

The Two Faces of Emergence in Economics

Mark Kuperberg

As this anthology makes clear, there is not one definition of Emergence that is universally agreed to, nor for progress to be made in the field does there need to be. For the purposes of this essay, however, I will use a stripped down definition which if not common to all emergent processes, does at least summarize what is at the core of most examples. These core characteristics include:

- 1) At least two levels of organization,
- 2) A multitude of individual agents at the lower level of organization who operate by following simple rules, and
- 3) An aggregate outcome at the higher level that results from the interaction of these individual agents, but which is not easily derivable from the rules that the individual agents follow. Many times, therefore, this aggregate outcome comes as a surprise to the observer because nothing in the rules at the lower level seem to predetermine the aggregate outcome.

If we take these three characteristics to be a canonical representation of emergent processes, then Economics was certainly the first discipline to be founded upon emergent principles. In 1776, Adam Smith wrote in The Wealth of Nations:

“It is not from the benevolence of the butcher, the brewer, or the baker that we expect our dinner, but from their regard to their own interest. We address ourselves, not to their humanity but to their self-love, and never talk to them of our own necessities but of their advantages (Book I, Chapter 2) ... every individual...neither intends to promote the public interest, nor knows how much he is promoting it...he intends only his own security; and by directing that industry in such a manner as its produce may be of the greatest value, he intends only his own gain, and he is in this, as in many other cases, led by an invisible hand¹ to promote an end which was no part of his intention. Nor is it always the worse for the society that it was no part of it. By pursuing his own interest he frequently promotes that of the society more effectually than when he really intends to promote it.” (Book IV, Chapter 2)

Not only was Economics founded on this principle, but it remains the central dogma of the discipline to this day. What distinguishes an economist from other social scientists (and other people in general) is a faith that self interest at the lower level will, when channeled through competitive markets, result in a beneficial outcome at the aggregate level. Modern economics has discovered many exceptions to this rule, but

¹ As an 18th century man, Adam Smith was referring to Providence, or God, when he used the phrase “invisible hand”. So Smith’s views cannot be considered fully modern in that emergence, as now understood, does not consider the phenomena that emerge at the higher level to be designed by anyone. For a long time now, economics has taken the “invisible hand” to refer to the impersonal forces of supply and demand which is more consistent with the modern meaning; but as will be discussed later, economics assumes the existence of a top-down coordinator which is somewhat a variance with the complete absence of a designer.

they remain the exception and Adam Smith's insight remains the rule². With the possible exception of evolutionary biology, there is no modern academic discipline that has a concept of emergence so at its core.

The Representative Agent

Robinson Crusoe has probably had a bigger influence on Economics than he has had on English Literature. Economists use Robinson Crusoe to derive conditions for economic efficiency. Economic efficiency is defined as an allocation of resources such that it is impossible to make someone better off without making someone else worse off. If you only have one someone, economic efficiency is synonymous with Robinson behaving sensibly and not wasting any of his resources³. For a Robinson Crusoe economy to find itself in an inefficient economic allocation, Robinson would have to be an idiot.

While there are conditions for economic efficiency which cannot be derived from an economy with only one person, it is surprising how many can. Adam Smith clearly understood that competitive markets benefitted society, but he did not know how to formalize the conditions that assured this beneficial outcome. It probably never occurred to him that such a formalization was possible. It took the economics profession more than one hundred years to derive what we now call the conditions for Pareto efficiency⁴.

What ties Adam Smith and Robinson Crusoe together is the First Theorem of Welfare Economics which states that a competitive market will create an allocation of resources that is Pareto efficient. This theorem enables economists to illustrate the benefits of competitive markets by studying what it is rational for Robinson Crusoe to do in isolation. I should emphasize that the Theorem is not a form of misplaced anthropomorphism, though the persistent use of the representative agent by modern economists may be. The Welfare Theorem is a rigorously proved proposition that does not conceive of the economy as one large individual. Still, the surprise is how much can be learned from such a conceptualization, but that surprise is completely different than the surprise contained in Adam Smith's original insight.

For Smith, what was surprising was that individuals motivated by self-interest could nevertheless promote the interest of society. In a Robinson Crusoe economy, there is no society, and it is completely *unsurprising* that Robinson Crusoe promotes his

² Technically, competitive markets only achieve economic efficiency under an additional set of assumptions, but for the purposes of this essay a discussion of these assumptions is unnecessary. I will, therefore, refer throughout the essay to the entire bundle of necessary assumptions by the shorthand "competitive markets".

³ It is important to note that with more than one person, there are an infinite number of efficient allocations of resources, so economic efficiency does not, in general, imply only one allocation.

⁴ Named after the economist Vilfredo Pareto who first systematized the conditions that are satisfied by an efficient allocation of resources in Manual of Political Economy, 1906.

own self interest. The representative agent embodied in Robinson Crusoe enables economists to turn what is a hard problem of market analysis into what is, in essence, an engineering problem as Robinson seeks to maximize his lifetime utility. The surprise that we can many times study a whole economy by looking at one isolated individual is, in a sense, diametrically opposed to Adam Smith's surprise. What made Smith's insight so remarkable was that there was a disconnect between the two levels of analysis: the rule at the level of the individual was self-interest, but what emerged at the societal level was what we now call economic efficiency. In the Robinson Crusoe correspondence, the rule at the level of the individual is optimization and the outcome at the societal level is what we now call a Pareto efficient/optimal allocation. There is nothing surprising about optimality flowing from the behavior of an individual to an entire economy in an economy with only one individual.

When I say that we can study the whole economy by looking at one individual, I mean this in a normative sense - we can study some of the efficiency conditions that can be achieved by competitive markets. The problem comes when economists start making *positive* statements about the real economy on the basis of a representative agent. The conditions under which the behavior of an entire economy can be predicted from the behavior of one individual are, not surprisingly, very severe. They basically amount to assuming that everyone in the economy is identical in terms of tastes and income. This, of course, is never true. So, for example, if it is the case that when the fish are "running", Robinson spends more time fishing because the price of fish in terms of foregone leisure has declined; we cannot conclude that for an entire economy, the demand for fish will go up as the price goes down. The First Theorem of Welfare Economics *does not* guarantee this. What the First Theorem guarantees is that if the economy is competitive, then whatever outcome emerges, when the price of fish falls, will be efficient.


The Dark Side of Emergence

So, the first face of emergence in economics, which comes down to us in a direct line from Adam Smith, is a very positive one: the road to heaven may be paved with bad intentions. Agents acting selfishly can, nevertheless, create an aggregate outcome such that it is impossible to make someone better off without making someone else worse off. This is an amazingly strong statement. The first economist, to my knowledge, to create an emergent model whose outcomes were not socially desirable was Thomas Schelling in Micromotives and Macrobehavior (1978). Schelling analyzed how neighborhoods would emerge given that people had some preference to live near people like themselves.

Figure 1⁵ below illustrates a "society" where people (the red and black dots) are distributed randomly throughout the space.


⁵ NetLogo Model created by Uri Wilensky 1998.

Figure 1



Each individual has 8 neighbors, and we assume that they will move unless $3/8$ th (37.5%) of their neighbors are of the same color as themselves. This is not a strong preference for segregation, and as a result, in the Figure 1, 72.1% of the people are happy - meaning that they have at least 3 neighbors of their same color. Nevertheless, when you move people around until no one is unhappy, Figure 2 emerges which has a substantial degree of segregation.

Figure 2




In Figure 1, approximately 50% of ones neighbors shared the same color, but in Figure 2 the number is over 80%. The surprise is that a relatively mild preference for living with people of the same color results in a substantial degree of segregation. So the rule at the lower level, “move if less than 3/8ths of your neighbors are of the same color” results in an aggregate outcome where more than 80% are of the same color.

Antz

The Schelling Model does not relate directly to economics. While the outcome is bad given a social preference for integration, one cannot say that the outcome is inefficient. In fact, since in the final equilibrium everyone is satisfied with their neighborhood, one could say the outcome is efficient. Of more relevance to economics is the model in Figure 3⁶ which is derived from a paper by Kirman⁷. Ostensibly, it is a model of ants who have a nest in the middle of the graph and forage for food from two equidistant food sources (red and blue) at the edges of the graph.

There are three kinds of ants: ants that have no source affiliation, blue ants who forage at the blue source, and red ants who forage at the red source. Initially all ants have no affiliation, but when they discover one of the sources they become that kind of ant and bring food back to the nest and then go out again to that source. If a blue ant encounters an unaffiliated ant (one that has not yet discovered a source), then that ant is recruited to become a blue ant (similarly for red ants). The final effect that makes the model interesting is that an affiliated ant that is not carrying food can be converted to the other color, with some probability, if it encounters an ant of the other color.

Figure 3



⁶ NetLogo Model created by author.

⁷ Alan Kirman, “Ants, Rationality, and Recruitment”, Quarterly Journal of Economics, February 1993.

Possible economic applications for such a model are choosing to adopt one of two alternative technologies, choosing to do business with one of two alternative firms, etc. The model then neatly illustrates two opposing views of how this competition will evolve:

- 1) Since the food sources are equidistant from the nest, equally plentiful, and the ants initially move randomly, one might think that 50% of the ants will be red and 50% will be blue. In the context of this model, this would be the “competitive” outcome, and it is what would be predicted by what economists call the Hotelling model⁸.
- 2) Since ants can recruit and convert other ants that they meet, one might think that if one source develops a lead in ant affiliation, it will build on that lead and ultimately all the ants will be of that color.

Which of these two outcomes emerges is not only of academic interest. One of the major driving forces behind the stock market bubble of the late 1990's was the belief that if a firm developed a lead in internet customers, it would lock in that customer base and have very high profits in the future even if it was currently suffering severe losses. What the model shows is that, as expected, if there is no recruitment or conversion of ants, the Hotelling result emerges: approximately 50% of the ants are red and 50% are blue. Surprisingly, this result is essentially unchanged if there is recruitment but no conversion. With recruitment and no conversion, the ability of ants affiliated with a given source to recruit other ants does not tip the scales irreversibly to the first source found. The fact that ants are randomly searching for a source at the beginning insures a nearly equal split between sources. In either case without the possibility of conversion, the model is in equilibrium when all ants are affiliated with some source, recruitment simply speeds this process up.

Figure 4 illustrates the case of recruitment with a conversion rate of 75% and plots the proportion of red ants. As can be seen, even after 15,000 periods the model does not settle into an equilibrium, the percentage of red ants fluctuates widely. Why is this? The reason is the complex interplay between positive and negative feedback that is at work in the model. Positive feedback results from the fact that when there are more ants of a particular color, it is more likely that an unaffiliated ant will meet an ant of that color and be recruited plus the fact that there are more “missionary” ants of that color to convert ants of the opposing color. If these were the only mechanisms in operation, eventually all ants would be of one color. Negative feedback results from the fact that when there are more ants of a particular color, there are necessarily more ants who are not carrying food of that color. For ants of the other color, therefore, there are many potential converts. For the ants of the minority color, the graph is a target rich environment. If the conversion rate is set high enough, these two forces are continually at war with one another.

⁸ Named after Harold Hotelling, who in "Stability and Competition", Economic Journal, 1929 first developed an equilibrium model of spacial competition. .

Figure 4

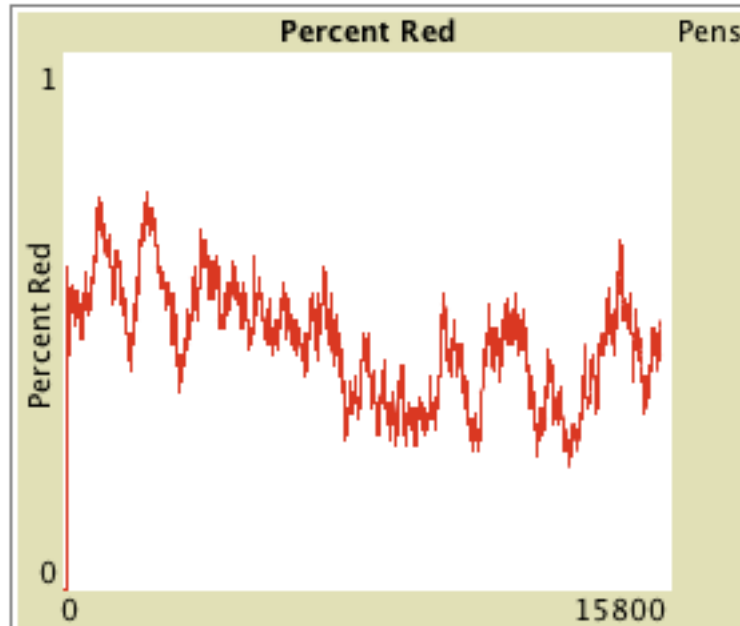


Figure 4 illustrates the danger in telling top-level stories or finding patterns in top-level phenomena when the underlying process is emergent. Looking at the time series in Figure 4, macroeconomists might analyze the tops and bottoms of the percent red ants as peaks and troughs of business cycles and seek macroeconomic explanations for their occurrence. Technical stock market analysts might look at the pattern of percent red ants and claim to be able to predict future movements⁹. But we know from how the model was constructed that telling aggregate stories of movements during particular time periods is nonsense because all of the observed phenomena were caused by interactions at the local level.

A New Kind of Economics

Just as there are multiple definitions of emergence, there are multiple descriptions of how an economics based on emergent principles differs from traditional

⁹ There is a lot of confusion in the emergence literature as to whether emergent phenomena are necessarily random and/or unpredictable. The Ant model is completely deterministic and in this sense completely predictable. Where technical analysts go wrong is in assuming that one can predict the movements in the graph based on past movements in the graph. In order to predict the movements of the percent red ants, one needs to know where every ant is and how it is interacting with every other ant.

The non-linear deterministic processes in Chaos Theory are also predictable in the above sense because they contain no random elements, but the computing power necessary to make these predictions may be so large that the distinction between “unpredictable in principle” and “unpredictable in practice” may be mute.

neoclassical/Walrasian¹⁰ economics. As with my definition of canonical emergence, I will state the minimum set of characteristics that distinguish what has come to be known as “agent based computational economics” from traditional economics.

A fortiori, agent based computational economics is populated by heterogeneous agents. I say *a fortiori* because it is in the very nature of emergence that agents interacting at a local level cannot be identical. So, for example, in the segregation model, agents differ by their initial position in the grid and therefore by who their immediate neighbors are. Even if agents were initially programmed to be identical, their local interaction with one another would be different and they would soon cease to be identical. This necessary lack of a representative agent means that one cannot in emergence adopt the Robinson Crusoe methodology where economic efficiency flows to the whole economy from the maximizing behavior of one individual. Still, heterogeneity in no way negates the First Theorem of Welfare Economics. A strength of the First Theorem and of Walrasian economics is that both are perfectly capable of dealing with any degree of heterogeneity. So, while some economists see heterogeneity as a hallmark of agent based computational economics, its real role is to eliminate the possibility of using the intellectually suspect Robinson Crusoe methodology.

What fundamentally differentiates agent based computational economics from traditional Walrasian economics is that economic activity occurs outside of equilibrium. Agent based models can certainly have an equilibrium. In the segregation model, for example, equilibrium occurs when everyone is content with the color distribution of their neighbors. This equilibrium is not unique, if by unique we mean an identical final pattern of dots; it depends critically on the initial placement of the dots and also on the order in which people get to move. It is generally the case that once out of equilibrium trades or economic activity are allowed, the ultimate equilibrium, if there is one, will not be unique.

While the existence of an equilibrium is important for the analysis in this essay, uniqueness of the equilibrium is not a central concern. The essential question is whether the equilibrium will be efficient. What assures efficiency in traditional Walrasian economics is the Walrasian auctioneer who aggregates all supply and demand information and allows trading only at equilibrium prices. The Walrasian auctioneer is the economics version of a top-down coordinator, and it is a hallmark of emergent processes that there is no such coordinator. Without the auctioneer, one must generate the final equilibrium from the local interaction of the individual agents. Under what circumstances such an equilibrium will be efficient is an open question.

Macroeconomics

Macroeconomics is the study of the economic activity of the economy taken as a whole. To carry out this study, macroeconomists create economywide aggregates of individual real world variables. Some of these aggregates are sums of individual variables, such as gross domestic product, which is the sum total of the economy's

¹⁰ Named after Leon Walras who first formalized the economy as a general equilibrium system in Elements of Pure Economics, 1874.

production of goods and services for a given time period; other aggregates are averages of individual variables, such as the price level or the inflation rate, which average individual prices and their percentage changes. The goal of macroeconomics is to understand the movements of and the relationships between these various aggregates. If we conceive of the economy as an emergent system, then from this description, it should be obvious that macroeconomics is inherently studying its top-level behavior.

Modern macroeconomics began with the publication in 1936 of The General Theory of Employment, Interest, and Money by John Maynard Keynes. Virtually since its inception, there has been a research agenda to provide microfoundations for the relationships between the macroeconomic aggregates. For the most part, this research program has used traditional neoclassical/Walrasian economics to provide the microfoundations. This approach contained within itself an internal contradiction which only became fully obvious in the 1970's with the advent of what has come to be called New Classical economics. We have seen that traditional Walrasian economics shares with Adam Smith an optimistic view of the workings of the economy. The central message of The General Theory, however, was that the performance of the economy would many times be sub-optimal. Because of this internal contradiction, the effort to provide microfoundations for Keynesian macroeconomics has yet to produce a model that is convincing to most economists.

What I wish to argue here is that the reason for this failure to provide adequate microfoundations may be that we are using the wrong microeconomic paradigm. Instead of using traditional neoclassical/Walrasian analysis perhaps we should be using the kind of emergent processes that are not so biased toward producing rosey outcomes. Returning to the ants model, we can see that it exhibits internally generated apparent cyclical behavior. I say apparent because there unquestionably is not a mechanism generating a fixed periodicity to these cycles. The cyclical behavior emerges from the local interaction of the ants. This is in stark contrast with the standard macroeconomic explanation for apparent cyclical behavior which is that the economy is hit by an exogenous disturbance, that the internal mechanisms of the economy may initially augment but ultimately dampen down this disturbance, and that the only reason there appears to be business cycles is that the economy is hit later on by another disturbance.¹¹

A key premise behind the standard view is that macroeconomic events must have macroeconomic causes: changes in the macroeconomic aggregates must be the result of macroeconomic disturbances. This is precisely what an emergent perspective calls into question. What the standard view calls a macroeconomic disturbance can be, as in the ants model, the bubbling up to the macroeconomic surface of small events at the local level. Some events at the local level are nullified at the local level: so, for example, an ant not carrying food converts to the opposing color, but then meets an ant of its original color and converts back. We never see these events at the top-level and are completely unaware of their existence. But sometimes, a local interaction, or the random occurrence of many local interactions of the same type, is propagated by

¹¹ There is a large literature dealing with endogenous business cycles, but this is not the majority view among macroeconomists.

positive feedback into a bigger and bigger event until it emerges at the top-level as a macroeconomic event.

In the give and take between micro and macroeconomists, a standard line by microeconomists is, "there is no such thing as macroeconomics", by which they mean that all that really exists is individual behavior and its aggregation into markets by standard Walrasian methods. The same statement could be made about a macroeconomics based on emergent principles. In fact, the same statement could be made about all emergent phenomena. The case could be made that all of what is observed at the top-level is epiphenomenal and that the only reality is the local interactions. I would argue, however, that one does not need to give up on aggregate relationships and finding higher level laws in emergent processes in general and in macroeconomics in particular, but our conception of what those relationships and laws will look like may need to change. The paradigm should be Boyle's Law where an aggregate equation describes the top-level behavior of a gas with no reference to the interactions of the individual gas molecules. The aggregate relationship must, of course, be consistent with what is happening at the micro level; but in an emergent system, there is no presumption that the aggregate outcome will be a mirror image of the micro rules.

With respect to macroeconomics, for example, the accounting identities continue to hold whether the phenomena is emergent or not. In addition, thinking of the macroeconomy as an emergent system does not preclude macroeconomic disturbances, just as putting a burner under a balloon will both expand the balloon and increase the pressure within according to Boyle's Law. What we need to be leary of is the presumption that all macroeconomic events have macroeconomic causes, not that none do. Finally, what needs to be abandoned is the naive presumption embodied in the Robinson Crusoe methodology that optimality at the individual level implies optimality for the entire economy.