

A Study of Severe Disruption in an Artificial Economy

Gustavo Recio¹, Wolfgang Banzhaf² and Roger White³

¹Universidad Carlos III, Madrid, Spain, gustavo.recio@uc3m.es

²Michigan State University, East Lansing, USA, banzhaf@msu.edu

³Memorial University of Newfoundland, St. John's Canada, roger@mun.ca

Abstract

We analyze an artificial economy model designed to handle severe non-equilibrium situations. This agent-based model is intended to allow innovation in the form of new technologies, producers and consumers entering (and leaving) the system. Here we examine a disruption of consumption patterns akin to the economic crisis brought about in the real economy through the corona virus and the following Covid-19 pandemic.

Introduction

Classical economics was developed on the basis of the idea that systems tend toward equilibrium, and that they are thus predictable, since the equilibrium state is in some sense stable. In particular, contemporary microeconomic theory is a static equilibrium theory, formulated as a set of linear equations. The basic question it is designed to answer is how Adam Smith's legendary invisible hand manages to achieve a stable set of prices and quantities of goods in the absence of any global coordination. In this approach change is normally addressed by comparative statics. However, in the real economy innovation in the form of new products and processes means that the economy is continually transforming itself, and so the equilibrium state is a goal that is continually displaced. Consequently, static theory is of little use (Schumpeter, 1961; Gualdi and Mandel, 2019) in helping us to understand the unfolding of major economic events like the long term economic consequences of the current pandemic crisis.

Fortunately, with advances in methods for treating self-organising and far from equilibrium systems, it has become possible to model the self-organising dynamics and structural evolution of economic systems. To that end, we develop an artificial economy as an agent-based model of an economy consisting of producers (firms) and consumers, as well as a set of algorithms that capture the behaviour of these economic agents (Holland and Miller, 1991; Farmer and Foley, 2009). The system undergoes continuous endogenous fluctuations, but also grows and changes because of innovation. In previous work we have studied the ability of such an

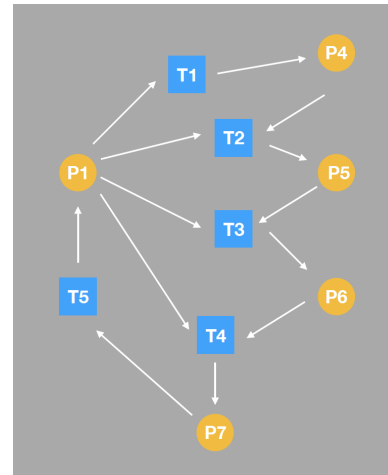


Figure 1: Bi-partite graph representing the structure of a simple economy consisting of technologies T1 ... T5 and products P1 ... P7 produced through these technologies.

economy without innovation to return back to equilibrium values during simulations (Recio et al., 2020a). The general tendency in such a system will be toward increasing complexity, as well as increasing efficiency when less productive technologies are replaced by more effective ones. These fluctuations and trends represent the structural evolution of the system (Recio et al., 2020b; Straatman et al., 2008).

Model of an Artificial Economy

An example of a very simple economy might look like that in Figure 1. Note that we depict here the structure of this economy, with technologies and resulting products depicted, not the actual agents (firms/businesses) that are in possession of a technology and use it and its stock of materials to produce a certain amount of products.

System structure

The structure of this economy can be represented by a pair of von Neumann technology matrices (von Neumann, 1946)

such as that shown in Table 1. The first matrix represents inputs to the production process and the second matrix the outputs, with each row representing a product and each column a technology. Thus, for example, Table 1 shows that 2/3 units of product 1 and 2/3 units of product 4 are used by technology 2 (input matrix) to produce one unit of product 5 (output matrix). We designate product 1 to represent labour inputs, product 2 to represent money, and product 3 a waste product. At present, for simplicity, we restrict technologies to one target product, but in addition every technology produces waste (this will allow us to incorporate aspects of recycling or the circular economy later). Note that there are some products defined as capital goods. For a manufacturing technology making use of these type of products the capital good appears on both input and output matrices. The difference on the input and output quantities of the capital goods indicate depreciation, as shown in Table 1 by the net amount of P5 ($0.00844 = 0.07818 - 0.07036$) used by T5 to produce P8.

The von Neumann technology matrices have two major advantages over the standard Input-Output matrix. First, they allow the possibility of having multiple technologies producing the same product. This is particularly important when modelling innovation, as we do here. Second, they do not aggregate producers into sectors with a resulting loss of specificity. This facilitates an agent-based approach and is important in modelling innovation and evolution, where small individual differences may be the origin of major structural changes.

System dynamics

The execution of economic activity is carried out by a population of agents of two types: producers and consumers. Each producer is characterised by an identification number, location, a stock of products, and a technology. (Currently each producer is limited to a single technology, but this restriction is easily relaxed.) A producer decides how much to produce on the basis of anticipated profit rate relative to a 'normal' rate, with higher anticipated profit rate leading to greater production, as specified in equation 1.

$$Q = Q_{min} + (Q_{max} - Q_{min}) \left(\frac{1}{1 + e^{-\alpha(\pi - \pi_0)}} \right) \quad (1)$$

The estimated profit rate is the ratio of the value of a unit of output to the cost of inputs required to produce it. Thus, as shown in Figure 2, $\pi = 1$ is the break-even point. However, the agent evaluates profit against an expected or normal level, π_0 , where it is frequently the case that $\pi_0 \neq 1$. Equation 1 gives the desired level of production, but the agent may not be able to acquire the necessary inputs to achieve it, either because of insufficient money to buy the inputs, or because they are not available in the system. Also, because the plan is based on observed prices, actual profit once the

Input matrix					
	T1	T2	T3	T4	T5
P1	1	$\frac{2}{3}$	$\frac{1}{2}$	1	0.25
P2	0	0	0	0	0
P3	0	0	0	0	0
P4	0	$\frac{2}{3}$	0	0	0.9
P5	0	0	$\frac{1}{2}$	0	0.07818
P6	0	0	0	1	0
P7	0	0	0	0	0
P8	0	0	0	0	0

Output matrix					
P1	0	0	0	0	0
P2	0	0	0	0	0
P3	0.01	0.01	0.01	0.01	0.01
P4	1	0	0	0	0
P5	0	1	0	0	0.07036
P6	0	0	1	0	0
P7	0	0	0	1	0
P8	0	0	0	0	1

Table 1: Technology matrices of a typical manufacturing economy. The economy being modelled here consists of eight products and five technologies.

production plan is executed will in general differ from what was estimated because prices will have changed since the plan was made. If an agent fails to achieve at least 30 % of the normal profit level, $\pi = 0$, for five consecutive iterations, it will risk being removed from the system.

Each consumer is defined by an identification number, a location, and a stock of consumable products. When more than one consumer good is produced in the system, each consumer has an individual profile of desired consumption. Consumer agents are required to consume a certain amount at each iteration; this comes from current purchases, or to the extent that that is not possible because of insufficient money, from a reserve of consumable goods that each consumer is required to maintain. If the reserve is below the required level, then in addition to current consumption the agent is required to add to the reserve to the extent that cash is available. If a consumer agent is unable to consume the required minimum for five iterations, it also risks being eliminated from the system. Agents do not buy from and sell into a generalised market; rather they search out other agents that can fulfill their needs, or respond to requests from other agents. Each agent interacts with the closest agents that can satisfy their needs. Update is sequential rather than simultaneous. Agents further down the update list may find it impossible to fulfill their requirements either for inputs or sale of output. In order to eliminate this phenomenon as a systematic bias, at each iteration a new update list is established in random order. At each iteration, the price of each product is adjusted on the basis of the ratio of the actual current stock of the product to the desired stock, which is a running average of stock levels over the past five iterations, adjusted for

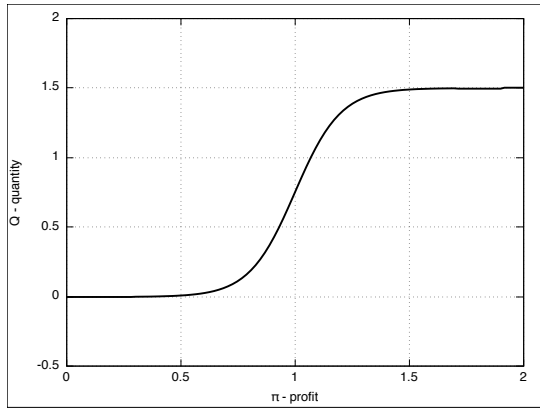


Figure 2: Production function $Q(\pi)$, i.e., the amount of production Q over profit π used to simulate producer agent behaviour. A profit level of 1 represents the break-even point, where costs are equal to revenue. $Q_{min}=0$, $Q_{max}=1.5$, $\alpha = 10$, and $\pi_0 = 1$.

the number of producers. Prices are adjusted according to equation 2 that gives the sigmoidal function shown in Figure 3 when the following parameters are used:

$$x_0 = 0.15, \Delta p_m = 0.2, \Delta p_a = -\left(\frac{\Delta p_m}{1+e^{-\alpha x_0}}\right) + \Delta p_m.$$

$$\Delta p = \begin{cases} \frac{-\Delta p_m}{1+e^{-\alpha(x-(1+x_0))}} + \Delta p_m - \Delta p_a & \text{if } x \leq 1 \\ \frac{-\Delta p_m}{1+e^{-\alpha(x-(1-x_0))}} + \Delta p_a & \text{if } x > 1 \end{cases} \quad (2)$$

System evolution

The model as described so far generates the dynamics by which an unperturbed economic system would move to equilibrium. In that sense it is a fully dynamic analog of standard microeconomic theory, similar to that in (Straatman et al., 2013). However, as previously emphasized, modern economic systems are profoundly innovative; that is, they transform themselves through the addition of new agents, new products, and new technologies (Arthur, 2009). In order to model these transformations we continually add agents, products, and technologies. Every change in the structure of the system will add or remove rows or columns of the von Neumann matrix describing the system. These may be adopted, in which case the system is structurally transformed, or they may be rejected. In the very long run most entities will be replaced by newer ones. Introductions occur irregularly, however, with the per-iteration probabilities of an introduction being specified as modifiable parameter values, so that the effect of varying rates of innovation can be investigated. If the system is not at its equilibrium size in terms of numbers of agents, and only new agents are introduced, it will approach its equilibrium size and oscillate around that size.

If new technologies and products are introduced, these will change the equilibrium system size; therefore, as long

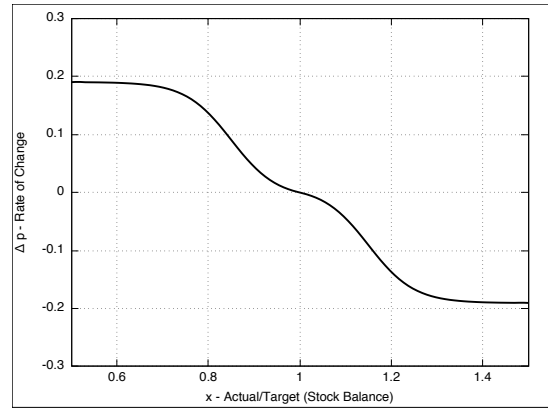


Figure 3: Price adjustment function $\Delta p(x)$ used to adjust prices as a function of the balance ratio x of actual vs. target stocks. Note that in the neighbourhood of the equilibrium point at (1,0) price adjustments vanish.

as introductions continue the system will chase a continually shifting equilibrium. A more significant consequence of the product and technology introductions is that they will force the elimination of some existing agents that are not able to compete; and since producer agents host products and technologies, occasionally, when all agents using a particular technology or producing a particular product disappear, so, too, will the technology or product tied to it.

In terms of system dynamics, the structural evolution driven by these introductions is equivalent to an exogenous perturbation. But now the perturbation is generated by the model itself. A thorough investigation of these phenomena will be published in (Recio et al., 2020b). However, the system may still be subject to exogenous events, just as our actual economy is currently (2020) being perturbed in a major way by the Covid-19 pandemic. Here we examine the effect of such a perturbation on the evolution of the system.

Counterfactual Experiments

One major issue with complex systems like the model of the economy under examination here is that different instantiations of these complex systems tend to behave in completely different ways. That is, whether a new product or technology, or even a producer or a consumer is introduced or removed at iteration i or at iteration j can make a large difference in the further development of the system dynamics. Since our model is abstract and does not have prescribed moments for the introduction of structural change, we have to consider them random events happening at externally determined rates.

In order to understand the effects of such random events on the development of the dynamics, one is left with historical contingencies which are not easy to disentangle. The social sciences have for a long time struggled with such sys-

tems (Wenzlhuemer, 2009), but the phenomenon is present in all fields where system history plays an important role. Even evolutionary and developmental biology, and the life sciences in general fall into this category.

Counterfactual strategies ask the 'what if' question. What if, in the course of a system's dynamics, at a certain point a certain perturbation would have happened? How would the system behaviour deviate from the ordinary behaviour of the system had it not been perturbed in this way?

Here we adopt this strategy by running our economic system simulation for a number of iterations (1,000, in this case), and then run it again from the same initial conditions, with an intervention partway through the dynamics. Thus the dynamics of the two systems will be identical until the perturbation (we use the same seed for the random number generator), but after that both system dynamics and the very structure of the system will be different.

We consider the differences in system behaviour in order to understand what different kinds of responses a complex dynamical system may exhibit. In order not to fall for anecdotal evidence, an artificial economy like the one we examine here offers the advantage that we can repeat experiments and at least cluster the outcomes qualitatively into different behavioural categories.

Central to our attempts is the hypothesis that such a complex system will not be able to return to its previous structural state, but will instead be driven by the disruption (and its consequent system response) to a completely different system structure. This should put to rest the idea that our real economies can return to the status-quo-ante if given enough time.

There is a deep connection between this hypothesis about the behaviour of complex systems and our understanding of the nature of time (White and Banzhaf, 2020). Essentially, what we are saying is that a system like this complex artificial economy, as well as other complex systems, cannot "time-travel" back to its previous state, once it has ventured into a different direction. Due to the contingencies of changes it suffered, it can only travel forward in time as a different system. Thus, the dynamics of such systems is not reversible, a fact that – given the ubiquity of complex systems in our world – sheds some harsh light on the idea of time reversibility that is so prevalent in simple systems traditionally examined in physics.

Experimental Protocol

Inspired by recent events in the real world – the spread of the corona virus in the covid-19 pandemic – with its accompanying severe disruption to economies world-wide, we study a treatment of our systems that mimics severe disruption.

In this investigation we examine the difference between run pairs of 1,000 iterations, a normal (control) run and another one with a treatment of disruption. Under treatment conditions we first run the unperturbed model for 250

Input matrix

	T1	T2	T3	T4
P1	1	$\frac{2}{3}$	$\frac{1}{2}$	1
P2	0	0	0	0
P3	0	0	0	0
P4	0	$\frac{2}{3}$	0	0
P5	0	$\frac{2}{3}$	0	0
P6	0	0	$\frac{1}{2}$	0
P7	0	0	0	1
P8	0	0	0	0

Output matrix

P1	0	0	0	0
P2	0	0	0	0
P3	0.01	0.01	0.01	0.01
P4	0	0	0	0
P5	1	0	0	0
P6	0	1	0	0
P7	0	0	1	0
P8	0	0	0	1

Table 2: von Neumann technology matrices used in current experiments. Products are as follows: 1=labour; 2=money; 3=waste product; 4=raw material; 5-7=intermediate products; 8=consumption good.

iterations and then stop consumer behaviour for 25 iterations so that from iterations 250 to 275 no final consumption occurs. Since final consumption ultimately provides the entire demand that drives the economy, a progressive absence of incentive to produce cascades through the system, starting with consumer good producers and spreading to suppliers and suppliers of suppliers, and so on. While consumption stops immediately and completely, production contracts more gradually, as determined by the price adjustment mechanism (equation 2) and the production function (equation 1). Consumers, who also provide labour, lose income but are not otherwise affected; they retain their reserves of money and consumables and are initially not at risk for being removed from the system. When the intervention period is over at iteration 275 they resume normal activity, but may now experience reduced income and begin to run down their reserves, and some may ultimately be removed.

For this experiment all runs are made starting with the von Neumann matrices shown in Table 2 as initial condition. The per-iteration probabilities for the introduction of new agents, technologies, and products are shown in Table 3. The values shown here provide the growth characteristic of the system. For the executions used in this work, these parameters were arranged to allow general growth of the system with a constant (probabilistic) introduction rate. A removal of entities takes place due to economic indicators, mainly based on ceasing activity – production or consumption.

Fifteen pairs of runs were executed in OCTAVE in two sets. The first run of each pair was unperturbed. The second run was subjected to the perturbation in which final con-

Type of Dynamics	Probability
Producer agent:	0.075
Consumer agent:	0.02
New technology for an existing product:	0.005
New technology for new consumable goods:	0.001
Inactive agents are subject to removal in 5 iterations	
Technologies are subject to removal when not being used	

Table 3: Per-iteration probabilities for the introduction of new agents, technologies and products into the system dynamics, resulting in varying numbers of those entities.

sumption was halted from iteration 250 to 275. Both runs in a pair used the same seed for the random number generator, so that output was identical up to the beginning of the perturbation. After that, while the sequence of random numbers is the same, the two runs of the pair are in different states, so the random numbers lead to different effects.

For reproducibility purposes we divided the runs, one set of 5 runs was done under GNU OCTAVE v3.8.2 on MacOS Catalina, v10.15.3. The other set of 10 runs was done under GNU OCTAVE v4.0.3 on MacOS Sierra, v10.12.3. Our data and the code for all runs discussed here will be made available on GitHub.

Results

Output for each run includes, for each iteration, the number of producer agents using each technology, number of consumer agents, price for each product, profit for each technology, and available supply of each product. The best global indicators of system behaviour are number of consumers and total number of producers. These measures, aggregated over the 15 pairs of runs, show the system to have, on average, a well behaved dynamics: the normal runs show steady growth, but at a decreasing rate, in both number of producers and number of consumers; while the disrupted runs show a sharp drop in the number of producers followed by a steady recovery after the crisis. The number of consumers is not immediately affected by the crisis since their number is frozen, but after the intervention, while the economy is still in recovery mode, the loss of employment income leads to a slower net increase in their numbers compared to the normal (unperturbed) case (Fig. 4).

While the means are very well behaved, the very large standard deviations show that the underlying systems exhibit a wide variety of behaviour. We therefore look next at individual pairs of runs. These can be classified into four categories in terms of the overall behaviour during the run:

Category 1 Ten of the 15 pairs show relatively consistent growth in both the unperturbed system and the perturbed one – with of course a temporary collapse following the perturbation (Fig. 5). In other words, the system is relatively robust since it typically recovers from the severe perturbation.

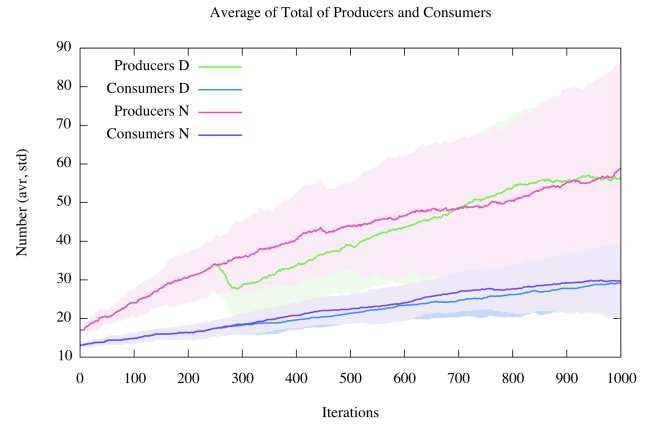


Figure 4: Mean number (lines) and standard deviation (shaded areas) of producer and consumer agents for 15 unperturbed normal (N) and perturbed or disrupted (D) runs. Color shades are indicating 1 standard deviation for each class of runs.

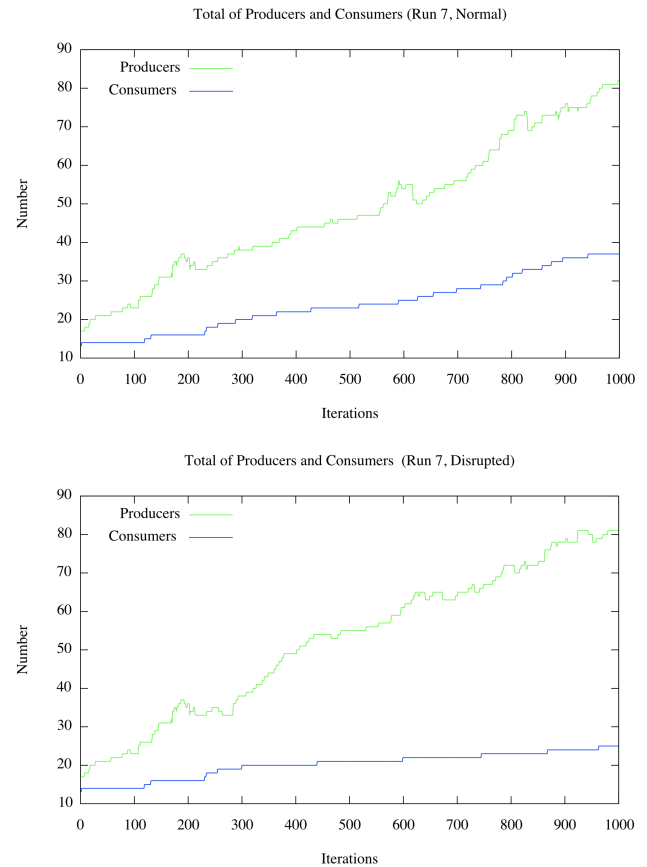


Figure 5: Sample runs of category 1, total number of producer and consumers, normal conditions (top figure); disrupted conditions (bottom figure). Both runs exhibit growth behaviour.

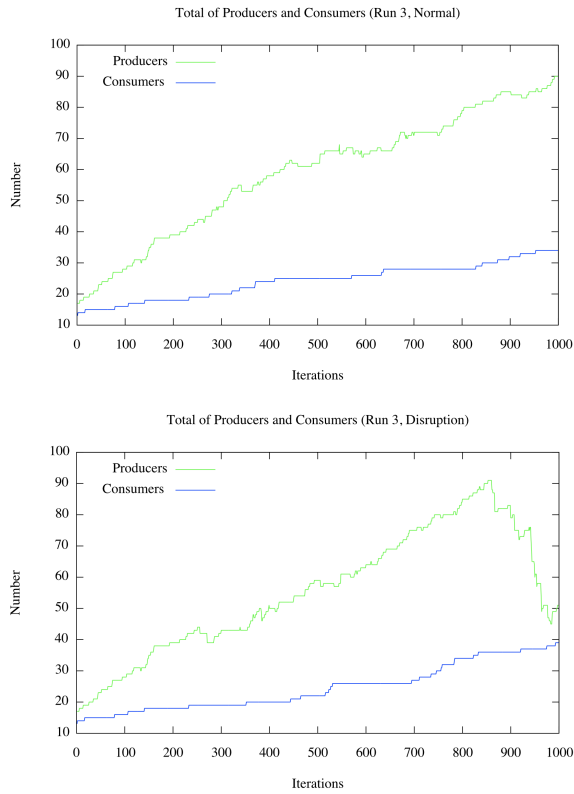


Figure 6: Sample runs of category 2, total number of producer and consumers, normal conditions (top figure); disrupted conditions (bottom figure). Major decline in disrupted system after a long delay.

Category 2 In one case the unperturbed system grows, but when perturbed it ultimately experiences a major decline (Fig. 6).

Category 3 In three cases, the unperturbed system declines or stagnates, whereas the perturbed version grows strongly, i.e., the perturbation in some sense 'cures' an unhealthy system (Fig. 7).

Category 4 In one case both the unperturbed and perturbed systems fail to show consistent growth (Fig. 8).

These variations in the patterns are ultimately due to the stochasticity of the evolutionary component of the model. Some runs see more technologies and more products added than others, and the particular diversity of technologies and products in some systems may constitute a more effective or efficient ecosystem than is found in others. This phenomenon may lie behind the apparent 'cure' provided by the exogenous shock in the category 3 systems – after the crash a less effective economic structure is replaced by a more functional one.

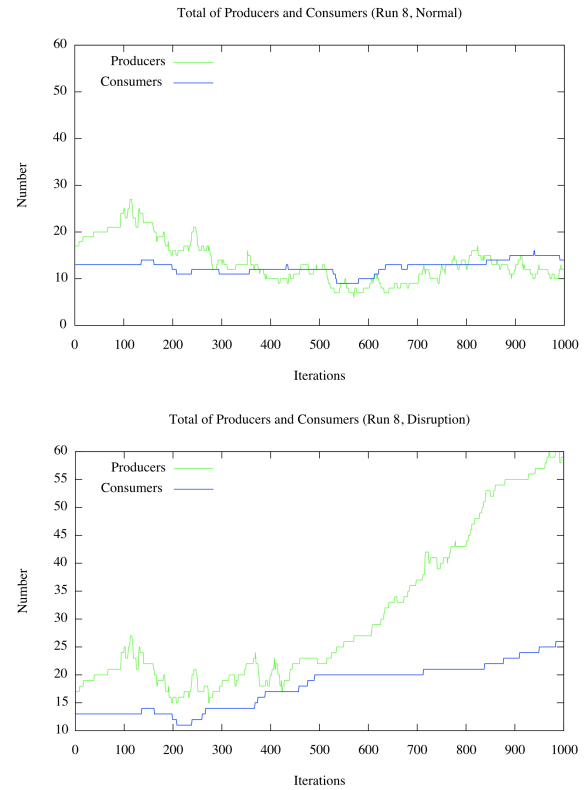


Figure 7: Sample runs of category 3, total number of producer and consumers, normal conditions (top figure); disrupted conditions (bottom figure). Normal run does not exhibit regular growth pattern and instead stagnates or even declines, but is stabilized after the disruption into a consistent growth pattern.

It is not possible at present to thoroughly analyse this phenomenon, but indications may be found in the graphs for the behaviour of individual technologies. In the case of the category 3 example displayed in Fig. 7, we show in Fig. 9 the profit rate for each technology. The most striking difference between the two cases is the number of technologies – 21 in the normal and 12 in the disrupted case. But the number of technologies is not the determining factor in the relative success of the disrupted run; taking all 15 run pairs together there is no relationship between number of technologies and growth, and in another of the category 3 runs the more successful, disrupted, run has more technologies than the normal run. What seems commonly to be the case is that a technology that consistently has a very high profit level depresses the profits of other technologies and this leads to decline of those sectors.

Thus, it seems that serious instabilities have developed in that normal system over the course of its evolution, which were in some sense prohibited by the disruption in the dis-

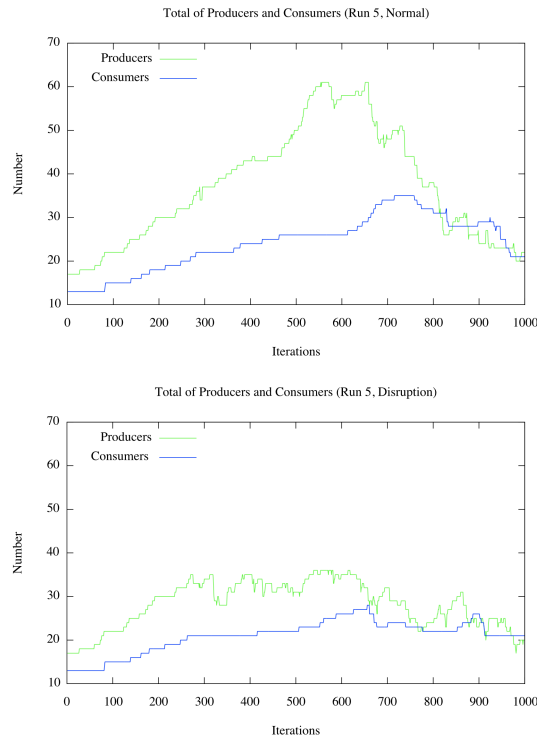


Figure 8: Sample runs of category 4, total number of producer and consumers, normal conditions (top figure); disrupted conditions (bottom figure). Both normal and disrupted runs do not exhibit regular growth patterns.

turbed system. This is corroborated by a look at the graphs for the structure of these economies. Figure 10 shows the structure of the economy *before* the disruption. This is the same structure for both the non-treated control run as for the treated (disrupted) run.

Figure 11 then shows the two different economic structures evolved over a further 500 iterations, the first *without* experiencing a disruption, the second *with* experiencing a disruption. Multiple loops have developed in the former that are capable of driving instabilities. Some of them are conspicuously absent in the latter structure.

It should be noted that the runs we discuss here are all enabling innovation to happen by the introduction and removal of technologies and products. That is the reason why the economic networks can look so different between runs starting from the same structure: These economies have responded differently to the dynamics of agents (producers and consumers) that acted on this structure.

We did a set of control runs in which the ability to produce innovation was turned off. Thus, no new products or technologies were allowed into the system, and in turn, no products or technologies could disappear. These runs never resulted in a serious destabilization of the dynamics, from

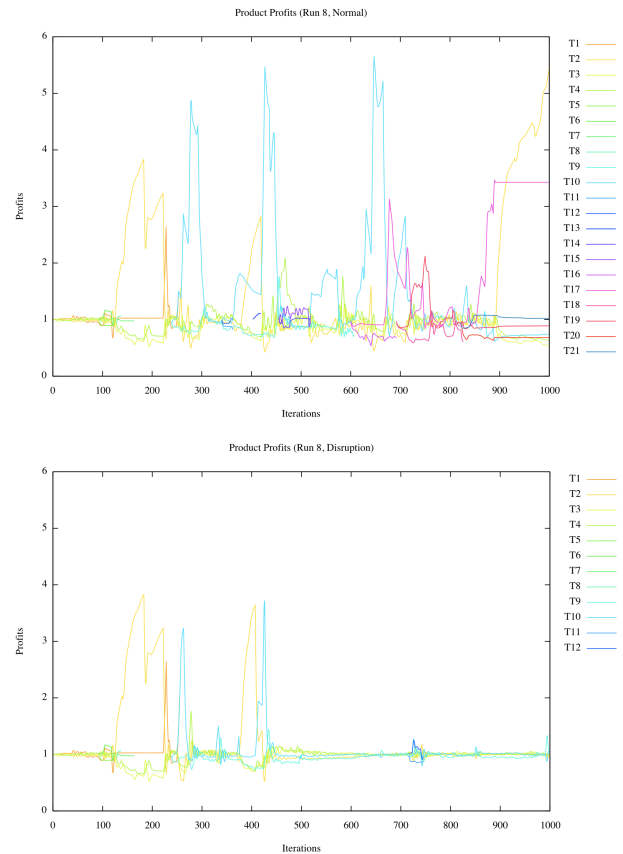


Figure 9: Sample runs of category 3, profit for each product, normal conditions (top figure); disrupted conditions (bottom figure).

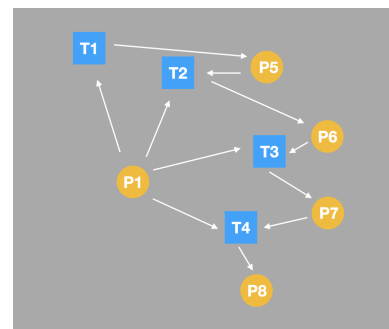


Figure 10: Sample runs of category 3, structure of the economy before potential disruption.

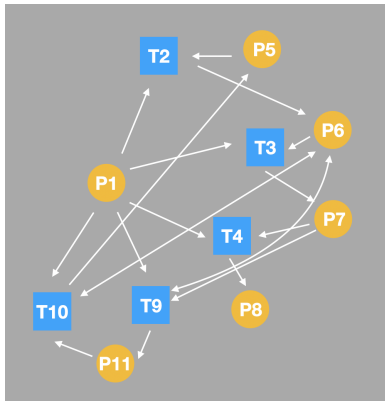
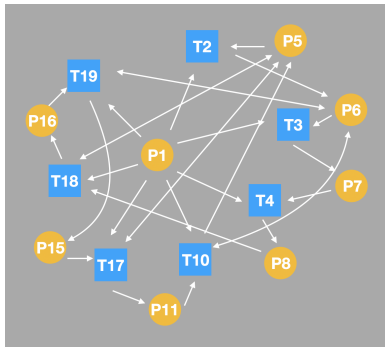


Figure 11: Sample runs of category 3, structure of the economy without (top) and with (bottom) disruption. The situation is depicted at iteration 700, i.e., 500 iterations after the disruption in the second run.

a normal growth economy, nor did they stabilize from an unstable economy. Thus the transitions characterized with categories 1-4 above do not exist in such simpler economic systems.

Conclusions and Generalizations

The overall effect of a severe disruption of a complex system modelling economic activity does not follow the same playbook for all cases. It heavily depends on the structure of an economy whether it recovers from the disruption, or remains unstable and is ultimately doomed. In fact, we found that even under normal conditions some of these artificial economies crash, a testimony to the fact that these complex systems are delicate and need to follow particular trajectories to remain stable.

One important aspect of the instabilities potentially developing in such systems is the necessity of innovation. Without innovation, the structure and the dynamics of these artificial economies remain separated and cannot influence each other. Without such interaction, dynamics is limited to small deviations from an otherwise prescribed way of working.

From a statistical point of view, average behaviour is only of limited use in these systems. The individual behaviour of these systems is so diverse and growing more diverse over time that statements about averages are becoming statistically irrelevant after some simulation time has passed.

The way to examine such diverse behaviours is to attempt to qualitatively cluster systems into different behavioural classes and to compare individual runs with counterfactual experiments.

Acknowledgements

This work was supported by the John R. Koza endowment fund to W.B. at Michigan State University.

References

- Arthur, W. B. (2009). *The Nature of Technology: What it is and how it evolves*. Simon and Schuster, New York, NY.
- Farmer, D. and Foley, D. (2009). The economy needs agent-based modelling. *Nature*, 460:685–686.
- Gualdi, S. and Mandel, A. (2019). Endogenous growth in production networks. *Journal of Evolutionary Economics*, 29(1):91–117.
- Holland, J. H. and Miller, J. H. (1991). Artificial adaptive agents in economic theory. *American Economic Review*, 81(2):365–370.
- Recio, G., Banzhaf, W., and White, R. (2020a). From dynamics to novelty: An agent-based model of the economic system. *Artificial Life*, submitted.
- Recio, G., Banzhaf, W., and White, R. (2020b). From dynamics to novelty: Innovation and complexity growth. *arXiv, in preparation*.
- Schumpeter, J. A. (1961). *The Theory of Economic Development: An Inquiry into Profits, Capital, Credit and the Business Cycle*. Harvard University Press, Cambridge Massachusetts.
- Straatman, B., Marceau, D. J., and White, R. (2013). A generic framework for a combined agent-based market and production model. *Computational Economics*, 41(4):425–445.
- Straatman, B., White, R., and Banzhaf, W. (2008). An artificial chemistry-based model of economies. In Bullock, S., Noble, J., Watson, R., and Bedau, M., editors, *Artificial Life XI: Proceedings of the eleventh International Conference on the Simulation and Synthesis of Living Systems*, pages 592–599. MIT Press.
- von Neumann, J. (1945-1946). A model of general economic equilibrium. *Review of Economic Studies*, 13(1):1–9.
- Wenzlhuemer, R. (2009). Counterfactual thinking as a scientific method. *Historical Social Research*, 34:27–54.
- White, R. and Banzhaf, W. (2020). Putting natural time into science. In Adamatzky, A. and Kendron, V., editors, *From Astrophysics to Unconventional Computation: Essays Presented to Susan Stepney on the Occasion of Her 60th Birthday*, pages 1–22, Cham, Switzerland. Springer.