To appear in *Foundational Papers in Complexity Science*,
D. Krakauer, ed., Santa Fe Inst. Press, 2022.

Simple Dynamics with Rich Implications [1]

W. Brian Arthur, PARC, Palo Alto & Santa Fe Institute

March 3, 2022

Every so often a paper comes along with a simple setup but a richness of implications. Kristian Lindgren's is such a paper [1]. It appeared thirty years ago and I still find it striking.

Lindgren constructed a computerized tournament where strategies compete in pairs to play an iterated Prisoner's Dilemma game. Players play one-against-one 100 times and each time have two options: "cooperate" or "defect." There's a tension here. If you and your opponent choose to cooperate you'll both do quite well. But if you defect you can do better than that. This seems fine, but if your opponent defects as well (and you don't know in advance what they'll choose) you will both be harmed. Cooperating, defecting, and retaliating are all possible; you need to choose judiciously if you want to do well.

The number of players in the tournament is kept fixed, and each plays a given strategy in the repeated game, a set of fixed instructions for how to act given its own and its opponent strategy's immediate past actions. All strategies play all strategies in a given round or "generation." If a strategy performs well over its encounters, it gets to replicate; if it does poorly, it dies and is removed. Existing strategies can mutate their instructions with a small probability, thus creating new ones. Lindgren adds two important ingredients to the mix. Strategies can make occasional mistakes, so there is some "noise;" and they can occasionally "deepen": mutate by using deeper memory of their opponent's immediate past moves and their own. This allows them to "read" their opponents' moves better, anticipate them, and become smarter.

Evolution here arises in a natural way. We can think of a given strategy as a species, each with its own numbers of players. Species (strategies) occasionally mutate, so they compete for survival and co-evolve in an "ecology" of other species (strategies) competing for survival and occasionally mutating.

The model's dynamics are simple enough that Lindgren can write them as stochastic equations, but these give far from a full picture; we need computation to see how things might unfold over time.

§

In Lindgren's computerized tournament 100 strategies play each other each for 100 moves, in one generation, and he runs the system for 60,000 generations. In each run—I'll call it an experiment—the outcome differs. Typically at the start, simple strategies such as tit-for-tat dominate, but over time, more sophisticated ones appear that exploit them. In some experiments complicated strategies show up early on, in others only later. Figure 1 shows a typical experiment. We see long epochs of dominance by one or a few strategies, with periods of turbulence in between. In this experiment there's a period of Lotka-Volterra oscillations (around time 33,000) between strategies.

What I find interesting is that if Lindgren had stopped the experiment toward the end of a given epoch you might suppose the resulting pattern to be the final one, the final word, be it equilibrium or

[1] This paper is an introduction to Kristian Lindgren's 1991 paper Evolutionary Phenomena in Simple Dynamics, to appear in the same SFI volume.

chaos; yet beyond it the system keeps evolving and changing. What we might accept as a status quo proves to be merely an introduction to the next epoch.

Other experimental runs of course randomly tell different stories. But in spite of variations, Lindgren's world shows consistent phenomena: emergence of mutual support among strategies, exploitation of strategies by other strategies, sudden large extinctions, periods of stasis followed by ones of tumultuous change.

A distinct biological, evolutionary theme emerges. In fact the overall scene, if it resembles anything at all, reminds me of species competition in paleozoological times.

Lindgren's paper appeared in 1991, and three years later he and Mats Nordahl [2] generalized it to a spatial version.

Lindgren's paper wasn't the first to explore prisoner's dilemma tournaments. Robert Axelrod had proposed such a tournament in 1980, where invited scientists sent in computer-coded strategies to compete, and a small but significant collection of studies had followed [1]. What interests me is not where Lindgren's paper fits into this literature, but what it has to say about economics and the doing of economics.

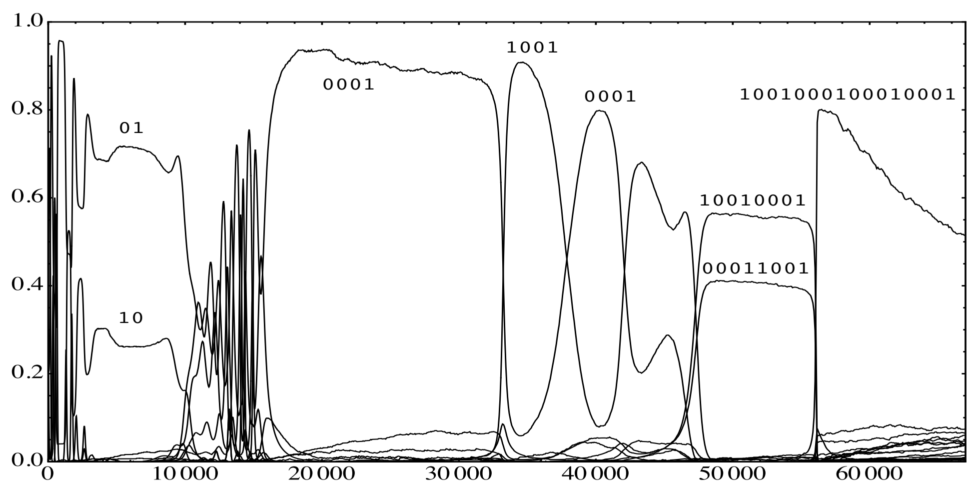

Fig 1. **Prevalence of strategies in a computerized tournament of the Prisoner's Dilemma**

The horizontal axis denotes time, the vertical axis numbers using a particular strategy. The length of labels indicate the memory-depth of strategies, i.e. how many previous moves in the game they take into account.

I would classify this paper as an early example of complexity economics [3,4], though the term wasn't used until 1999. Complexity economics relaxes the main assumptions of standard economics. It views the economy as not necessarily in equilibrium, its agents as not hyper-rational, the problems they face as not necessarily well-defined, and the economy not as a perfectly purring machine but as an always changing ecology of actions, forecasts, and strategies. When we were developing this approach at the Santa Fe Institute in the late 1980s, Lindgren's paper wasn't written, but I've found it useful ever since because it illustrates much of what complexity economics is about: agents updating their beliefs and actions in a setting they mutually create and one that changes as they update. I would also say Lindgren's setup is an early agent-based model, though that term arrived only around the same time as Lindgren's paper.

What is striking about this paper is that from a very few assumptions it constructs a computational world, an ecology of strategies competing, deepening and changing. The world created is an economic one: it is one of competition for scarce resources (here, membership in the set of strategies that survive), though there is no equilibrium reached, no rationality assumed, no "optimality" arrived at.

This world displays a number of properties, all of which are now familiar in complexity economics, but still remarkable. I'll draw attention to three.

1. It shows that significant structural change can arise endogenously. The model assumes no outside shocks, no novel technologies, no new territory discovered. Yet the "structure" or character of the economy constantly changes, and often significantly. Change happens from within as exploration uncovers better strategies, and the system lurches between epochs of relative equilibrium and epochs of turbulence. It might seem that small changes in a strategy should cause small changes, but occasionally they cause large changes, new epochs in this strategy world. Each epoch, if it happened in real life, would have its own character, its own culture, its own explanation. History forms, slowly and endogenously, but not always smoothly. Nature does make leaps.

2. It shows "intelligence" emerging without any built-in assumption of rationality. The strategies automatically get deeper and smarter as the computation progresses. The method by which they get smarter, random trial and survival of successful ones, may appear clunky. But it is the backbone of methods used these days to teach computers to play Go at a grandmaster level. The kind of intelligence arrived at in Lindgren is neither perfectly rational nor necessarily optimal. It is merely pretty good. But it is not forced, and not assumed. It emerges naturally.

3. It shows that computation in economics can sometimes be regarded as "theory." Lindgren's study needs to use computation to track outcomes of his dynamical system. He calls these "simulations," so it would be easy to dismiss this as an exercise in "mere simulation." But his model is rigorously described mathematically; it is a well-specified stochastic process with well-described realizations (particular series of events that result from a run of this stochastic process). As such the process and results are just as mathematical—just as theoretical—as any stochastic process and its implications. Yet the implications are arrived at computationally. The boundary here between computation and mathematics becomes fuzzy. This isn't unusual in modern mathematics. Theoretical systems like the one that generates Mandelbrot's set are undoubtedly mathematical, yet their implications need to be explored computationally. In this sense I would say that Lindgren's system is theory, explored computationally.

Lindgren's paper is now over 30 years old and I still find it fascinating. One reason I gave earlier. From simple assumptions come complicated outcomes and a plethora of lessons. The paper is an allegory, a demo, a parable worth pondering. The other reason is that the outcomes give us—or me at least—a feeling of economic "realism." What emerges is not a world of stasis and perfection, but a world of exploratory trials, "discoveries" and setbacks, endogenous adjustment, openness, long-lasting epochs followed by sudden collapses.

Good theory is like a Chekov play. It sets up characters that interact within a context which contains some tension, some unresolved issue. The interest is to see what plays out and how things play out. We learn not by extrapolating the tendencies of the individual characters, but by seeing them reacting and re-reacting to the situation their actions mutually bring forth.

References

[1] Lindgren, K. Evolutionary Phenomena in Simple Dynamics. In *Artificial Life II*, Addison-Wesley, Ma., 1991.

[2] Lindgren, K. and Nordahl, M.G., *Physica D: Nonlinear Phenomena*, Evolutionary dynamics of spatial games. 75, 292-309, 1994.

[3] Arthur, W. Brian, "Inductive Reasoning and Bound-ed Rationality: the El Farol Problem," *Amer Econ Rev Papers & Proc.*, 84, 406-411, 1994.

[4] Arthur, W. Brian. "Foundations of Complexity Economics." *Nature Reviews Physics*, 3, 136-45, 2021.