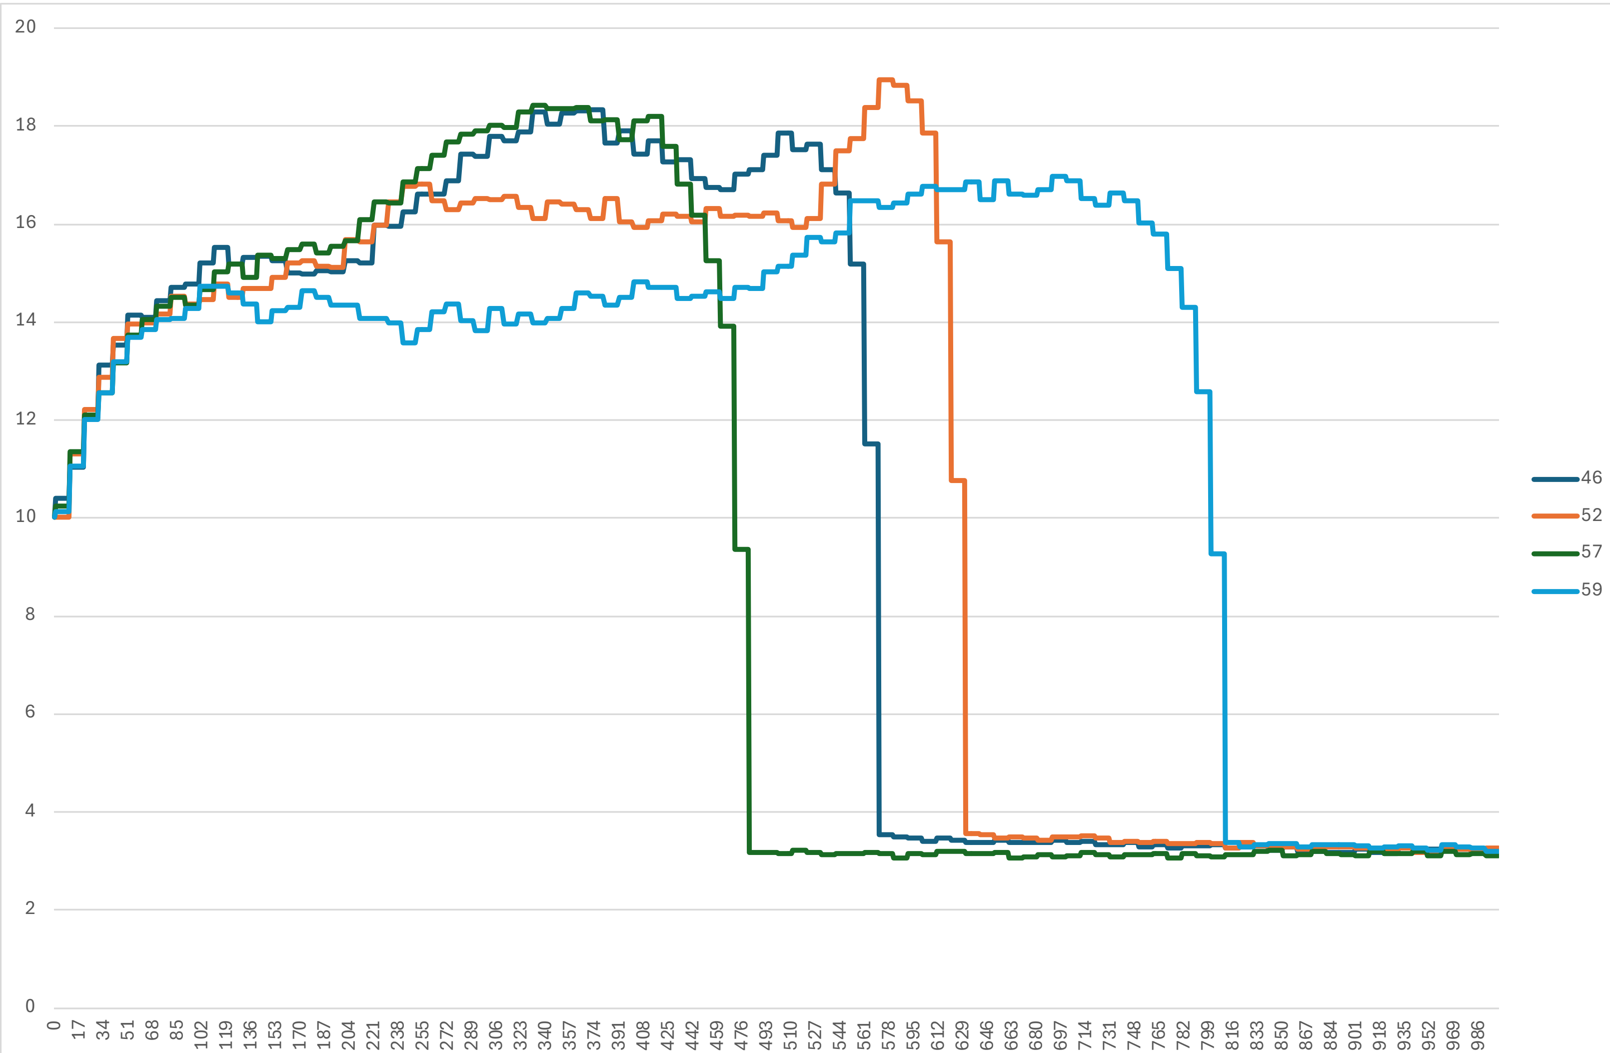

10000 – cont fishing, 110 – num fish species (extinct only)

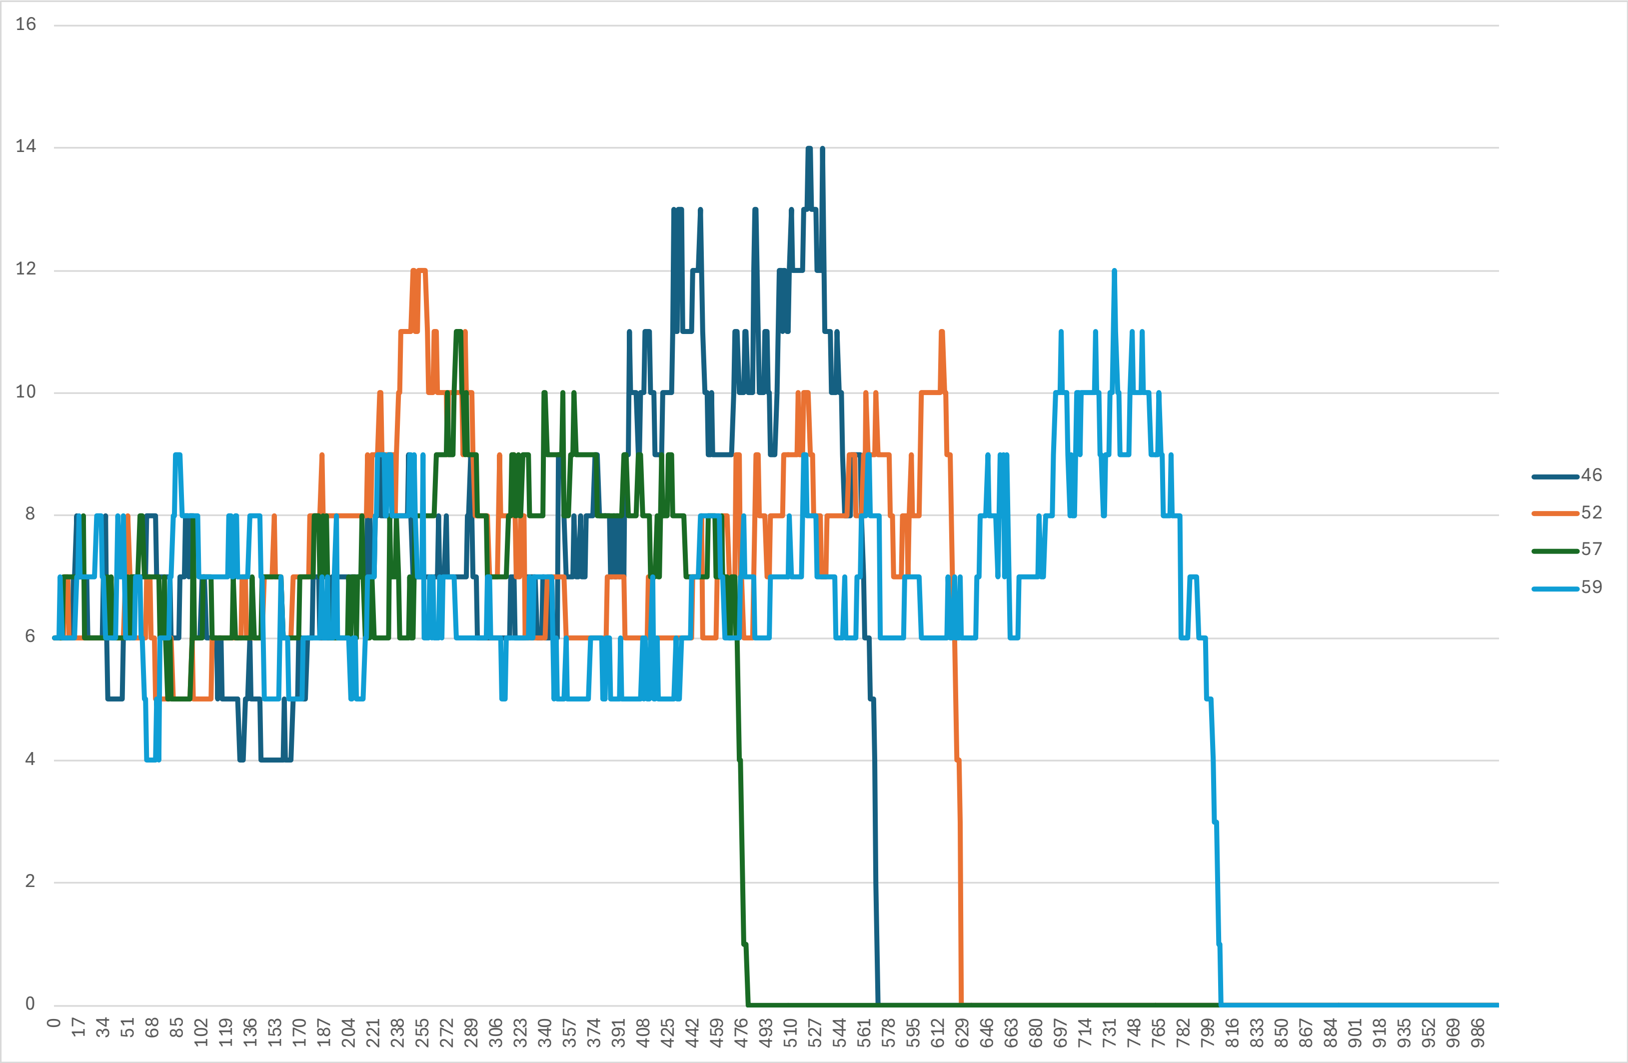

catch level 110 – particular runs that went *extinct* – number of each species

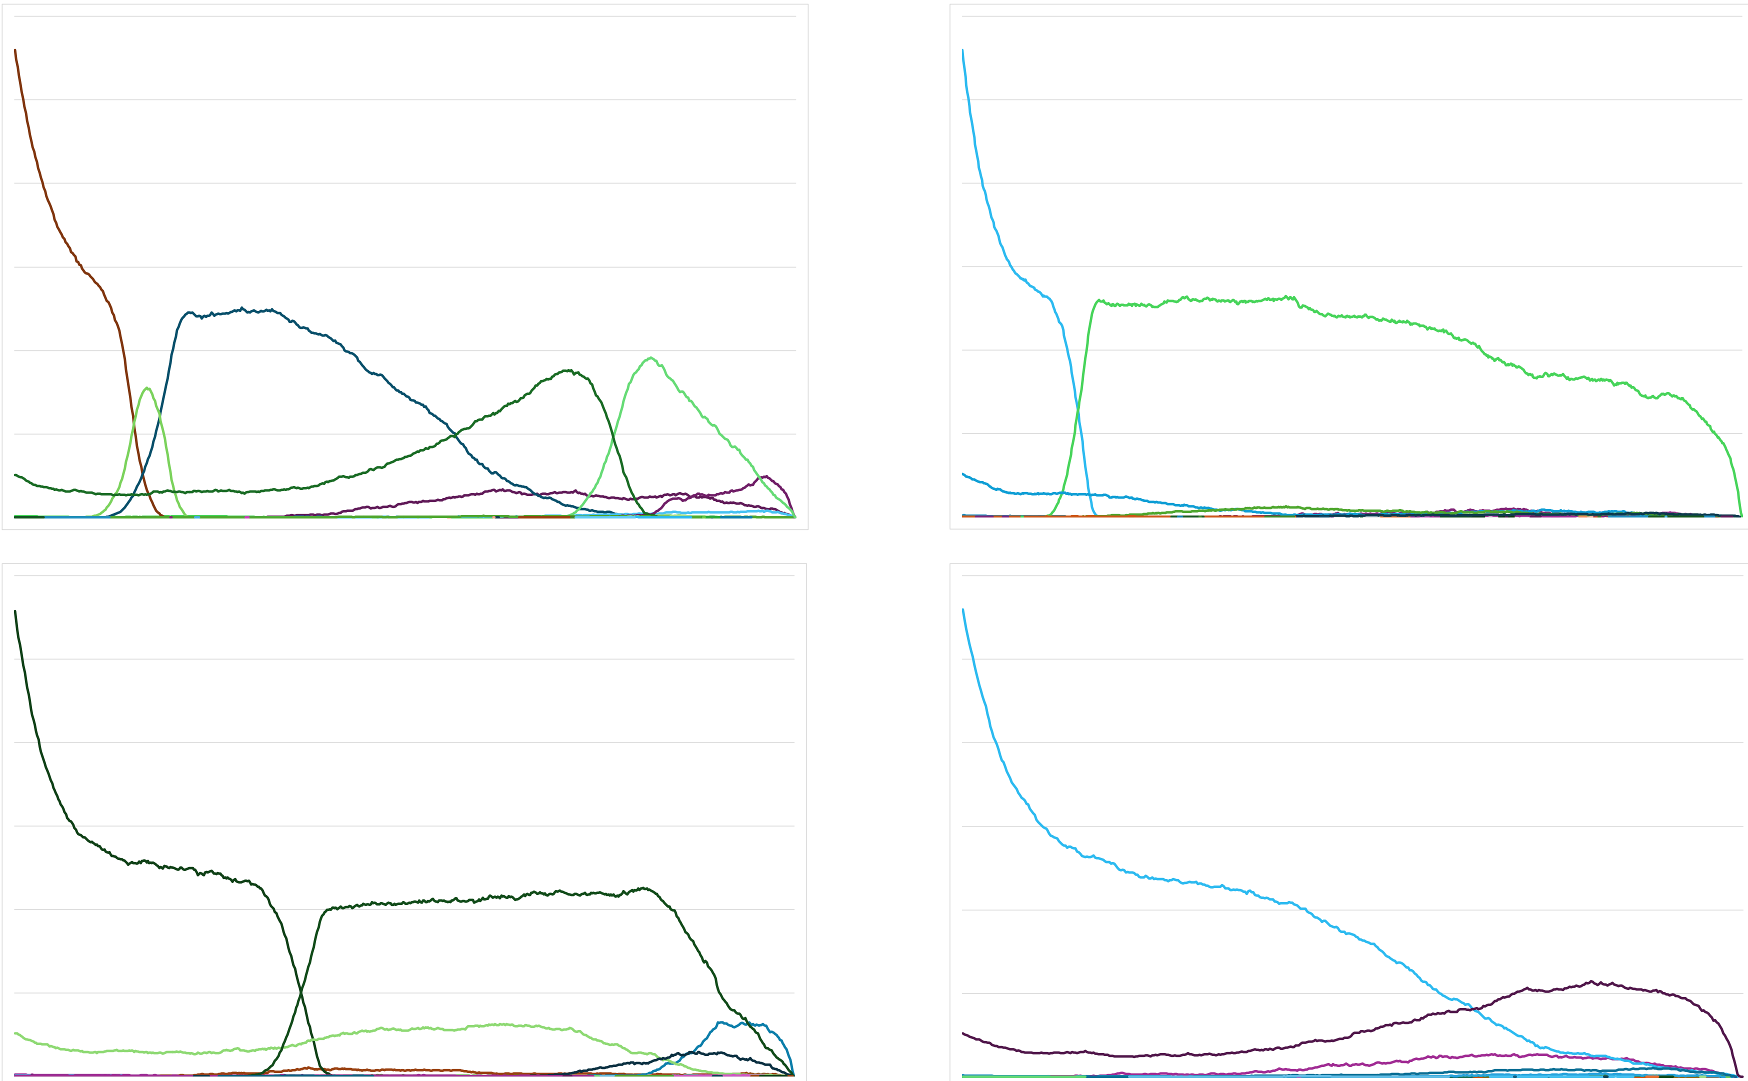

catch level 110 – particular runs *not* extinct – number each species

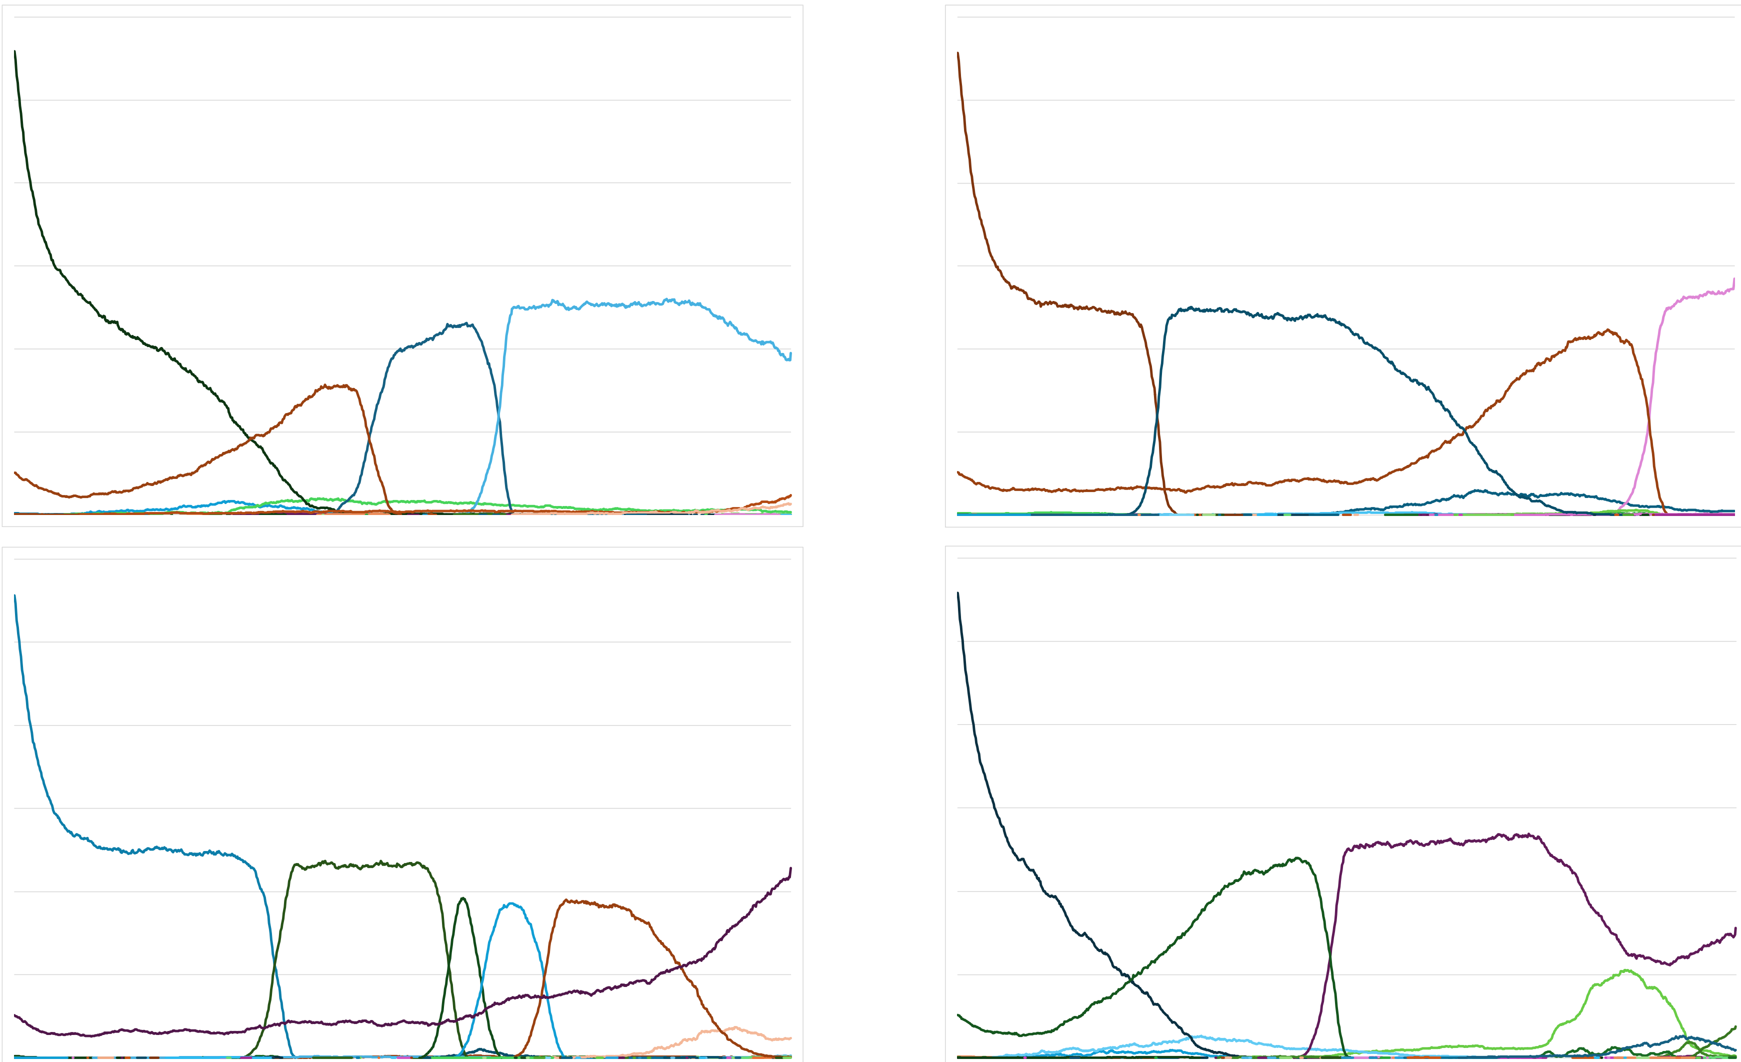

10000 – cont fishing, run 59 (110 catch) diversity & fish numbs (both scaled so max=1)

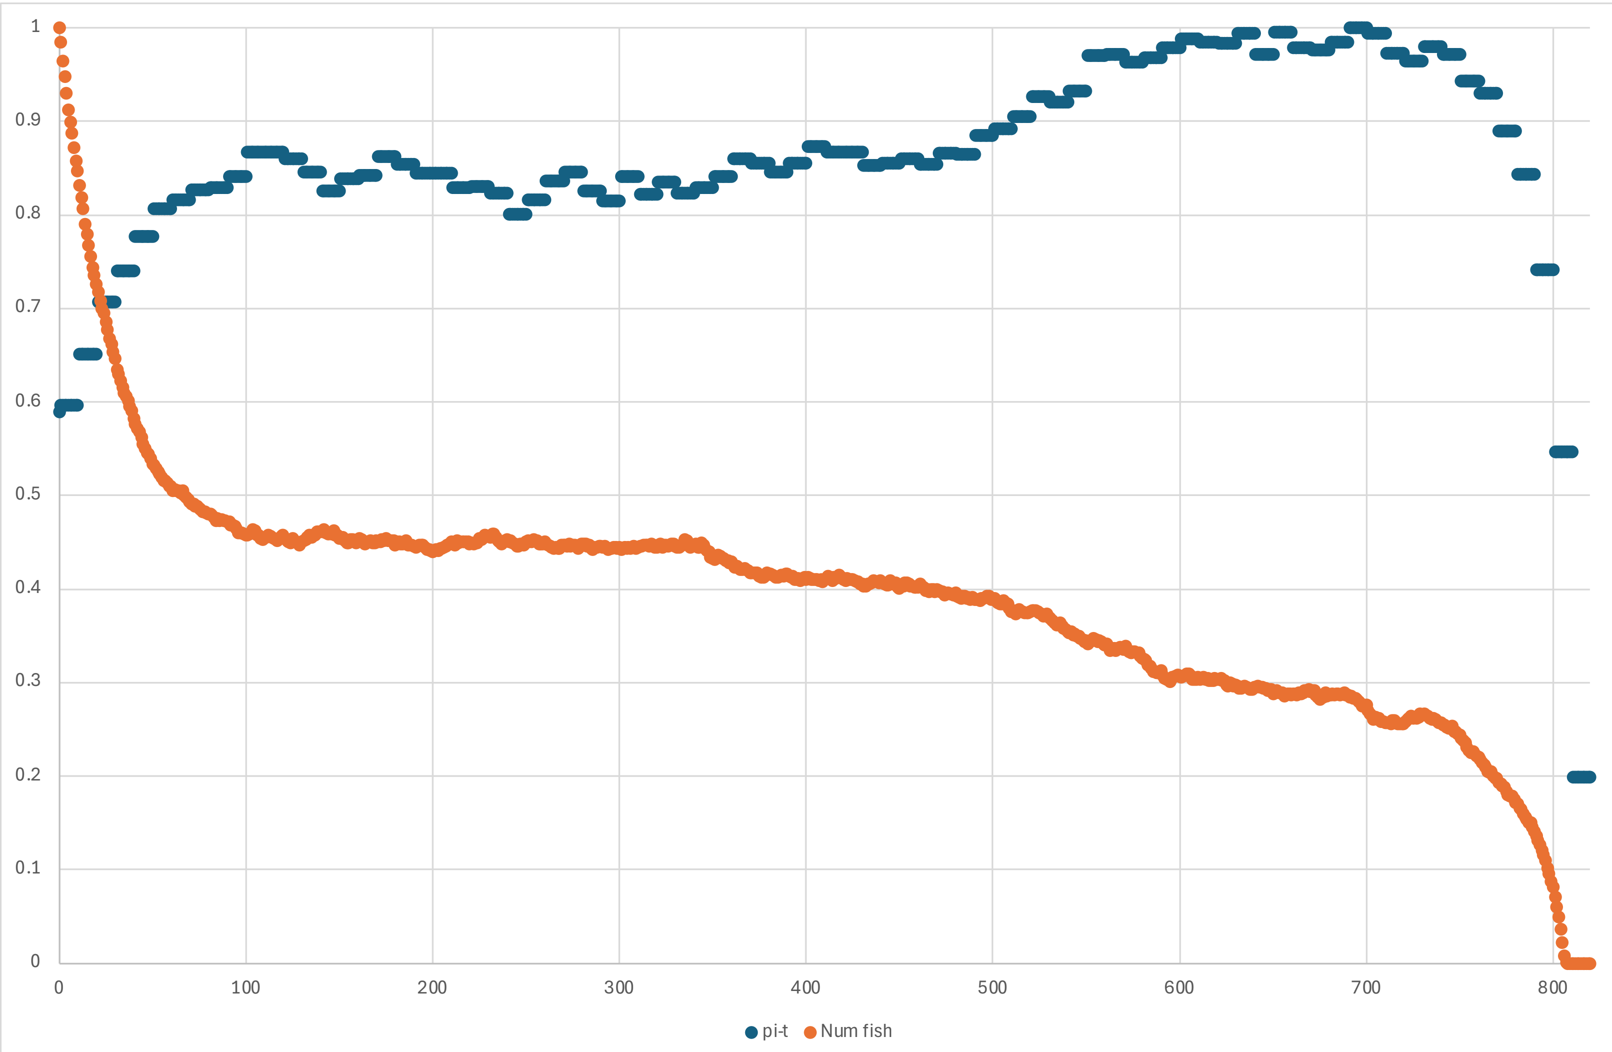

10000 – cont fishing, run 59 (catch 110), species numbers

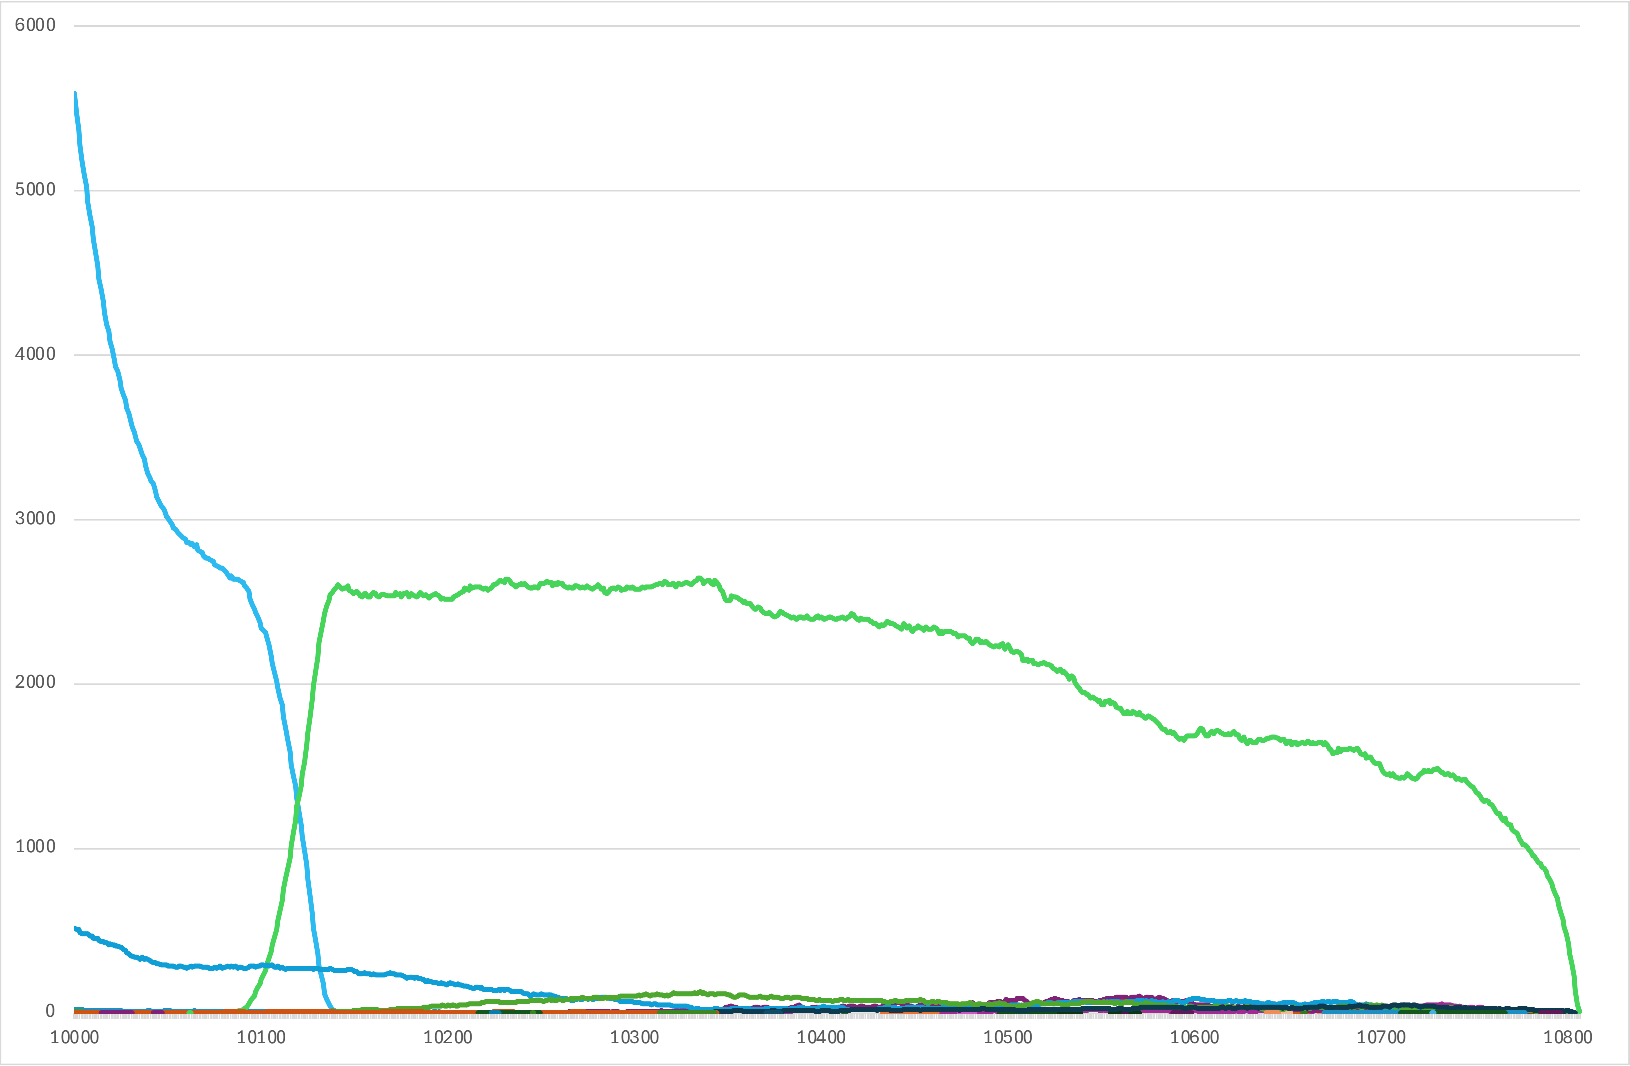

10000 – cont fishing, run 53 (catch 110), diversity & fish numbs (both scaled so max=1)

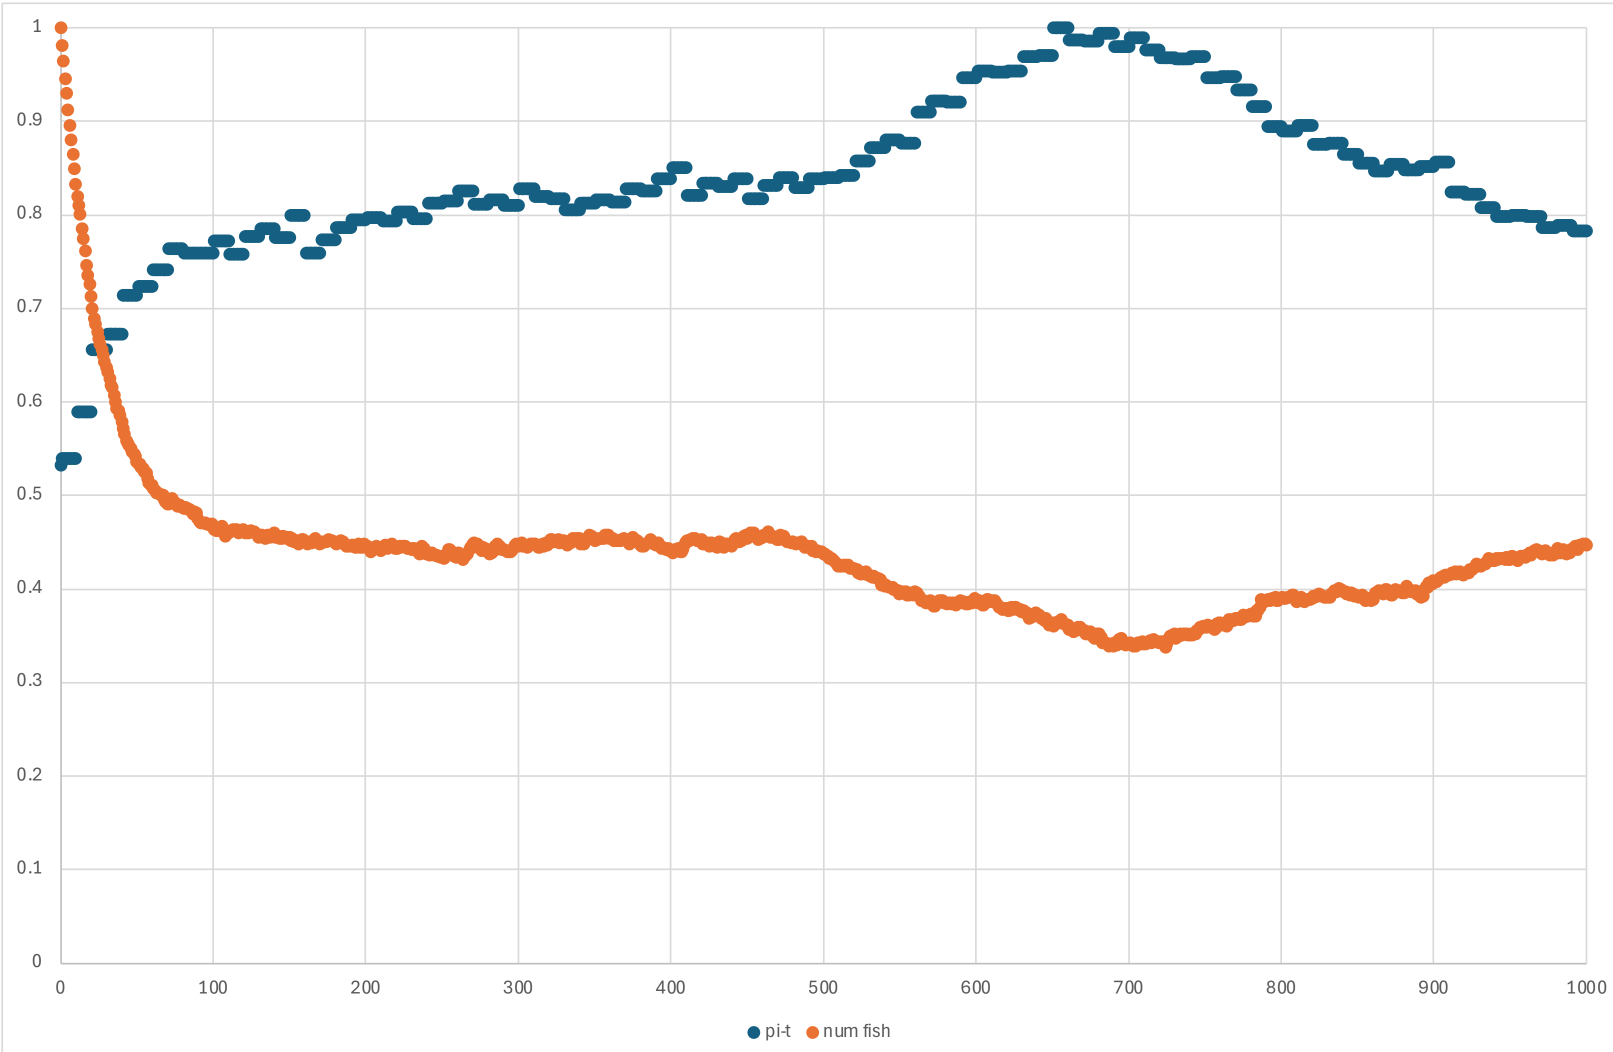

10000 – cont fishing, run 53 (catch 110), species numbers

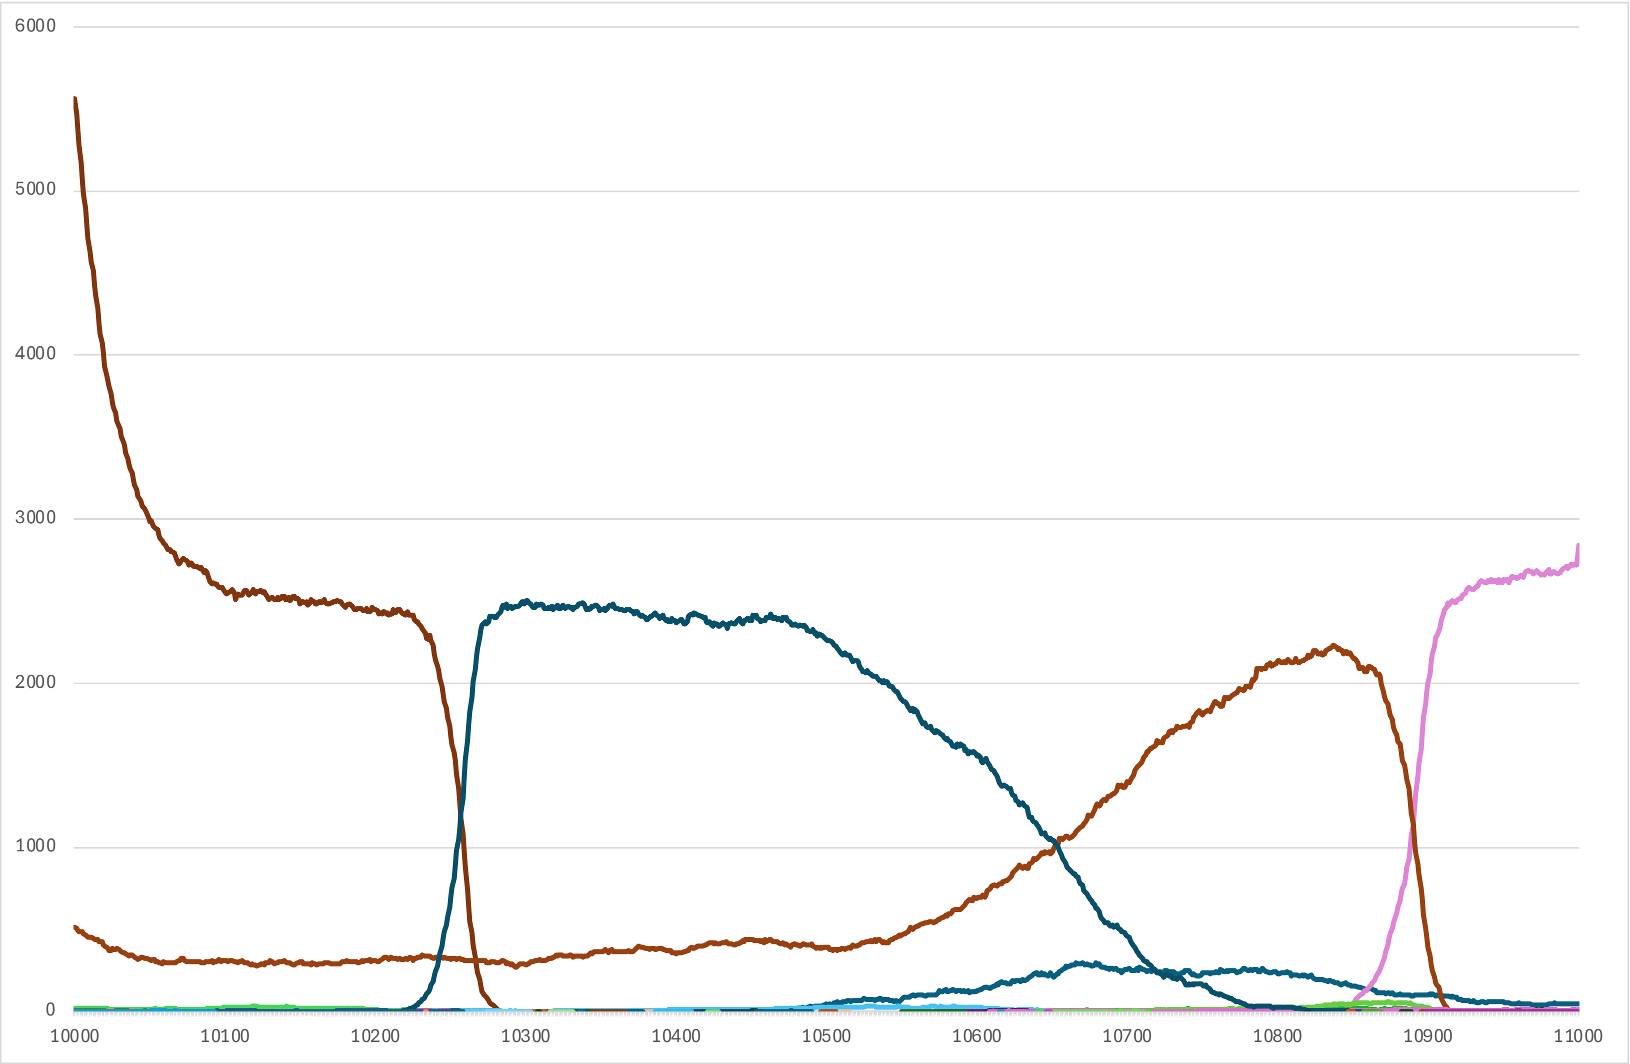

Conclusions

Non-simplistic ecosystems are complex dynamic – never in an ‘equilibrium’

Model suggests:

- that they are vulnerable at catch levels way below the “optimal” level
- Any level of extraction reduces number of species
- Seem to be particularly vulnerable during ecosystem transitions

Possible measures:

- variance in total catches over comparable areas but not over time (too rapid) and only shows close to collapse
- breakdown in number fish vs. diversity relationship

Part 5

Concluding Discussion

Conclusions

- Complex systems can not be relied upon to behave in regular ways
- Often averages, equilibria etc. are not very informative
- Future levels can not meaningfully be predicted
- Simpler models may well make unreliable assumptions and not be representative
- Complex models can not predict probable outcomes but can be part of a risk-analysis

Suggested approach

- Use a variety of inputs (stakeholder, model variants, early signals, theories etc.) to suggest ways in which things might go wrong
- Model these to understand some of the ways these processes might work and interact…
- …and thus identify what needs to be measured in order to get the earliest possible warning that they are emerging
- Do not rely on a few measurements – get lots of data about critical systems
- React fast at many levels, change your mind it these are not working

References

Bithell, M., Chattoe-Brown, E., & Edmonds, B. (2022). The Large-Scale, Systematic and Iterated Comparison of Agent-Based Policy Models. In: Czupryna, M., Kamiński, B. (eds) *Advances in Social Simulation*. Springer Proceedings in Complexity. Springer. https://doi.org/10.1007/978-3-030-92843-8_28

Calder, M. & al. (2018) Computational modelling for decision-making: where, why, what, who and how. *Royal Society Open Science*, https://doi.org/10.1098/rsos.172096

Edmonds, B. (2017) System Farming. In -García-Díaz, C. & Olaya, C. (eds.) *Social Systems Engineering: The Design of Complexity*. Wiley, 45-63. http://onlinelibrary.wiley.com/book/10.1002/9781118974414

Edmonds, B. (2018) A bad assumption: a simpler model is more general. *Review of Artificial Societies and Social Simulation*, 28th August 2018. https://rofasss.org/2018/08/28/be-2/

Edmonds, B. (2019) A Socio-Ecological Test Bed. *Ecological Complexity*, Vol 40 part B, 157-170, https://doi.org/10.1080/13645579.2022.2137921

Edmonds, B. (2020) Good Modelling Takes a Lot of Time and Many Eyes. *Review of Artificial Societies and Social Simulation*, 13th April 2020. https://rofasss.org/2020/04/13/a-lot-of-time-and-many-eyes/

Edmonds, B. (2023) The practice and rhetoric of prediction – the case in agent-based modelling, *International Journal of Social Research Methodology*, 26:2157-170. https://doi.org/10.1080/13645579.2022.2137921

Edmonds, B. & Adoha, L. (2019) Using agent-based simulation to inform policy – what could possibly go wrong? In Davidson, P. & Verhargen, H. (Eds.) (2019). *Multi-Agent-Based Simulation XIX, 19th International Workshop*, MABS 2018, Stockholm, Sweden, July 14, 2018, Revised Selected Papers. Lecture Notes in AI, 11463, Springer, pp. 1-16. https://doi.org/10.1007/978-3-030-22270-3_1

Edmonds, B., Carpentras, D., Roxburgh, N., Chattoe-Brown, E. & Polhill, G. (2024) Delusional Generality – how models can give a false impression of their applicability even when they lack any empirical foundation. *Review of Artificial Societies and Social Simulation*, 7 May 2024. https://rofasss.org/2024/05/06/delusional-generality

Edmonds, Bruce (2014, May 4). "A test-bed ecological model" (Version 1). *CoMSES Computational Model Library*. Retrieved from: https://www.comses.net/codebases/4204/releases/1.1.0/

Gotts, N.M., van Voorn, G.A.K., Polhill, J.G., de Jong, E., Edmonds, B,. Hofstede, G.J., & Meyer, R. (2019) Agent-based modelling of socio-ecological systems: Models, projects and ontologies. *Ecological Complexity*, Vol 40 part B. https://doi.org/ 10.1016/j.ecocom.2018.07.007

The End!

Bruce Edmonds' Publications:
http://bruce.edmonds.name/pubs.html

These slides at: http://cfpm.org/slides

Centre for Policy Modelling: http://cfpm.org

The basic evolutionary model (without "fish" or "humans") is available at:
http://comses.net/model/4204

Bonus Part

Adding in human agents explicitly to explore social-ecosystem co-evolution

An Example of Adding Pretty Simple "Human" Agents

- The agents representing humans are "injected" (as a group) into the simulation into a pre-evolved ecology with complex food webs
- The state of the ecology is then evaluated some time later or over a period of time
- These agents are the same as other individuals in most respects, including predation but "humans":
  – can change their bit-string of skills by imitating others on the same patch (who are doing better than them)
  – might have a higher "innovation" rate than mutation
  – might share excess food with others around
  – might have different migration rates etc.
- Could have many other learning, reasoning abilities

Human migr. rate vs. diversity (all with humans, other entities having 0.1 migration rate)

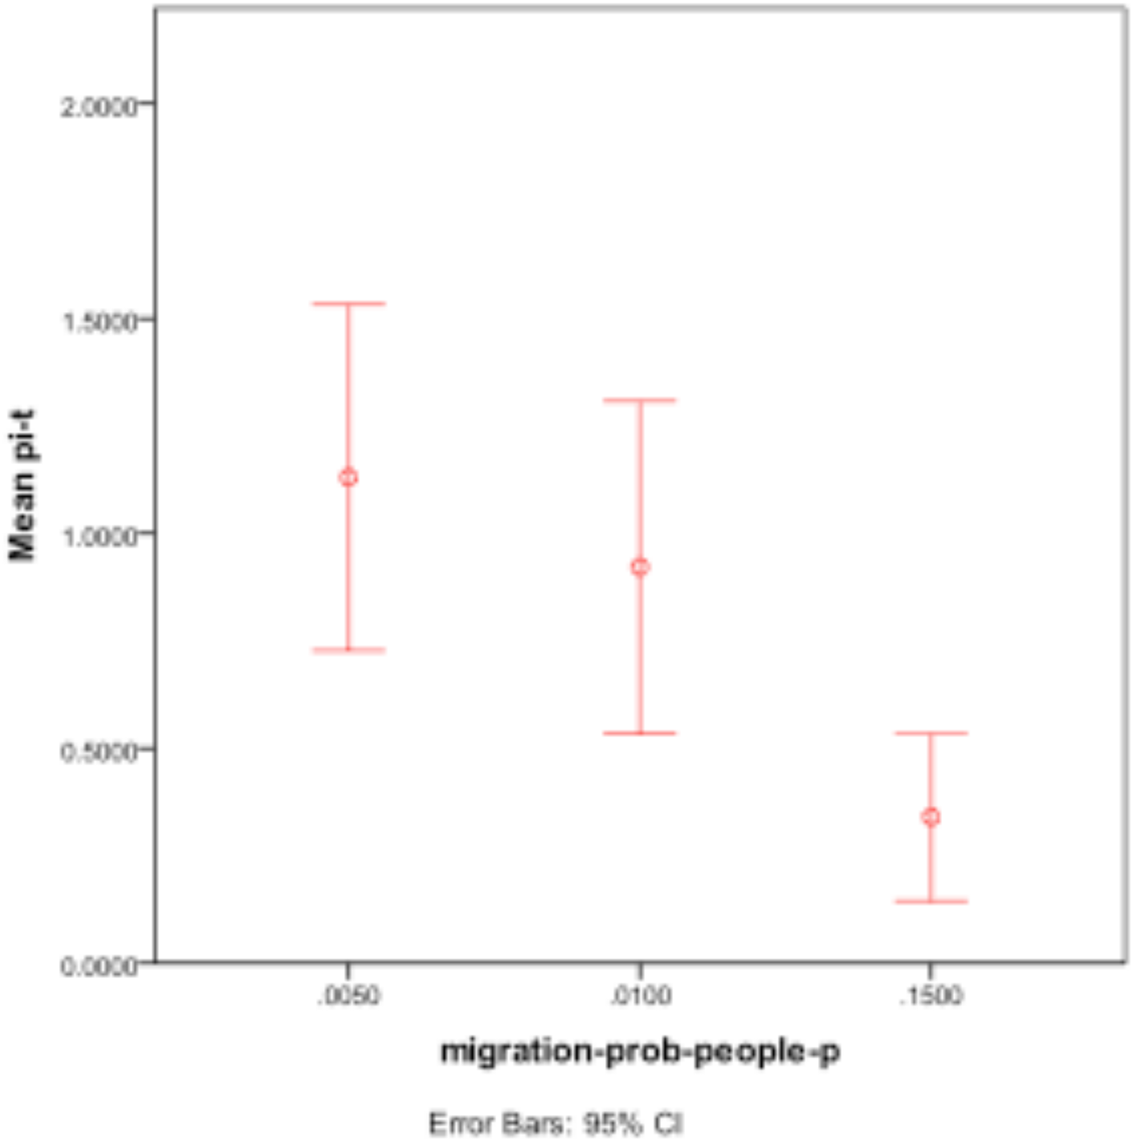

Effect of humans vs. food input to world

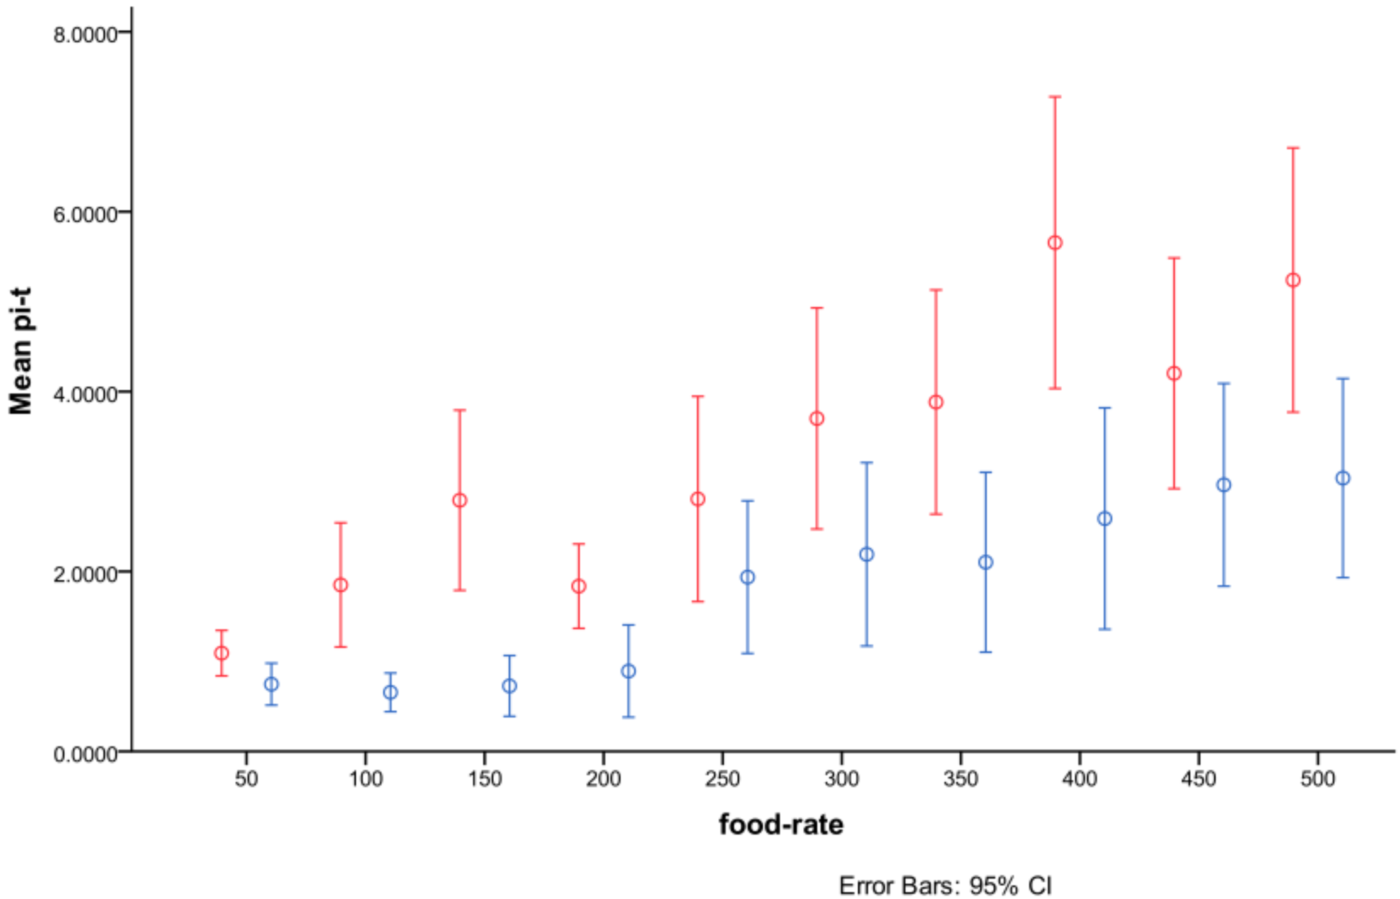

diversity of ecology, blue=with humans, red=without

Effect of humans vs. food input to world

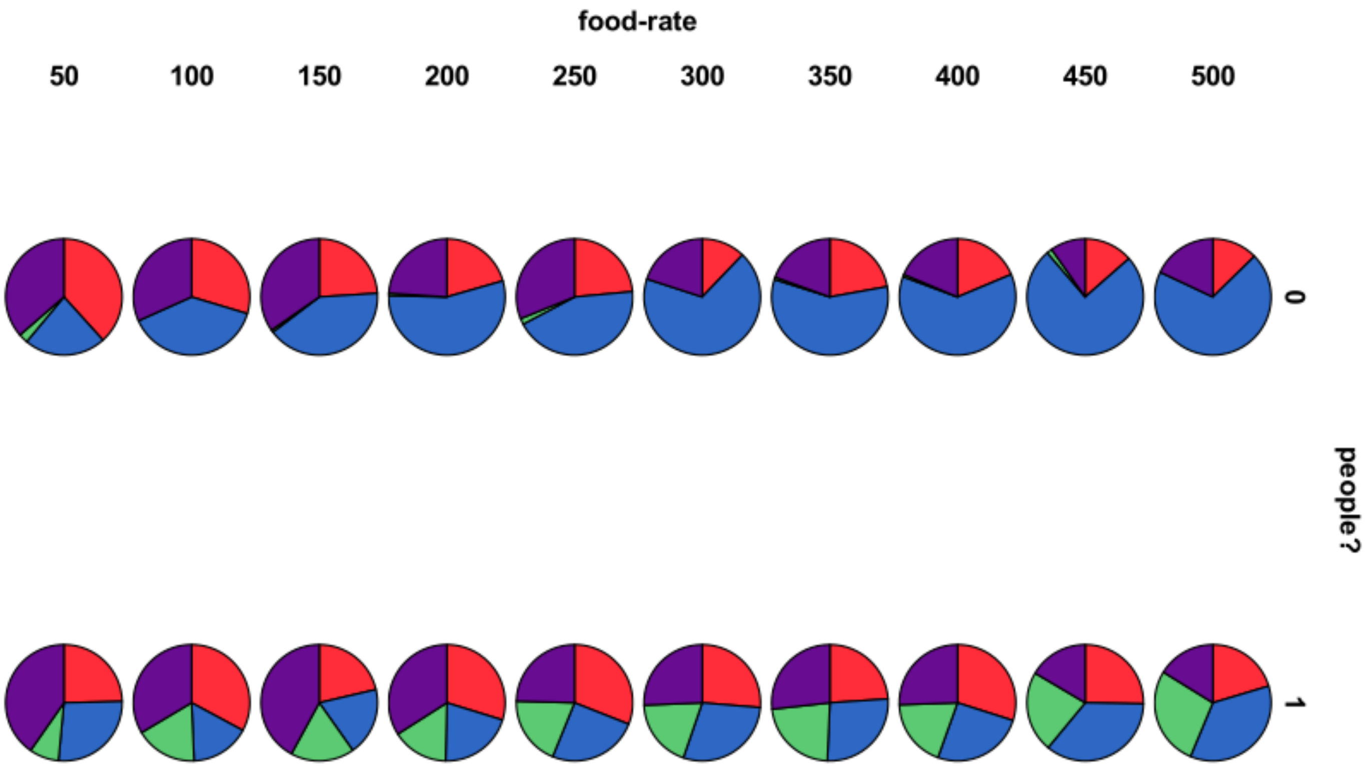

proportion of ecology types, red=plant, blue=mixed, purple=single species, green=non-viable

Migration rate people
vs migration rate others

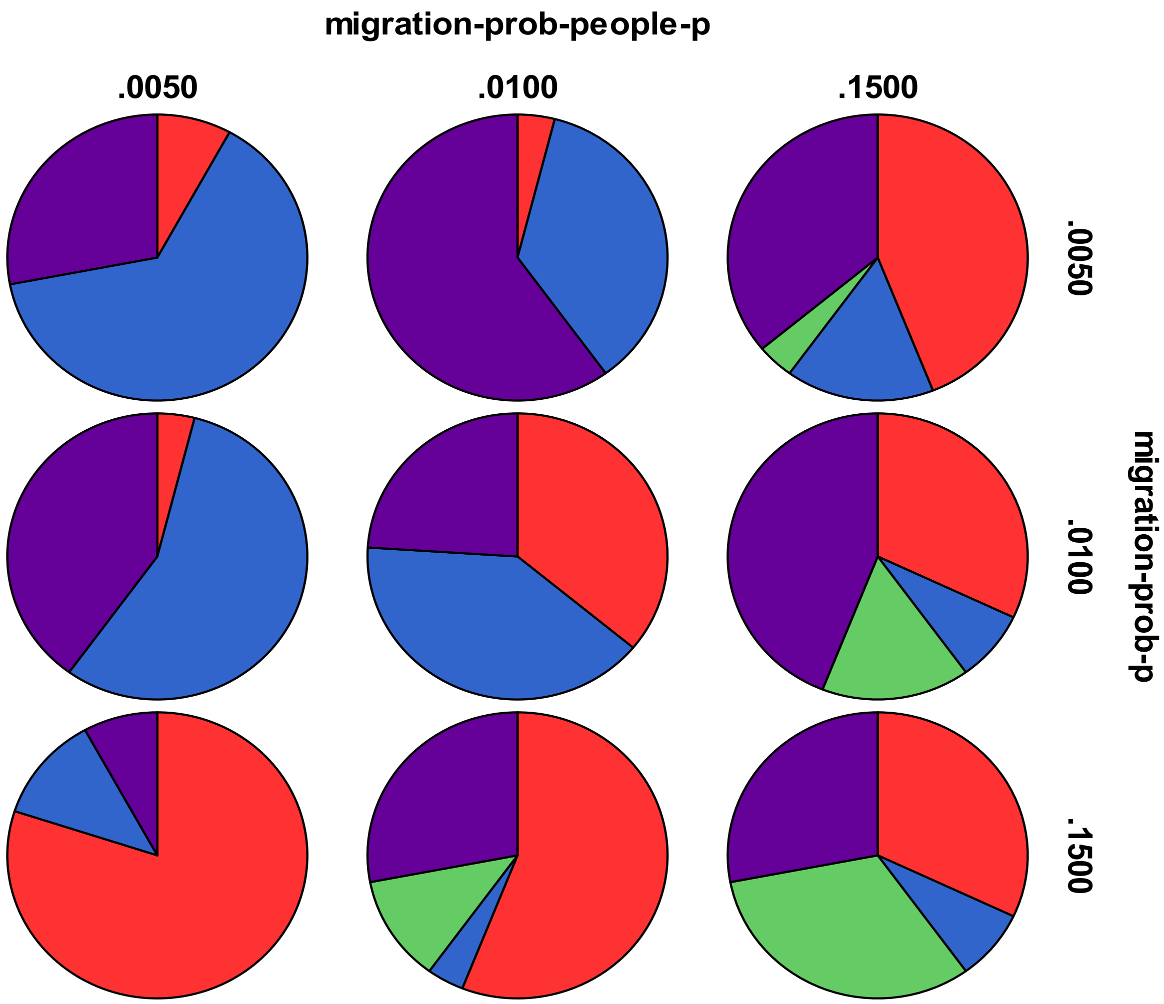

proportion of ecology types 25 simulations each treatment

red=plant, blue=mixed, purple=single species, green=non-viable

Migration (all) vs. food rate (all with humans)

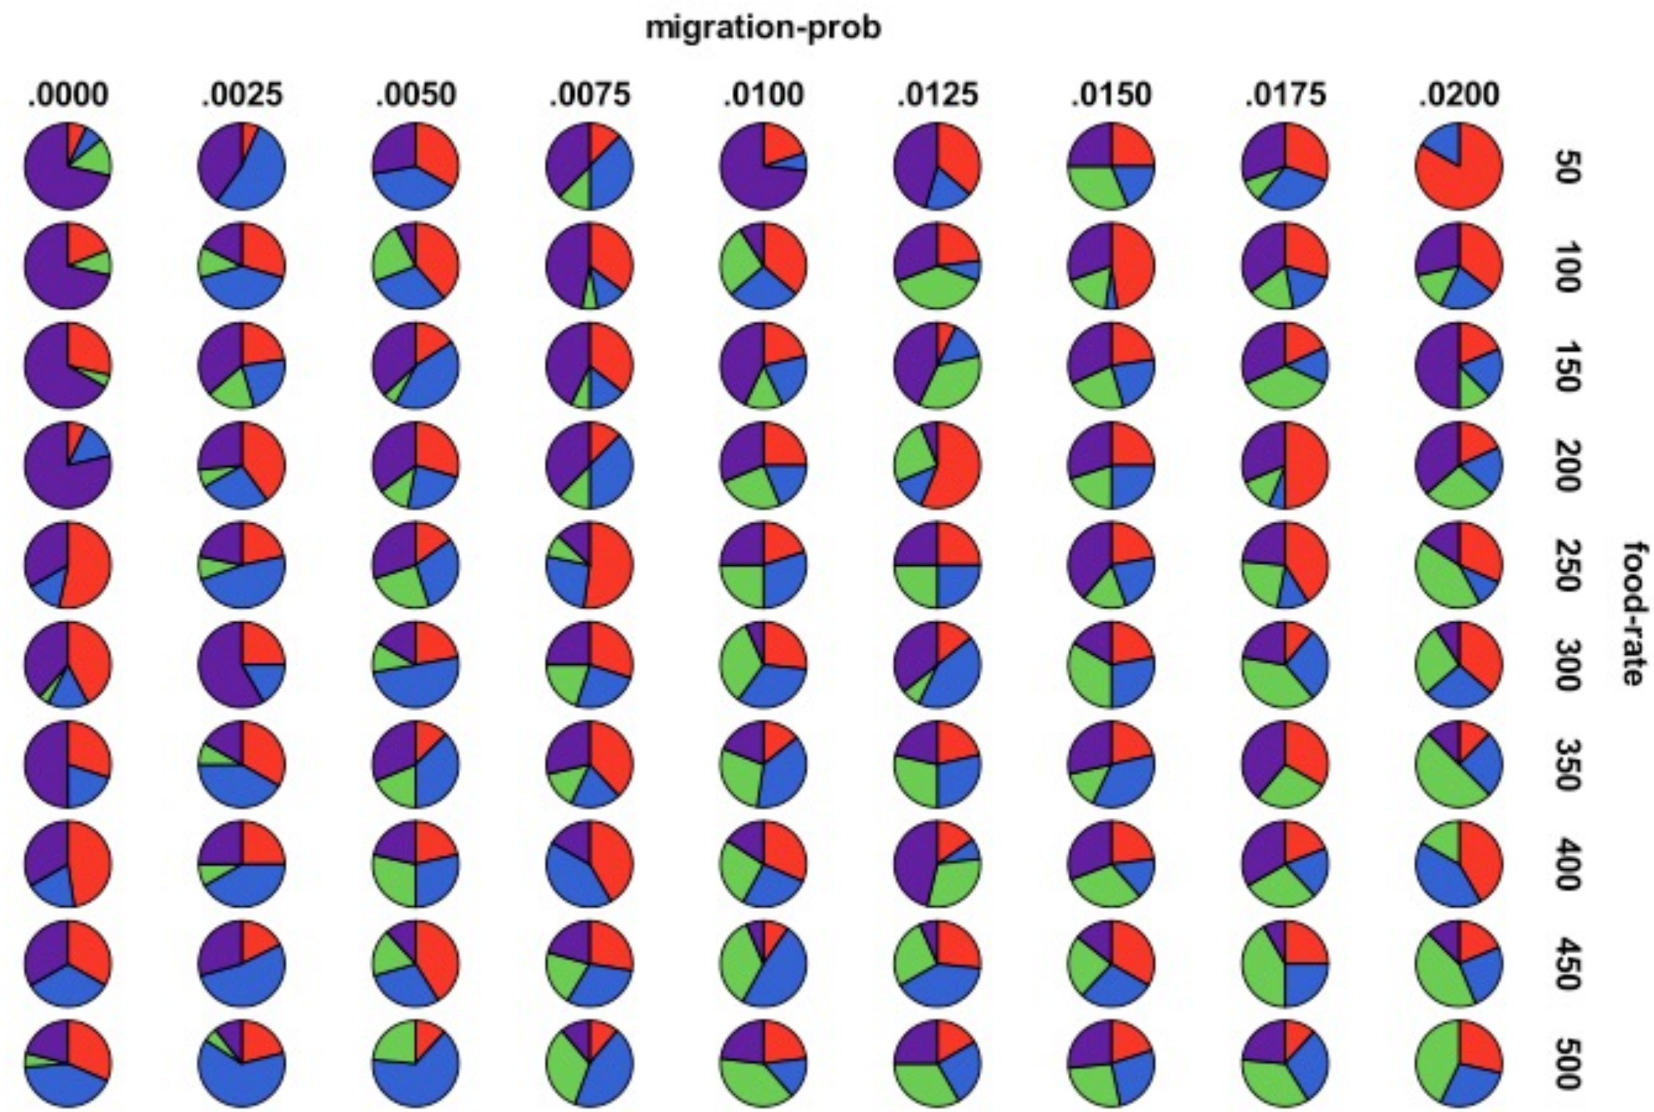

Some observations

- It does not ever get to a ‘steady state’ but is constantly changing and co-adapting
- So approaches to assessing resilience that assume this are not easily applicable
- But we can compare with and without “humans” after a long period of time
- In this model, the way “humans” adapt seems to be more significant that which particular adaption is adopted
- This is only a simple kind of society
- Competition among human groups and their general social evolution is also significant here