

Emergence of Autocatalytic Sets in a Simple Model of Technological Evolution

Wim Hordijk^{1*}, Stuart Kauffman² and Roger Koppl³

¹SmartAnalytiX, Vienna, Austria.

²Institute for Systems Biology, Seattle, WA, USA.

³Syracuse University, Syracuse, NY, USA.

*Corresponding author(s). E-mail(s):

wim@WorldWideWanderings.net;

Contributing authors: stukauffman@gmail.com;

rogerkoppl@gmail.com;

Abstract

Two alternative views of an economy are combined and studied. The first view is that of technological evolution as a process of combinatorial innovation. Recently a simple mathematical model (TAP) was introduced to study such a combinatorial process. The second view is that of a network of product transformations forming an autocatalytic set. Autocatalytic (RAF) sets have been studied extensively in the context of chemical reaction networks. Here, we combine the two models (TAP and RAF) and show that they are compatible. In particular, it is shown that product transformation networks resulting from the combinatorial TAP model have a high probability of containing autocatalytic (RAF) sets. We also study the size distribution and robustness of such “economic autocatalytic sets”, and compare our results with those from the chemical context. These initial results strongly support earlier claims that the economy can indeed be seen as an autocatalytic set, and reconcile seemingly opposing views of evolution vs. mutualism in economics.

Keywords: TAP, RAF, Adjacent possible, Combinatorial innovation

JEL Classification: C63 , O31

1 Introduction

From the start, evolution and mutualism have been rival economic “visions” [1] that have often been viewed as incompatible. In our mathematical model of economic evolution, however, mutualism is a product of evolution. The two visions can be reconciled. In the evolutionary vision, the goods produced and the exchanges made both change over time, with each change building on earlier changes [2–7]. It is “evolutionary” because change is cumulative, progressive, or ramifying. In the mutualist vision, the production and exchange of commodities go around and around in a self-replicating circuit such as the textbook model of the circular flow [8–12]. It is “mutualist” because each part of the system is supported by every other part of the system in a mutually reinforcing pattern.

We show that a recent model of combinatorial evolution known as the TAP model brings evolution and mutualism together. TAP stands for *theory of the adjacent possible* [13–15]. As we explain below, the TAP process almost certainly creates mutualistic networks that are “*reflexively autocatalytic and food-generated*” (RAF) [16, 17]. The reconciliation is possible because the economy’s RAFs change over time as the system evolves.

Each vision reflects a common-sense observation about social life which is then transformed into an analytical engine [1]. The evolutionary vision may reflect the common-sense observation of cumulative change in the set of goods available for trade, in how they are produced, and in the institutional environment within which making and trading occur. David Hume gave great attention to these changes in his *History of England*. Later, Adam Smith provided a theoretical account of them. The mutualist vision may reflect the common-sense observation that we are all dependent on one another. Year after year, the farmer grows wheat and sells it to the miller, who grinds it and sells the flour to the baker, who bakes bread and sells it to the farmer to eat. Quesney’s famous *Tableau* is not so far from our simple observation about wheat, flour, and bread.

With the repeated mutualistic cycle of wheat, flour, and bread it might seem that the same thing happens over and over. With the evolutionary emergence of new goods and methods of production, however, new things are happening all the time. It may be understandable, then, that evolution and mutualism have often been seen as inconsistent in economics [18–21]. In biology, however, the *evolution of mutualism* has been a topic of research since at least 1862, when Darwin published *On the Various Contrivances by which Orchids are Fertilised by Insects* [22]. And yet, even in biology it has not been clear how mutualism fits with other biological concepts such as competition and diversity [23]. It is thus of interest within economics and, perhaps, beyond to demonstrate that evolution and mutualism both emerge from one strikingly simple combinatorial model, namely, the TAP model. We provide such a demonstration.

We do not wish to suggest that evolution and mutualism are the only two visions in economics. And we recognize that both visions have generated a

variety of models. But we believe it is significant that RAFs emerge almost certainly from the TAP process. Thus, we reconcile evolution and interdependence in a simple unified model. We place this reconciliation in its scholarly context within economics after developing our technical result.

Technological and cultural evolution are driven by a process of combinatorial innovation [24, 25]. As noted above, a new mathematical model (TAP) was introduced to describe and study such a process [13]. This model accurately reproduces the well-known “hockey stick” phenomenon observed in economic growth: a long period of slow growth followed by explosive growth in just a short time [26, 27].

The TAP model is based on the assumption that a simple cumulative combinatorial process underlies this pattern of technological and cultural evolution. At its core is the following equation:

$$M_{t+1} = M_t + \sum_{i=1}^{M_t} \alpha_i \binom{M_t}{i}, \quad (1)$$

where M_t is the number of different types of goods in an economy at time t , and α_i is a decreasing sequence of probabilities (i.e., real numbers between 0.0 and 1.0).

The main idea of this equation is that new goods (or tools, or artifacts) are created by combining any number i of already existing goods. Arbitrary combinations of i existing goods have a small probability α_i of resulting in a useful new one. The TAP model was recently studied both theoretically and with computer simulations [13]. Note that this equation represents a deterministic version that does not guarantee M_t to be an integer value, and it only serves to convey the general idea behind the model. A stochastic implementation that does guarantee integer values is presented below.

Alternatively, it has been suggested that an economy can be seen as an instance of an autocatalytic set [14, 15, 28, 29]. The notion of an *autocatalytic set* originally comes from chemistry, where it is informally defined as a chemical reaction network in which the molecules mutually catalyze (i.e., speed up, or facilitate) each other’s formation, and which can sustain itself from a given set of basic building blocks, or “food set” [30, 31]. This notion was formalized and studied in more detail as reflexively autocatalytic and food-generated (RAF) sets [16, 17], and has also been implemented with real molecules in the lab [32–35] and shown to exist in the metabolic networks of prokaryotes [36–38].

A chemical reaction is basically a *transformation* of a set of input molecules (reactants) to a set of output molecules (products). Similarly, in an economy, input goods are transformed into output goods, such as wood and nails (input) being transformed into tables (output). Furthermore, and again similar to chemistry, some of these goods can act as “catalysts”, in that they facilitate the production (or transformation) of other goods. Examples are hammers, conveyor belts, and computers. These goods are not used up in a production process, but facilitate the production of other goods such as tables, cars, or

animated movies. Yet, such catalysts are themselves products of the same economy. As such, an economy can be seen as a product transformation network in which the goods mutually catalyze each other’s formation, sustained by a basic food set of raw materials. In other words, an economic autocatalytic set.

An interesting question is whether these two alternative views of an economy, as a process of combinatorial innovation on the one hand and as an autocatalytic set on the other, are somehow compatible. We show that the answer to this question is a resounding *Yes*, and that this positive answer reconciles evolution and mutualism. Using the original TAP model, we extend it with the assignment of (random) catalysis, as was hinted at recently [39]. We then show that autocatalytic (RAF) sets have a high probability of emerging in product transformation networks that result from a process of combinatorial innovation. Additionally, we also study the size distribution and robustness of these economic RAF sets, and discuss connections with and implications for existing economics models.

2 Methods

An implementation of a stochastic discrete-time version of the TAP model based on an earlier version [13] is used here, with $\alpha_i = \alpha^i$ (i.e., α to the power i , for some given value of α). The creation of a new good from a combination of already existing “parent” goods is then interpreted as the introduction of a new product transformation, where the parent goods are the inputs and the new good is the output. In addition, each newly created good is assigned as a catalyst to the already existing product transformations with a given catalysis probability p_c . Similarly, a new product transformation is catalyzed by any of the already existing goods also with probability p_c . The resulting model is implemented as described in Algorithm 1.

Note that an upper limit K on the number of parents is set for numerical and computational reasons. It was already shown earlier that this does not significantly affect the overall behavior of the TAP model [13]. Also note that it is possible that r_i in line 7 of the algorithm is assigned a value of zero, in which case the *for*-loop of lines 8–16 is not executed (in most programming languages, a loop from 1 to 0 is simply not executed).

At each time step t the M_t existing goods thus form a product transformation network with a food set consisting of the M_0 initial goods, where each product transformation consists of a good being produced from its parents. In addition, the goods catalyze each other’s production according to the catalysis probability p_c . Note that each product transformation can thus have none, one, or multiple catalysts, depending on the random catalysis assignments (with probability p_c). Similarly, each good may catalyze none, one, or multiple product transformations.

Figure 1 shows a simple example of such a product transformation network, resulting from a TAP process, in a graph representation. The black dots represent the different types of goods created at each time step, and the white

Algorithm 1 TAP with catalysis

Require: $M_0, K, \alpha, \mathbf{M}, p_c$

```

1:  $t \leftarrow 0$ 
2: Create  $M_0$  initial goods labeled  $1, \dots, M_0$ 
3: while  $M_t < \mathbf{M}$  do
4:    $t \leftarrow t + 1$ ;  $M_t \leftarrow M_{t-1}$ 
5:   for  $i = 1, \dots, K$  do
6:      $s_i \leftarrow \alpha^i \times \binom{M_{t-1}}{i}$ 
7:      $r_i \leftarrow \text{Poisson}(s_i)$ 
8:     for  $j = 1, \dots, r_i$  do
9:       Create a new good  $x$  labeled  $M_t + 1$ 
10:      Select  $i$  random “parents” for  $x$  from  $1, \dots, M_{t-1}$ 
11:       $M_t \leftarrow M_t + 1$ 
12:      for  $y = 1, \dots, M_t$  do
13:        With probability  $p_c$  assign  $x$  as catalyst to production of  $y$ 
14:        With probability  $p_c$  assign  $y$  as catalyst to production of  $x$ 
15:      end for
16:    end for
17:  end for
18: end while

```

boxes represent their product transformations. Solid arrows indicate inputs and outputs of product transformations. Dashed gray arrows indicate which goods catalyze which product transformations (assigned randomly).

Note that this graph representation of a product transformation network is similar to that of a chemical reaction network, where dots represent molecule types and boxes represent chemical reactions [40], and with catalysis added as dashed gray arrows. Therefore, one could ask whether this network contains a subset that is reflexively autocatalytic and food-generated, or RAF [17]. Formally, a RAF set is a set R of chemical reactions (or, in this case, product transformations) and the molecule types (here: goods) involved in them such that:

1. Each reaction (product transformation) in R is catalyzed by at least one of the molecules (goods) involved in R .
2. Each molecule type (good) involved in R can be produced from the food set through a sequence of reactions (product transformations) from R itself.

The food set is a subset of molecule types (goods) that is assumed to be available from the environment (in this case the M_0 initial goods).

Figure 2 presents an illustrative economic RAF, where the food set consists of wood, metal, and bricks. These can be combined in different ways to produce hammers, wheelbarrows, and brick ovens. Each of these products can also act as a catalyst: a hammer speeds up the production of wheelbarrows (without being used up in the process), the wheelbarrow speeds up the transportation of bricks to make ovens, and the oven speeds up the production of hammers by




Fig. 1 An example product transformation network resulting from a TAP process, with random catalysis assignments (dashed arrows).

melting the metal so it can be molded into the right shape. This “economy” as a whole (within the solid oval) thus forms an autocatalytic set. Once this set is established, it may allow the emergence of new goods such as nightstands and pizzas that were not possible before (the dashed ovals). We discuss other economic RAF examples in the discussion section below.

An efficient computer algorithm exists to find such RAF sets in arbitrary chemical reaction or product transformation networks, or determine that no such subset is present [16, 41, 42]. The algorithm works by iteratively removing reactions (product transformations) that by definition cannot be part of a RAF set, until no more reactions can be removed. This RAF algorithm actually finds the *largest* RAF (maxRAF) present in the network, which in the example in Figure 1 is the entire network itself.

However, such a maxRAF could consist of the union of several smaller RAFs (subRAFs), including minimal, or *irreducible*, RAFs (irrRAFs). Indeed, there are several (smaller) subRAFs within the maxRAF of Figure 1, such as the two product transformations forming the goods $\{5, 6\}$, or the four product transformations forming $\{5, 6, 7, 8\}$. Such subRAFs and irrRAFs can be identified by (repeatedly) applying the RAF algorithm after removing one or more (random) elements from the maxRAF.

Recently, some fundamental theorems and results on the combination of the TAP and RAF models were presented, but in the context of chemical evolution [43]. Here, we investigate this model combination in more detail, but




Fig. 2 An illustrative economic RAF set. From [14].

re-interpreting it explicitly in the context of technological evolution. In particular, using the TAP simulation model including catalysis assignments as described in Algorithm 1, many runs are performed using different catalysis probabilities p_c . The RAF algorithm is then applied to the product transformation networks resulting from these TAP processes to see how often a RAF set emerges, and how large they are.

In the results presented below, most TAP model parameters are fixed as follows:

- $M_0 = 10$
- $K = 4$
- $\alpha = 0.01$
- $\mathbf{M} = 1000$.

Rather than running a TAP simulation for a fixed number of time steps, a run will end once a threshold of $\mathbf{M} = 1000$ goods is crossed. So, different runs might terminate at different time steps, and result in a different number of final goods M_t (but with $M_t \geq 1000$), due to the stochastic nature of the simulations. The catalysis probability p_c is then allowed to vary between batches of runs, to see how it influences the possible emergence and sizes of RAF sets.

Jain and Krishna (J-K) provide an alternative model of autocatalytic sets, which also combines mutualism and evolution [44]. The J-K model represents a so-called *elementary* RAF, where all products are directly formed

from the (implicitly assumed) food set [45]. These products then act as catalysts for each other’s formation (mutualism), and a form of evolution results from random changes in these catalytic connections. Because the products never act as inputs to (new) reactions, linear methods like calculating eigenvalues and eigenvectors can be used to analyze the network mathematically [44, 45]. Unfortunately, our context of combinatorial innovation makes such linear methods inapplicable.

3 Results

3.1 Probability of RAFs

First, the probability of finding RAFs at the end of a TAP run (i.e., when at least 1000 goods have been produced) is investigated. The catalysis probability p_c was increased from $p_c = 0.0030$ to $p_c = 0.0060$ with increments of 0.0002 (16 different values). For each value of p_c , 1000 runs of the TAP simulation were performed. Figure 3 shows the probability of finding a (max)RAF in the final step of a TAP run (i.e., the fraction of the 1000 runs that contained a RAF) for each value of p_c .




Fig. 3 The probability of RAFs (Pr[RAF]) for different values of the catalysis probability p_c .

The figure shows a typical S-shaped curve, going from never finding a RAF for $p_c < 0.0030$ to almost always finding one for $p_c > 0.0060$. For a value of $p_c = 0.0044$ the probability of finding a RAF is close to 0.5 (i.e., about half of the runs result in a RAF). This S-shaped curve is similar to what is observed in the standard binary polymer model used in previous studies on

RAFTs in a chemical context [16]. However, a higher level of catalysis is required in the TAP model to find RAFTs. For example, in the binary polymer model it suffices to have each polymer catalyze (on average) between one and two chemical reactions to get a probability of RAFTs around 0.5 (at least for model instances with a maximum polymer length of up to 50) [16]. But in the TAP model it requires (on average) between five and six product transformations catalyzed per good. The average size of the product transformation networks at the end of the runs is about $\overline{M}_t = 1250$, so the average number of product transformations catalyzed per good is $p_c \cdot \overline{M}_t = 0.0044 \cdot 1250 = 5.5$ to get $\Pr[\text{RAFT}] = 0.5$.

An important difference between these two models is that in the binary polymer model the ratio between the number of reactions and the number of molecules grows linearly with n (with n being the length of the largest polymers), whereas in the TAP model the ratio between the number of product transformations and the number of goods is constant (in particular, the number of product transformations is exactly the same as the number of goods). The reason that there are roughly n time more reactions than molecules in the binary polymer model is that each polymer of a given length n can be produced from $n - 1$ ligation reactions between smaller polymers [16, 31].

3.2 RAFT sizes

Next, the sizes of the RAFTs are investigated. Figure 4 (left) shows a scatter plot of the maxRAFT sizes (vertical axis) against the size of the full product transformation network (i.e., the number of goods M_t at the end of the run) for the 1000 runs with $p_c = 0.0044$ (which resulted in $\Pr[\text{RAFT}] \approx 0.5$). The line of dots at the bottom of the graph represents the close to 500 runs that did not result in a RAFT (i.e., the RAFT size is zero), or where there may have been a small number of product transformations where one of the M_0 initial goods is transformed into a new good, also catalyzed by one of the M_0 initial goods. In the current implementation the initial items (i.e., those in the food set) were allowed to be catalysts as well (see below for more on this).

In those roughly 500 cases where there was a “real” RAFT, the maxRAFT size is anywhere from about half the total number of items M_t up to the size of the full reaction network (i.e., $\text{size}(\text{RAFT}) = M_t$). The far majority of RAFTs are actually close to the full network in size. Figure 4 (right) shows a histogram of the relative maxRAFT size, i.e., the RAFT size divided by M_t . This histogram again shows the close to 500 instances with no RAFT (or just a few product transformations with all inputs and at least one catalyst in the food set), and around 330 instances where the RAFT size is close to or equal to the full network. In other words, if a RAFT exists, it immediately tends to be quite large, often consisting of the entire network. This is unlike the binary polymer model, where RAFTs (when they occur with a probability of about 0.5) are usually only around 10% of the full network (in terms of number of reactions). But note again that in the binary polymer model there is redundancy in the reactions in the sense that multiple reactions will have the same product, which




Fig. 4 Left: The size of a (max)RAF against the size of the full network it is part of. Right: A histogram of the relative RAF sizes (as a fraction of the full network).

is not the case in product transformation networks resulting from the TAP model.

For the range of p_c values used so far, a RAF usually only shows up in the final step, when $M_t > 1000$ is reached. And as Figure 4 shows, they tend to be very large immediately. To see if it is possible that RAFs might first emerge as smaller sets and then grow over time, one single run of the TAP model was done with a larger catalysis probability $p_c = 0.0100$. In this case, a RAF shows up a few time steps before the end of the run. Figure 5 (left) shows the last 9 steps ($t = 188$ to $t = 196$) of this particular run, with the number of goods M_t represented by the solid line and the maxRAF size represented by the dashed line.

As the figure shows, at time step $t = 190$ a RAF shows up and immediately consists of a large fraction of the full network ($M_{190} = 405$, $\text{size}(\text{RAF}) = 337$). In just a few more time steps the RAF then grows to encompass virtually the entire network, which it does at the final time step (both being of size 1144).

Finally, the importance of individual product transformations is investigated in the RAF that was found in the final step of the single run just presented (with $p_c = 0.0100$). Previously, a version of the TAP model that also includes a death rate μ was investigated [13]. In other words, goods can “die” with a given probability, meaning they cannot be used anymore as a parent to create new goods. Here, a death rate of $\mu = 0.0$ was used in all experiments. However, the impact of such “death” events can still be inferred by removing each individual product transformation from the maxRAF, and then applying the RAF algorithm again. This will obviously result in a smaller RAF. In some cases, the reduction may be only by one product transformation (the one just removed). But in some cases this affects other product transformations in the RAF as well (as they may now have lost their catalyst, or an input), and the reduction in RAF size could be much larger. The initially removed product transformation is then put back into the original RAF, and the next product




Fig. 5 Left: The number of goods M_t (solid line) during the last few steps of a single run, and the size of the RAF (dashed line) it contains. Right: A histogram of the number of product transformations by which the maxRAF from the final step ($t = 196$) decreases upon removal of each product transformation individually.

transformation is removed and its impact on the size of the remaining RAF is measured, and so on for all individual product transformations.

Figure 5 (right) shows a histogram of the RAF size reductions (measured in number of product transformations) caused by these individual removals. As the histogram shows, the far majority of product transformations have little impact on the RAF size. However, there are also a few that have a very large impact, sometimes even reducing the RAF size to almost nothing. But overall, the RAF seems to be quite robust and resilient against random removal of product transformations. This property was also observed in the binary polymer model.

3.3 Irreducible RAFs

As mentioned above, the M_0 initial goods (i.e., the food set) were allowed to be catalysts. This often gives rise to a small number of product transformations where an initial good is transformed into a new one, also catalyzed by one of the initial goods. Such a product transformation by itself forms a (sub)RAF of size one. When searching for the smallest RAF subsets (i.e., irreducible RAFs, or irrRAFs) in a larger (max)RAF, it is consequently always these irrRAFs of size one that are found.

To get a better idea of the presence (and sizes) of irrRAFs, one simulation run was performed where the initial goods were *not* allowed to be catalysts. A catalysis probability $p_c = 0.0044$ was used, which still gives a probability of finding a RAF of close to 0.5 over a large set of runs. For one particular run, which resulted in a (max)RAF of 1101 product transformations, a random sample (of size 100) of irrRAFs was obtained.

The distribution of irrRAF sizes from this sample is presented in Figure 6. These sizes range from about 360 to 470, with an average of 412 (i.e., slightly

more than one-third of the maxRAF). Although all 100 irrRAFs in the sample are different, they have an average overlap of close to 80%, meaning that for any (arbitrary) pair of irrRAFs from the sample, about 80% of their product transformations are the same.




Fig. 6 A histogram of irrRAF sizes.

4 Discussion

As we indicated in our introduction, the emergence of RAFs from the TAP process reconciles evolution with mutualism. Emergent RAFs admit of change and dynamism as Ulanowicz explains in his analysis of the “centripetality” of autocatalytic processes [46]. Change and dynamism enter because one element in a RAF may replace another without change in the other elements. Figure 7 (drawn from [46]) illustrates. Technological change might produce a faster computer or, say, a more reliable internal combustion engine. This new element D may enter the into competition with B, the slower computer or less reliable engine, and replace it. “The elements in an autocatalytic set move in emergent paths of self-reinforcing interdependence. But this mutualistic self-reproducing nature of the process can be disrupted by the arrival of a competitive element that draws energy or resources toward itself and away from the incumbents it is competing with” [14, p. 113].

RAFs are emergent from a larger dynamic ecosystem. They are themselves, therefore, subject to change in the way Ulanowicz described. Modeling combinatorial evolution with the TAP equation gives a mutualism that is compatible with and emerges from a dynamic evolutionary system. But the mutualist economic vision is fully reflected in the relationships of the autocatalytic web of




Fig. 7 Centripetality induces competition.

production. Even in a dynamic evolving economy, coal and iron make coal and iron.

Our analysis suggests that in economy and ecology alike, “competition is *subsidiary* to centripetality”, which, in turn, “rests on *mutuality*” [46, p. 1888]. As Cazzolla Gatti *et al.* say, if markets are autocatalytic, then “the competitive element, though important and real, rests on mutuality. In other words, the ‘competitive market process’ is first and foremost a cooperative social process” [14].

The history of technology suggests that mutualism preceded competition. Our biological ancestors were making tools and cooperating long before they engaged in market exchange. It is difficult to say when in the long history of the species something like “market competition” first emerged. We can probably say that something like “market exchange” exists to at least some degree when there is trade beyond the small group. When that first began to happen is disputed. But the earliest it may have happened was between 500,000 years ago and 300,000 years ago [47], which is long after mutualistic cooperation among toolmaking hominins began. The production relations of that long earlier period were mutualistic even though life in the small group was far from idyllic and peaceful. These facts bolster Ulanowicz’s view that competition “rests on mutuality”.

When composite tools emerged about 500,000 years ago [48], the economic RAFs of our ancestors began to look like simple versions of the economic RAFs of today. Composite tools are “conjunctions of at least three technounits, involving the assembly of a handle or shaft, a stone insert, and binding materials” [49, p. 1751]. Figure 8 illustrates with a simplified representation of how archaic humans made the adhesive for spear making. These adhesives bound a stone point to a wooden shaft, thus producing a hunting weapon more deadly than a purely wooden spear. A plant resin would be combined with red ochre or other “loading agents” and heated [50].

We have included in our RAF an equally simplified representation of how fire might have been made with the aid of a friction stick that, in turn, would have been shaped with a sharp flake. Already at this still-early stage of the human Technosphere, a complete representation of it would be far too complex to provide here. And even the sliver we have given in Figure 8 is greatly simplified. When we move from the world of archaic humans such as *Homo heidelbergensis* to the modern world created by the Industrial Revolution, the complexity of the Technosphere rises greatly, as emphasized in [15].

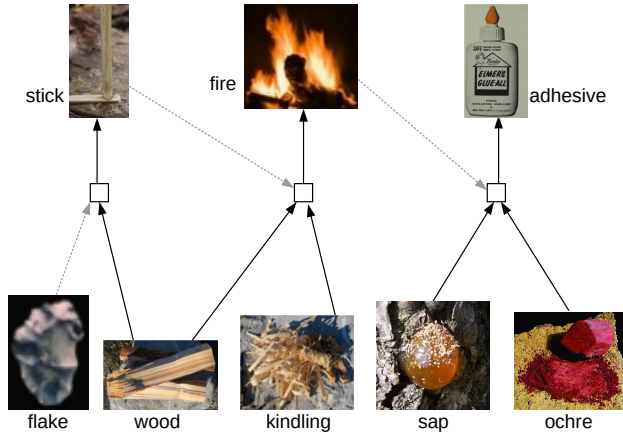


Fig. 8 An early economic RAF for making adhesive.

The TAP process and the concept of an economic RAF are helpful in understanding the evolutionary history of the human Technosphere. They are helpful in other applications as well as suggested by the results and discussions in [14, 15]. It is our belief that the TAP equation and the analytics of emergent economic RAFs can be useful tools for the development of evolutionary economics.

For example, TAP processes are non-ergodic, which suggests that they may be helpful in developing the idea of lock-in proposed by Arthur and others [24]. Furthermore, in general it cannot be predicted *which* new products will come into existence, because any product can, in principle, be used for indefinitely many functions [39]. However, we cannot deduce one (possible) function of a product from another (current) function of that same product. TAP implicitly recognizes this, as it is not explicitly stated what exactly each product is, so the same product can be used in indefinitely many ways in combination with others.

Finally, to cite one more example, we believe that our analytics may be helpful in studying the “twin hockey sticks” noted by Koppl *et al.* [15]. The hockey stick of economic growth is a good thing, which has lifted billions out of grinding poverty. But it has an evil twin, namely, the crisis of the Anthropocene. Is it possible to keep the good hockey stick without succumbing to its evil twin? If so how? These questions may be the most urgent of the question our analytics may help with. But there is, we believe, an indefinite host of further questions that our TAP analytics might illuminate.

5 Conclusions

In the true spirit of combinatorial innovation, we have combined two different visions of economic theory, one evolutionary and the other mutualistic. Using

the TAP and RAF models, we have shown that the two approaches to economic theory are compatible, and that autocatalytic (RAF) sets are highly likely to emerge in product transformation networks that result from a process of combinatorial innovation (TAP). We believe our approach to economic theory advances our understanding of economic development and adds valuable analytic methods to the economist’s toolbox.

We have shown that, depending on the probability p_c , any good can catalyze any product transformation, a product transformation network needs to be of a certain minimum size to have a RAF set emerge. But when it does, such a RAF set likely consists of almost the entire product transformation network. However, smaller autocatalytic subsets (irrRAFs) exist within the larger RAF, on average about one-third the size of the full network.

Many of the results obtained here are similar to those from earlier studies on RAF sets in a simple model of polymer chemistry, which is both reassuring and encouraging. However, there are some subtle differences, mostly resulting from the difference in the ratio between the number of reactions (product transformations) and the number of molecules (goods). In the binary polymer model, this ratio increases linearly with the length of the largest polymers, whereas in the TAP model this ratio is constant. In other words, in the polymer model there are many reactions that can create a given molecule, whereas in the TAP model there is always only one product transformation that creates a given good.

Of course there are many possible variations of the basic model as presented and analyzed here. For example, a power law distribution in catalysis could be used, where most goods do not catalyze anything and a few goods catalyze many product transformations. Or an economy could be “partitioned” into specialized sub-markets. For example, one market could specialize on everything to do with forestry, and another market on everything related to IT, but goods can act as catalysts both within and between markets (e.g., computers are helpful in any market, and wooden desks are needed for informaticians to work on). However, results on the binary polymer model have shown that such variations of the basic model often lead to very similar results, at least qualitatively. Moreover, sometimes the quantitative differences can be predicted mathematically from the basic model [17]. We expect the same to be true for the TAP model.

The initial results presented here are promising, and warrant further investigation. The resiliency of productive networks is a good illustrative candidate. Covid-era supply-chain problems have underlined the importance of the resiliency of production networks. We have shown that autocatalytic production networks are generally resilient to random removals of product transformations, but that the removal of certain linchpin transformations can induce a significant collapse in RAF size. (This result is reminiscent of [51].) It seems reasonable to expect that RAF models could be used to help identify when and how shocks such as Covid-era disruptions will produce relatively large collapses in production and when and how they will produce relatively

small disruptions. Presumably, the degrees of redundancy and “degeneracy” [52] in production networks will be an important factor. RAF models might be used to produce such measures in terms of the amount of overlap between different irrRAFs. We hope that the initial results we have presented here will stimulate further research in such directions.

Acknowledgments. The authors thank Mike Steel for helpful comments to improve an initial version of this article, and two anonymous reviewers for suggesting additional technical and contextual improvements.

Compliance with ethical standards

No funding was received to perform this work, and no conflicts of interest are declared.

Data availability statement

All data was generated with custom-made implementations of the TAP and RAF algorithms. This data has not been made publicly available, though, as it would be highly appreciated to have an independent implementation of the TAP + catalysis algorithm to verify the results presented here. An independent (but well-tested) implementation of the RAF algorithm is available in the program *CatlyNet*: <https://github.com/husonlab/catlynet>.

References

- [1] Schumpeter, J.A.: History of Economic Analysis. Oxford University Press, New York (1954)
- [2] Law, J.: Money and Trade Considered, With A Proposal for Supplying the Nation with Money. Augustus M. Kelley, New York (1705). (1966 edition)
- [3] Smith, A.: An Inquiry Into the Nature and Causes of the Wealth of Nations. Liberty Fund, Indianapolis (1776). (1981 edition)
- [4] Menger, C.: Principles of Economics. New York University Press, New York (1871). (1981 edition translated by J. Dingwell and B. F. Hoselitz)
- [5] Veblen, T.: Why is economics not an evolutionary science? The Quarterly Journal of Economics **12**(4), 373–397 (1898)
- [6] Young, A.A.: Increasing returns and economic progress. Economic Journal **38**, 527–542 (1928)
- [7] Nelson, R.R., Winter, S.G.: An Evolutionary Theory of Economic Change. Harvard University Press, Cambridge, MA (1982)

- [8] Quesnay, F.: *Tableau Oeconomique*. Macmillan and Co, London (1758). (1894 edition)
- [9] Walras, L.: *Elements of Pure Economics*. Richard D. Irwin, Inc., Homewood, IL (1874). (1954 edition translated from the Edition Définitive of 1926 by W. Jaffé)
- [10] Schumpeter, J.A.: *The Theory of Economic Development*. Oxford University Press, Oxford (1911). (1934 edition)
- [11] Leontief, W.W.: Quantitative input and output relations in the economic systems of the united states. *The Review of Economics and Statistics* **18**(3), 105–125 (1936)
- [12] Sraffa, P.: *Production of Commodities by Means of Commodities: Prelude to a Critique of Economic Theory*. Cambridge University Press, London (1960)
- [13] Steel, M., Hordijk, W., Kauffman, S.A.: Dynamics of a birth-death process based on combinatorial innovation. *Journal of Theoretical Biology* **491**, 110187 (2020)
- [14] Cazzolla Gatti, R., Koppl, R., Fath, B.D., Kauffman, S., Hordijk, W., Ulanowicz, R.E.: On the emergence of ecological and economic niches. *Journal of Bioeconomics* **22**, 99–127 (2020)
- [15] Koppl, R., Cazzolla Gatti, R., Deveraux, A., Fath, B.D., Herriot, J., Hordijk, W., Kauffman, S., Ulanowicz, R.E., Valverde, S.: *Explaining Technology (Elements in Evolutionary Economics)*. Cambridge University Press, Cambridge (2023)
- [16] Hordijk, W., Steel, M.: Detecting autocatalytic, self-sustaining sets in chemical reaction systems. *Journal of Theoretical Biology* **227**(4), 451–461 (2004)
- [17] Hordijk, W., Steel, M.: Chasing the tail: The emergence of autocatalytic networks. *BioSystems* **152**, 1–10 (2017)
- [18] Witt, U.: How evolutionary is Schumpeter’s theory of economic development? *Industry and Innovation* **9**(1–2), 7–22 (2002)
- [19] Witt, U.: The future of evolutionary economics: Why the modalities of explanation matter. *Journal of Institutional Economics* **10**(4), 645–664 (2014)
- [20] Carpintero, O.: When heterodoxy becomes orthodoxy: Ecological economics in ‘The new Palgrave dictionary of economics’. *The American*

- Journal of Economics and Sociology **72**(5), 1287–1314 (2013)
- [21] Wagner, R.E.: A macro economy as an ecology of plans. *Journal of Economic Behavior & Organization* **82**, 433–444 (2012)
- [22] Darwin, C.: *On the Various Contrivances by Which British and Foreign Orchids Are Fertilized*. Murray, London (1862)
- [23] Bascompte, J.: Mutualism and biodiversity. *Current Biology* **29**, 425–473 (2019)
- [24] Arthur, W.B.: *The Nature of Technology*. Free Press, New York, NY (2009)
- [25] Enquist, M., Ghirlanda, S., Eriksson, K.: Modelling the evolution and diversity of cumulative culture. *Philosophical Transactions of the Royal Society B* **366**, 412–423 (2011)
- [26] Roser, M.: Economic growth. *Our World in Data* (2013). (<https://ourworldindata.org/economic-growth>)
- [27] Karakas, A.D.: Destiny of the 'hockey stick' economic growth. *The Economics Review* (2019). (<https://theeconreview.com/2019/09/12/destiny-of-the-hockey-stick-economic-growth/>)
- [28] Kauffman, S.A.: Economics and the collectively autocatalytic structure of the real economy. *NPR 13.7 Cosmos & Culture* (2011). (<https://www.npr.org/sections/13.7/2011/11/21/142594308/economics-and-the-collectively-autocatalytic-structure-of-the-real-economy>)
- [29] Hordijk, W.: Autocatalytic sets: From the origin of life to the economy. *BioScience* **63**(11), 877–881 (2013)
- [30] Kauffman, S.A.: Cellular homeostasis, epigenesis and replication in randomly aggregated macromolecular systems. *Journal of Cybernetics* **1**(1), 71–96 (1971)
- [31] Kauffman, S.A.: Autocatalytic sets of proteins. *Journal of Theoretical Biology* **119**, 1–24 (1986)
- [32] Ashkenasy, G., Jegasia, R., Yadav, M., Ghadiri, M.R.: Design of a directed molecular network. *PNAS* **101**(30), 10872–10877 (2004)
- [33] Vaidya, N., Manapat, M.L., Chen, I.A., Xulvi-Brunet, R., Hayden, E.J., Lehman, N.: Spontaneous network formation among cooperative RNA replicators. *Nature* **491**, 72–77 (2012)
- [34] Arsène, S., Ameta, S., Lehman, N., Griffiths, A.D., Nghe, P.: Coupled

- catabolism and anabolism in autocatalytic RNA sets. *Nucleic Acids Research* **46**(18), 9660–9666 (2018)
- [35] Miras, H.N., Mathis, C., Xuan, W., Long, D.-L., Pow, R., Cronin, L.: Spontaneous formation of autocatalytic sets with self-replicating inorganic metal oxide clusters. *PNAS* **117**(20), 10699–10705 (2020)
- [36] Sousa, F.L., Hordijk, W., Steel, M., Martin, W.F.: Autocatalytic sets in *E. coli* metabolism. *Journal of Systems Chemistry* **6**, 4 (2015)
- [37] Xavier, J.C., Hordijk, W., Kauffman, S.A., Steel, M., Martin, W.F.: Autocatalytic chemical networks at the origin of metabolism. *Proceedings of the Royal Society B* **287**, 20192377 (2020)
- [38] Xavier, J.C., Kauffman, S.A.: Small-molecule autocatalytic networks are universal metabolic fossils. *Philosophical Transactions of the Royal Society A* **380**, 20210244 (2022)
- [39] Kauffman, S.A., Roli, A.: A third transition in science? *Interface Focus* **13**, 20220063 (2023)
- [40] Temkin, O.N., Zeigarnik, A.V., Bonchev, D.: *Chemical Reaction Networks: A Graph-Theoretical Approach*. CRC Press, Boca Raton, FL, USA (1996)
- [41] Hordijk, W., Smith, J.I., Steel, M.: Algorithms for detecting and analysing autocatalytic sets. *Algorithms for Molecular Biology* **10**, 15 (2015)
- [42] Hordijk, W.: A concise and formal definition of RAF sets and the RAF algorithm. *arXiv* (2023). <https://arxiv.org/abs/2303.01809>
- [43] Hordijk, W., Steel, M., Kauffman, S.: Autocatalytic sets arising in a combinatorial model of chemical evolution. *Life* **12**(11), 1703 (2022)
- [44] Jain, S., Krishna, S.: Autocatalytic sets and the growth of complexity in an evolutionary model. *Physical Review Letters* **81**, 5684–5687 (1998)
- [45] Steel, M., Hordijk, W., Xavier, J.C.: Autocatalytic networks in biology: structural theorems and algorithms. *Journal of the Royal Society Interface* **16**, 20180808 (2019)
- [46] Ulanowicz, R.E.: The dual nature of ecosystem dynamics. *Ecological Modelling* **220**, 1886–1892 (2009)
- [47] Brooks, A.S., *et al.*: Long-distance stone transport and pigment use in the earliest middle stone age. *Science* **360**, 90–94 (2018)
- [48] Wilkins, J., Schoville, B.J., Brown, K.S., Chazan, M.: Evidence of early

- hafted hunting technology. *Science* **338**, 942–946 (2012)
- [49] Ambrose, S.H.: Paleolithic technology and human evolution. *Science* **291**, 1748–1753 (2001)
- [50] Zipkin, A., Wagner, M., McGrath, K., Brook, A.S., Lucas, P.W.: An experimental study of hafting adhesives and the implications for compound tool technology. *PLoS One* **9**(11), 112560 (2014)
- [51] Albert, R., Jeong, H., Barabasi, A.-L.: Error and attack tolerance of complex networks. *Nature* **406**, 378–382 (2000)
- [52] Edelman, G.M., Gally, J.: Degeneracy and complexity in biological systems. *PNAS* **98**(24), 13763–13768 (2001)