The Emergence of Specialization

Hugues Bersini

IRIDIA-CP 194/6
Université Libre de Bruxelles
50, av. Franklin Roosevelt – 1050 Bruxelles
Belgium - bersini@ulb.ac.be

Abstract

The emergence of specialization remains a true challenge. Suppose a world initially filled with specialized and generalist agents, and let's define these later as able to endorse the various competences characterizing the specialists and to endorse them as well as the specialists. In an evolutionary perspective, it is obvious to see why a specialist will always be less adapted than a generalist, which can indeed alternatively act as many experts. The generalists will meet much more agents and much more situations to which they are adapted to and then cumulate much more payoff (unless the tasks done by generalists are systematically of pitiful quality). This is indeed a paradox to face in order to make sense of a world nonetheless full of specialists. This paper will discuss various ways, beyond the obvious possibility of unfavoring multi-specialization by paying a high cost, to allow specialists to survive the presence of generalists.

Introduction: Specialist vs generalist

We will here define a specialist as an agent which, in contrast to a generalist, can only accomplish a subset of specific tasks. This restriction can be explained by either genetic or phenotypic specificity or any educational or cultural bias. As a direct consequence, a specialist suffers from lack of autonomy; it is never self-sufficient and requires complementary specialists to interact with in order to survive. A first version of this necessary specialists grouping with other complementary specialists is generally designed as "mutualism", for which each specialist contributes by its own competence to the interest of another. It is certainly one possibility among others to understand the "invisible hand" metaphor of Adam Smith, which famously claims that 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. The butcher enjoys good bread as much as the baker enjoys good meat. Naturally so, the baker will make the best bread to the advantage of the butcher who reciprocally will prepare the best meat to the benefit of the baker.

Another version or justification for this grouping, different (as figure 1 shows) from "mutualism", is "division of labor", for which specialists join together to the benefit of an external user, outside of the group. Only this external user, exploiting the group of specialists as a whole, can "make sense" of their joining together and rewards back any member of the group in the case it benefits from their collaboration.

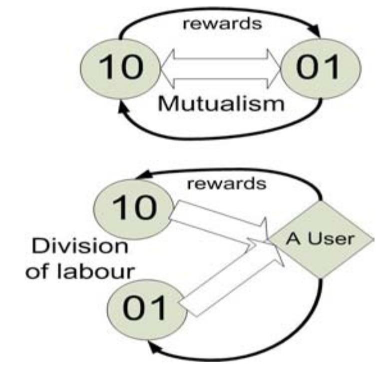

Figure 1: Mutualism vs "division of labor". "10" and "01" are just ways of identifying two specializations.

For specialization to succeed, cooperation is obviously needed, relating the problem of the emergence of specialization with the one of "cooperation", made so popular since the invention of the "prisoner dilemma" (Nowak, 2006). Here we initially see these two problems as orthogonal, assuming that cooperation needs to be resolved before specialization can take place. Without cooperation no specialization turns out to be possible and, indifferently, either generalist or specialist can choose to defect in an interaction. There will be more to say on that precise topic in the conclusions.

A generalist, on the other hand, can stand alone and execute any of the tasks, as a function of which specialist it interacts with. Both in the natural and in the human world, specialization seems to be the common case. Example of "old fashioned" mutualism in the human world could be "husband at work" and "housekeeping wife". Any form of win-win commercial exchange is a successful case of mutualism. Biology is full of example of mutualism, such as associations between plant roots and fungi, with the plant providing carbohydrates to the fungus in return for nitrogenous compounds and water. Other examples are the Lynn Margulis endosymbiotic theory (Margulis, 1998) or the co-virus replication, only possible due to the existence of complementary virions.

Regarding the division of labor, an entrepreneurial team of construction specialists, composed of plumbers, carpenters and electricians is an alternative to a single handyman. Another nice example of contrast between generalists and a "division of labor" team of specialists is illustrated in figure 2. In nature, examples of "division of labor" are numerous such as in insect societies or in the very well studied biological reality of slime mold moving (Alexopoulos et al, 2004). As will be discussed later, natural selection can favor grouping of complementary specialists to improve the fitness of the whole group.

Figure 2: A generalist vs a divided band of specialists

So either in a case of division of labor or in a mutualistic one, how could specialization be favored in the presence of generalists able to adopt any form of specialization? Obvious solutions are that either it costs too much to be generalist or the way a generalist accomplishes the task is always worse than when done by the agent specialized in that specific task. This is an obvious possibility and does not need any further justification. However, the main motivation of this paper is to propose and discuss further evolutionary roads for specialization to emerge in a world possibly inhabited by generalists.

We will first discard other versions of the problem that this paper is definitely not interested in and then describe the ways the issue will be addressed. The technical road taken to study this problem will be the use of evolutionary games, increasingly popular these days for studying the emergence of cooperation in a "genuinely" and "rationally" competitive world. A computer simulation of a spatial version of evolutionary games will be explained and tested on the problem. We will illustrate the different ways for specialization to emerge, insisting more on the original ones, requiring division of labor and the presence of an environment able of a large scale observation of the group of specialists. We will then conclude by relating this solution with a conceptualization of "emergence" that has been advocated in previous papers.

What the problem is definitely not

An interesting related problem is one of combinatorial optimization, as illustrated in figure 3. It can be enunciated as follows. A group of specialists is available and one needs to find the best sub-group obtained by sampling some of these specialists and grouping them together. A score is associated with each possible sub-group, so that in order to find the best solution (i.e. the best group of specialists) some form of grouping optimization algorithm is required.

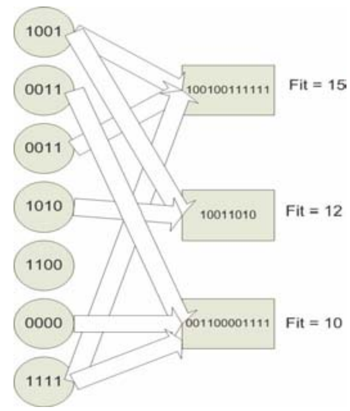

Figure 3: A combinatorial optimization problem to group the specialists in an optimal way.

While it is an interesting problem per se, the problem is definitely not what we are interested in since, in such a case, specialization is already there and does not need any form of justification. The only challenge is to make use of it in the best way but not to question its "raison d'être". However, many engineers remain confronted to this version of the problem and it is quite an interesting one for GA users.

What are the evolutionary driving forces of specialization? An evolutionary game study

The question addressed in this paper is: What can be the evolutionary driving forces for specialization to come out in a world in which generalists could be as likely and perform the task as well as specialists? In other words, what could drive a generalist to become specialized? This question could result to be unexciting because trivial to answer, or tedious, because dealing with a non realistic world (like for instance a world in which agent would have no other choice that popping up endowed with an innate specialization). We believe this not to be the case since both the natural and the human world offer many examples of generalist (stem cell, omnivore, medical generalists and handyman) and many opportunities for these generalists to turn into specialist, provided there is any advantage in doing so. So for what specific reasons could a generalist choose to specialize?

Our favorite way to face the problem is to rely on evolutionary games, pioneered by John Maynard-Smith (Maynard-Smith, 1989) and very actively studied and exploited by researchers such as Martin Nowak (Nowak, 2006). While the Darwinian inspiration of evolutionary games is obvious, these same mathematical and algorithmic tools can help to make sense of the social world in which men tend to imitate the most successful of their neighbors or colleagues.

In what follows, a specialized agent will be defined as a binary string of length n, allowing then 2^n types of specialization. The use of binary strings allows to easily define complementary specialization, by computing the Hamming distance between two specializations. For instance, the task “00” will perfectly complement the task “11”. Additionally, the generalists will be characterized by the presence of the John Holland’s don’t care symbols into the binary strings (Holland, 1975). For instance, a complete generalist would be “##”, while a partial one could be “#1”, making thus possible some degree of generalization.

A simple canonical case

Let’s focus on the simplest case with n=1 and 3 types of agents: “0”, “1” and “#”. The evolutionary two players game payoff matrix for such a simple situation could be given as:

	#	0	1
#	b-c	b-c	b-c
0	b-c	0	b
1	b-c	b	0

with $b > c$.

Table 1: The two players payoff matrix

“*b*” is the reward gained by an agent when benefiting from the task done by its complementary partner in the game. Thus when two similar agents meet, they gain nothing: “0”. On the other hand, in order to make the appearance of specialists a challenging issue, a cost “c” penalizes a task done by a generalist, so that any agent (either a specialist or a generalist) meeting a generalist for a coupled interaction will receive a “*b-c*” reward. The reasons for the “c” could be twofold: either generalists can do the same task as specialists but with much less competence, or they can perform this task as well but it simply costs to be a generalist capable of so many diverse specializations. When asking to the layman why they do believe that the world is filled with more specialists than generalists, this would generally be the kind of answers they give: being a generalist means a lack of competence or it is simply much harder to achieve.

Many other matrices are possible, with easier to justify and in general more sophisticated reward values attributed to each of the nine possible entries. But for mathematical reasons to follow, this elementary matrix is enough to convey the sole effect of the cost “c” on the success or disappearance of generalists.

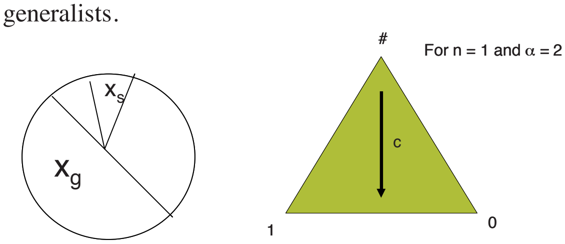

Figure 4: The basis of the mathematical analysis and its result illustrated for $\alpha = 2$ and n = 1.

Suppose a population of agents decomposed in the following way (as illustrated in figure 4), a part x_g of generalists and α subgroups of complementary specialists x_s. The time evolution of generalists and specialists can be mathematically described as follows, by supposing that any specialist can only interact with its unique complementary specialists (and we suppose a same concentration x_s for any group of specialists), and generalists can interact with all generalists and all specialists.

$$\dot{x}_s = x_s((b-c)x_g + bx_s - \phi)$$
$$\dot{x}_g = x_g((b-c)(x_g + \alpha x_s) - \phi)$$
$$x_g + \alpha x_s = 1$$

After some easy manipulations:

$$\dot{x}_s = (1 - \alpha x_s)x_s^2(b(1-\alpha) + c\alpha)$$

So that the two possible solutions are:

$x_s = 0$ or $x_s = 1/\alpha$ depending on the value of c as compared to $b(\alpha\text{-}1)/\alpha$.

We clearly understand here the two main reasons for the dominance of generalists over specialists: the cost c and the frequency α. The bigger the cost the harder for the generalist to survive and the steady-state solution in such a case turns out to be a subdivision of the population in α subgroups of complementary specialists (illustrated in fig. 4 for two groups of complementary specialists). The smaller the specialist frequency of encounters (i.e. the higher the “α”) the easier it is for generalists to replace specialists, since specialists have fewer agents to interact with and fewer opportunities of reward.

Such a simple mathematical analysis is the reason behind the simplicity of the payoff matrix as compared with much more sophisticated versions of it (Wahl, 2002, D’Orazzio and White, 2006) (as shown in figure 5 extracted from (Wahl, 2002)).

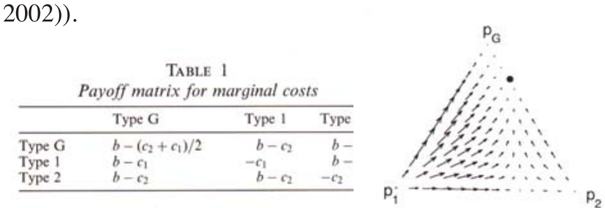
TABLE 1
Payoff matrix for marginal costs

	Type G	Type 1	Type
Type G	$b - (c_2 + c_1)/2$	$b - c_2$	$b -$
Type 1	$b - c_1$	$-c_1$	$b -$
Type 2	$b - c_2$	$b - c_2$	$-c_2$

Figure 5: A finer mathematical analysis of the evolutionary game between two groups of specialists and generalists.

Nevertheless, as a matter of fact, the main qualitative results remain unchanged through all the various versions of the problem. The only way to make specialists survive the dominance of generalists is either to increase the cost of generalization or to augment the possibility of specialist opportunities to cumulate payoff, here the number of complementary specialists they can interact with.

The spatial cellular automata simulation

Again largely inspired by Nowak's spatial version of the prisoner dilemma (Nowak, 2006), we propose a simple cellular automata type of simulation in which every cell contains one agent of length n, either specialist (without "#") or partially or completely generalist. The spatial environment is a 2-D toroidal cellular automata in which every agent/cell can interact with its 8 neighbors (generally called a "Moore neighborhood").

The simulation goes as follows. At every time step, each agent will interact with its 8 neighbors and cumulate the payoff perceived according to the two players payoff matrix. Then asynchronously (by means of a random selection iteratively performed over all agents a number of times equal to the number of cells), each agent is both selected and "reproduced" by being replaced by the best agent in its neighborhood. Thus, at each time step, all agents compute their payoff and reproduce according to their fitness value, the best agents locally invading the neighborhood. Two results are shown in figure 6, for n= 1 (so three types of agents: 1, 0 and #). The first result is shown for a low value of c (generalists win) and the second for a high value of c (the two groups of specialists win and equally divide the population).

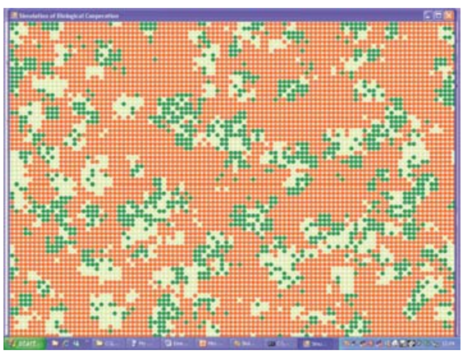

In this snapshot of the simulation, the generalists are invading the whole population (the large clusters of homogeneous cells – the generalists – are percolating through the small two-color clusters – the specialists).

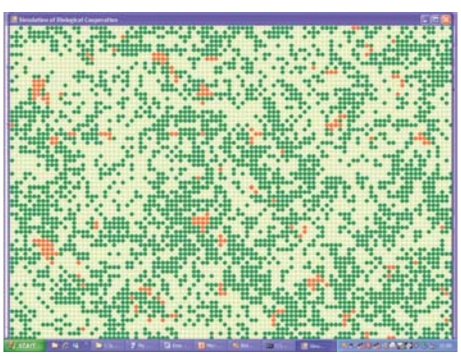

In this snapshot of the simulation, the specialists are invading the whole population. The two colors represent the two complementary groups of specialists: "1" and "0".

Figure 6: Results of the spatial simulation: above for a low cost of "generalization" and below for a high cost.

These simulations just reproduce the results anticipated by the mathematical analysis. The outcome of the simulation is binary and depends on the value of the cost, below or above the threshold. Above specialists win, below generalists win.

The question addressed in the following is: "Is there any less trivial driving factor than just the cost of generalization?" According to the mathematical analysis, another important contributing factor, beyond the cost, is the frequency of encounters. This frequency depends on the number of opportunities that a specialist has for cumulating reward, but that generalists miss for one reason or another. So is there any other original way for specialists to increase this frequency of encounters that would however escape the generalists? We admit to enter from now on in the realm of speculation.

Division of labor

In order to make a real profit from their specialization, specialists should better cluster together in the hope of creating by such grouping more opportunities of payoff. This is exactly the message provided by the idea of "division of labor" which definitely needs to be distinguished from the simple mutualism.

The following possibility will be added in the spatial simulation namely that complementary specialists can cluster together when they are neighbors. At every time step, before any reward gain and before any reproduction, a specialist is allowed to connect to at maximum two of these complementary neighbors, like illustrated in figure 7. Two is the minimal number for making possible groups of complementary specialists to appear.

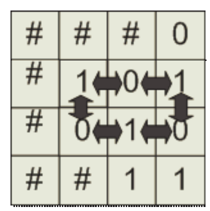

Figure 7: A possible cluster of complementary specialists

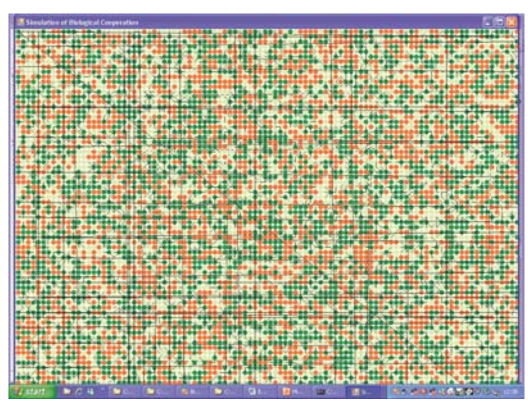

Figure 8: A snapshot of the simulation where the black lines reflect the clustering among the specialists.

The following snapshot of the simulation (figure 8) shows the presence of black lines which testify for the connections among specialists only. The clusters just survive one time step namely the time for the agents to interact (in drawing some payoff opportunity of their presence in the cluster) and then reproduce. Once reproduced, an agent looses its membership to any cluster and its connectivity is simply reset to null. The payoff gained by a specialist agent will be equally shared among all the members of the cluster it belongs to. The simulation at every time step thus takes place in three successive steps: 1) cluster, 2) get payoff by interacting and 3) reproduce.

An obvious concern could come from the impossibility for generalists to cluster. Why should it be so? Are we not resolving the problem in a too much ad hoc way? A simple reason is that, by clustering with, for instance, a specialist, a generalist would loose its capacity to adopt any other form of specialization since it would be forced to take one specific profile. It would sacrifice its "chameleon" side. A human generalist, once convinced of the advantages in being so, would decline any opportunity to freeze himself in one of his multiple possible profiles. In a more primitive world, we will take for granted that a non-specialized organism cannot connect with another non-specialist or with a specialist, since the pattern recognition ability which allows two specialists to recognize each other and to connect could be absent from the generalists. But remember that these are all speculations, wished for discovering more opportunities for a specialized world to replace one populated mainly by generalists.

If nothing is modified in the simulation and the specialist and generalists just interact as before, the results won't be affected by this added clustering possibility. The gain of the clustering and thus the increase opportunity of frequency is perfectly compensated by the sharing of the benefits. For instance, in figure 7, any one of the six agents composing the cluster will take six times more benefit, nevertheless always divided by six. Of course, increasing the benefit of the specialist with respect to the generalist (for instance, by not dividing the benefit by six and just distributing the whole "b" to any one) would favor specialization. However, such a move is far to be satisfactory since this benefit has to be equally shared through the cluster in a way or another. Additionally, this would reduce the whole problem to that same question of cost that we want to avoid. Is there any further way for specialists to increase their frequency of encounters which could make sense while giving them a natural advantage over generalists?

Simulation with varying length of agents

Suppose now a new version of the simulation in which agents can be of any length (with n the maximum length): "1", "00", "##1"… Suppose further that any agent of length x can only interact with and take profit from agent of equal or inferior length: x, x-1, x-2, … (Just as if it would swallow it). So at any time step, an agent looks in its neighborhood to discover equal or smaller agents. It thus computes its total profit only with this restricted part of its neighborhood. As a consequence, the payoff matrix is no more symmetrical, since an interaction between agent x and agent x-1 (if both specialists and complementary) will reward "b" to agent x but 0 to agent x-1. The complementarity between two agents is simply assessed by comparing the first common bits. An agent can be of length n in two different ways: either it is naturally of length n (from birth) or it takes part of a cluster of dimension n. Finally, only agents of length 1 can cluster (remember that, among them, only specialist can cluster) and reproduce.

This makes obviously a lot of assumptions, leaving a strong impression of arbitrariness. Generalists cannot cluster and only the smallest agents can reproduce. These last limitations are simply there to make easier to grasp the results of the simulation, which could similarly happen for less restrictive conditions (for instance bigger agents, due to their size, could simply be slower to reproduce). The simulations are performed with a low value of the cost c, which has always favored generalists in all previous cases. The results of the simulations are shown both for a maximum length of 2 and 3 in the two following plots of figure 9.

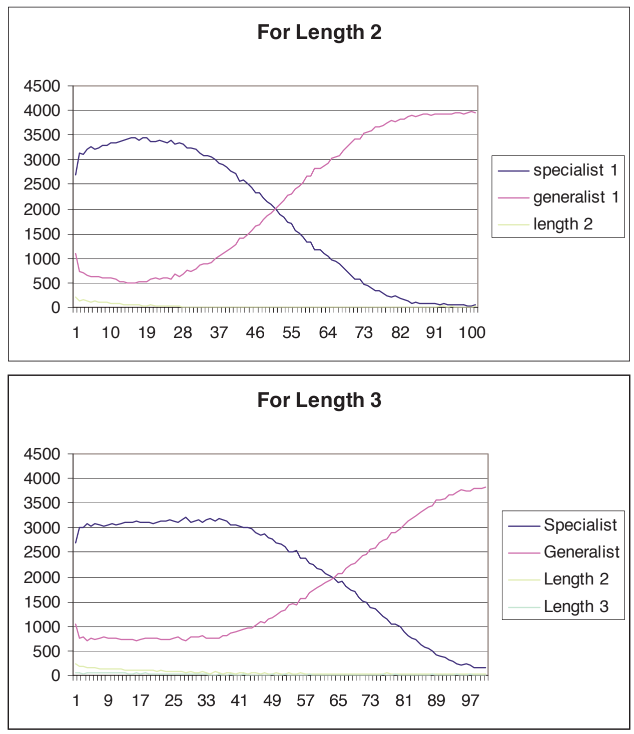

Figure 9: Results of the simulations for maximum length of 2 and then 3 and a low cost of generalization. It is important to look also at the decreasing of the curves at the bottom of the two plots, showing the decreasing concentration of agents of length 2 and 3. The life time of specialization increases with the length of the simulation.

These plots show the cumulated concentration of specialists of length 1, the concentration of generalists of length 1 and the concentration of agents of length 2 and 3. One can observe that during the first time steps, the number of specialists grow in size and easily win the game, the duration of this winning period depending on the agent maximum length: the greater the longer. The reason is easy to understand. Initially, the specialists, by clustering together, can access to sources of reward that are simply inaccessible to generalists (remember that the generalist cannot cluster so that they can only interact with agents of length 1). For instance, a cluster of two or more specialists can interact with agents of length 2, an impossible source of payoff for generalists.

Then, since the big agents cannot reproduce, they are overcome by the specialists until the instant where only agents of length 1 remain in the simulation. From this instant on, the simulation proceeds as usual (the small value of the cost favors generalization) with the generalists invading the world by benefiting from much more sources of rewards than the specialists. So besides the obvious possibility to excel in the task, another possibility for specialists to impose themselves is to create communities of complementary members which, as a whole, can achieve new functionalities much beyond the possibility of any isolate generalist. Here this is what happens, since the bigger agents in the simulation can be exploited by clusters of specialists. The generalists simply can't see them.

Emergence and conclusions

We were interested in this paper in investigating other evolutionary driving forces, beyond the simple cost, in order to favor the survival of specialists to the expense of generalists. Our mathematical analysis shows that, whatever formulation given to that problem, these forces will always have to do with either the cost (that we discard here because too obvious and commonly accepted) or the frequency of encounters. This last effect is more original since it might give rise to ways for specialist to create opportunities of rewards that escape generalists. In substance, this is what we have attempted to do through our last simulations.

Suppose for instance that in order to construct your house, you can choose between a handyman capable of optimally achieving all parts of the construction alone: bricklaying, carpentry, electricity and plumbing, and an entrepreneurial team composed of as many specialists as required. Whatever the quality of the handyman work, there is clearly one thing he will never be able to achieve alone: work on all parts in parallel, something obviously possible for a team. Therefore a group of specialists is able to work in parallel and achieve the same construction in a much shorter time, whatever the quality of each single part of the construction. This is a simple human example of a clustering benefit inaccessible to a single generalist. When adopting the "division of labor" perspective and by joining together, specialists can discover a whole new set of opportunities that a generalist (which has other good reasons to stay alone) can't even see. Specialists, by grouping together, access to opportunities that generalists simply miss.

As regards the problem of "cooperation vs defection", in a classical world where mutual defections turn out to be the only Nash equilibrium, this "joining together" might provide an extra road for cooperation to emerge. Only by connecting together, can the specialists achieve some reward. So defectors, just like generalists, by refusing to be part of the groups, would disappear to the advantage of cooperators. We further suppose here that all members of the group have to do their own, so that groups allowing for the integration of defectors will never succeed in accomplishing the task. In a mutualistic situation, defection wins over cooperation (as defectors don't give but just receive), while, in a division of labor situation, cooperation has obvious advantages (defectors don't receive any payoff).

In a very stimulating book of Matt Riddley (Riddley, 1998) entitled "The origins of virtue", in the chapter entitled "Division of labor", we can find these excerpts:

> "*In a phrase, therefore, the advantage of society to me is the division of labor. It is specialization that makes human society greater than the sum of its parts… It is this synergy between specialists that makes human societies tick, and it is this that distinguishes us from all other social creature … Adam Smith was the first to recognize that the division of labor is what makes human society more than the sum of its parts… As division of labor between specialists evolves, integration into higher unit systems also advances, and, as social homeostasis evolves, the individual human loses some portion of his self-regulation and becomes more dependent for his existence upon the division of labor and the integration of the social system*".

The repetition of the famous expression "The whole is more than the sum of its parts" relates the concept of "division of labor" with the concept of "emergence". In previous publications (Bersini, 2004; Bersini and Philemotte, 2006), I have discussed this later concept and the necessary three ingredients which together allow a collective phenomenon to be described as "emergent". First the phenomenon, as usual, requires a group of agents entering in a non-linear relationship and entailing the existence of two semantic descriptions depending on the scale of observation: micro or macro. Second, the macro phenomenon (the one that raises philosophical debate) has to be observed and "objectivised" by a mechanical observer, which has the natural capacity for temporal and/or spatial integration. This mechanical observer positively substitutes for the human one, which is blameworthy of endowing "emergence" with an unacceptable dose of subjectivism. Finally, for this natural observer to detect and to select the collective phenomenon, it needs to do so in rewards of the adaptive value this phenomenon is responsible for. Basically, physics can simply explain the proximate causes of the phenomenon while natural selection provides the "adaptive", "engineering", "ultimate" causes of the functionality. The presence of natural selection brings me to defend in these previous papers the idea that emergent phenomena can only belong to biology.

In the European Swarm-bots project, which is being coordinated in our laboratory (Gross et al. 2006), largely inspired by the capacity of some insect species (such as ants) to assemble in order to accomplish tasks that none of them, alone, is able to accomplish, small robots connect together to do as well. For instance, two robots join together in order to pass over a gap that would make any of them fall down if trying alone. I claim that, in the case of real insects, "passing over that gap" is an emergent behavior, since it requires a cluster of agents, an external observer integrating the agent's behavior in space and time (the gap here plays the role of this required natural observer) and natural selection to favor the agents able to pass over the gap.

In the last simulations presented in this paper, again the same three ingredients are present: the agents cluster together; once they do cluster they can exist for and be observed by the "bigger agents", and once they do exploit these resources, they are favorably selected. Consequently, one extra road for both cooperation and specialization to appear, beyond the quality of the task or the cost of generalization, might really be any new opportunity made possible by groups of specialized agents which can do more than their parts. Specialization really emerges in a genuine sense once specialized agents decide to enter, not in a mutualistic relationship, but in a genuine strategy of "division of labor", so as to make it relevant for an external observer: human or biological, whose presence justifies the fitness increase of the specialists. This jump in fitness provided by the grouping of single specialized agents into new types of organisms is one possible way to construe the concept of the major transitions in evolution (Maynard-Smith and Szathmary, 2005).

References

Alexopolous, C.J., Charles W. Mims, M. Blackwell et al (2004)., *Introductory Mycology, 4th ed*. John Wiley and Sons, Hoboken NJ.

Bersini, H. (2004). Whatever emerges should be intrinsically useful. In *Artificial life 9*, pages 226--231. The MIT Press, 2004.

Bersini, H. and Philemotte, C. (2006). Emergent phenomena only belong to biology. In Proceedings of the 9h European Conference on Artificial Life.

D'orazzio, A.E. and White, T.A. (2006). Incomplete Division of labor. Error-prone multitaskers coexist with specialists . J. Theoret. Biol. 2007

Groß R., Bonani M., Mondada F., Dorigo M. 2006. Autonomous Self-Assembly in Swarm-Bots. *in IEEE Transactions on Robotics*, volume 22(6).

Holland, John H (1975), *Adaptation in Natural and Artificial Systems*, University of Michigan Press, Ann Arbor

Margulis, Lynn, (1998). *Symbiotic Planet : A New Look at Evolution*, Basic Books

Maynard Smith, J. (1989) *Evolutionary Genetics*. Oxford: Oxford University Press.

Maynard Smith, J. and E. Szathmary (1995). *The Major Transitions in Evolution*, Freeman press, Oxford.

Nowak, MA (2006). *Evolutionary Dynamics*. Cambridge, MA: Harvard University Press.

Riddley, M. (1998). *The origins of virtue*. Penguin (Non-Classics);

Wahl, L.M., 2002a. Evolving the division of labor: generalists, specialists and task allocation. *J. Theoret. Biol.* 219, 371-388