WASPS: A Weight-Allocated Social Pressure System for the Emergence of Agent Specialization

Denton Cockburn[1] and Ziad Kobti[1]

[1] University of Windsor, School of Computer Science, Windsor, Ontario, Canada
cockbu3@uwindsor.ca; kobti@uwindsor.ca

Abstract

Division of labour, or specialization, is common in many types of insect colonies. It emerges in some of these societies as a result of age polyethism, whereby the division of labour is tied to the age of the individuals. One known method that explains this is social inhibition. Individuals release pheromones when they interact with other agents. The strength of their pheromones is tied to their age. These pheromones inhibit the desire of other agents to perform the same task. Using social inhibition, individual agents can be allocated among the available tasks to be performed related to the colony. We apply a variation of this approach to the problem domain where agents can divide their time among multiple tasks. While age is not a factor, agents differ in their skill at performing each task. We create a weight-allocated social inhibition approach whereby more skilled agents inhibit the desire of less skilled agents to perform a task. We are able to see that this approach drives agents toward tasks where they have comparative advantages. This leads to an increase in the division of labour within the population. While inspired by social insects, this approach is easily applicable to agents in other domains.

Introduction

Specialization is where individuals produce goods and services beyond local or personal need, depending upon other individuals to supply other needed goods. There are many varying definitions of specialization, with most taken from the archaelogical, biological and economic fields. One definition from archaeology is that specialization is "the production of substancial quantities of goods and services well beyond local or personal need, and whose production is generally organized, standardized and carried out by persons freed in part from subsistence pursuits" (Arnold and Munns, 1994). By choosing to specialize, specialists must obtain some or all of their subsistence goods through exchange with others (Evans, 1978). There are varying levels of specialization, ranging from being able to sustain oneself, while simultaneously producing goods for the consumption of others, to complete dependency upon exchange with others for subsistence goods. Dependence upon others for subsistence was viewed by Childe as the essence of economic specialization (Childe, 1951).

Specialization allows individuals to maximize their productivity by exploiting their environment (Murciano, 1997), and occurs because entities belong to a community of mutual interest, cooperating to serve that mutual self-interest (Spencer et al., 1998). Specialization may be assigned, as in caste systems, or chosen by an individual driven by varying means, including genetic, social and economic. Another term for specialization is division of labour, which is defined by Hollbloder as "...when individuals can be turned into specialized working machines, an intricate division of labour can be achieved and a complicated social organization becomes attainable even with relatively simple repertory of individual behaviour" (Hölldobler and Wilson, 1990).

There are both internal and external factors that influence an individual's choice of specialization (Beshers and Fewell, 2001). Internal factors include genetic, neural, hormonal and experience elements. External factors include economic factors such as demand (stimulus) and social influences (Julian, 1999; O'Donnell, 1996; Robinson et al., 1989). It seems that no single behavioural model may fully explain division of labour in complex systems (Traniello and Rosengaus, 1997). Different models and approaches have different assumptions, which makes it particularly difficult to compare the effects of factors across different approaches.

The study of specialization is important to several fields. For instance, archaeologists study specialization to understand the changes in societies as a result of the emergence of specialization. It also gives insight into why individuals would choose to produce certain goods over others. From the biological perspective, specialization helps to explain the behaviour of biological creatures such as ants, wasps and bees (Larsen, 2001; Page et al., 1998; O'Donnell, 1996), which have been empirically shown to specialize based on tasks. Economically, specialization is studied to understand its effect upon a society's economy. It further serves to study how a market may grow or contract based on the specializations present, as these specializations lead to increases in the productivity of market systems (Murciano, 1997). Allyn Abbot Young points out that a productive individual increases the supply of certain commodities, while simultaneously increasing the demand for others (Young, 1928). In spite of its role in economics and biology, little is known of the origins and causes of specialization and exchange (Beaudreau, 2003).

In this paper we focus on the social approaches to artificial agent specialization. Here we define an agent as an autonomous social party that can perform several tasks with varying levels of skill. Being social, these parties can also be influenced by their peers across their social networks. It is our hypothesis that competition will drive agents to allocate more of their resources to produce goods with which they possess a comparative advantage in relation to their competitors. As the primary differentiator of efficiency in our model is skill, it can be assumed that more skilled agents will have a comparative advantage over their less skilled competition. In this individual based model, these self-interested agents will be influenced towards performing tasks that will maximize their own productivity. We believe this approach will lead to significant increases in the overall level of specialization within an agent population. In the next section we introduce the social inhibition model from which this work is primarily inspired. We then describe our generic model that uses weight-based allocations. Finally, an experimental setup is presented and discussed with concluding remarks.

Social Inhibition

There are several social models for the emergence of agent specialization. One such method is social inhibition, which implies that as agents choose their specialization, they notify other agents that they have done so, reducing their desire to also choose this specialization. To put that idea in economic terms would be that choosing a specialization reduces the demand (stimulus) for that specialization. Social inhibition aims to explain concepts such as temporal polyethism, which is division of labour based on age, as a result of the interaction between behavioural development and the inhibitory effects of other workers (Huang and Robinson, 1992; Naug and Gadagkar, 1999; Beshers and Fewell, 2001). Temporal polyethism can also be explained experientially, as older agents would have more experience, and thus more knowledge upon which to base their actions (Ravary et al., 2007). This model is more concerned with the physiology of workers and their interactions. Initially, the model took the form of an activator-inhibitor approach, whereby all agents would eventually mature to perform specific tasks, but inhibitors from current performers of these tasks would slow their activation.

Naug and Gadagkar presented a social inhibition model that aimed to explain the age polyethism in wasp species (Naug and Gadagkar, 1999). Their model was in turn based on the verbal model of Huang and Robinson (Huang and Robinson, 1992). In Naug and Gadagkar's model, each agent has two pods: one that increases its own preference for a task, and another that inhibits the preferences of agents it interacts with for the same task. Their model claimed that individual specialization is emergent from the increase in activator due to age, as well as the amount of inhibitors exchanged when agents interact. The model assumes that all agents possess the same preference and skill level for task performance, which makes it difficult to adapt to situations such as those we aim to address.

The effect of competition on task specialization was examined in (Merkle and Middendorf, 2004). Competition was shown to lead to the occurence of specialists as an emergent phenomenon dependent on the size of colonies. Their model was based on a genetic preference model though, whereas our model is based on social interactions. They also studied differing demands for tasks, something which we do not explore here.

Another social interaction model was explored in (Gordon and Trainor, 1992). Agents had an active and inactive state for the four tasks in the model. The agents communicate with each other, giving them some idea of how many other agents are performing the same task. These interactions between agents is designed such that the system will trend toward a stable set-point where there is a balance of active and inactive agents for each task. Like the above mentioned models, they also assume that agents do not possess an innate preference or skill for tasks.

A non-social model that is also relevant is (Lavezzi, 2003). Lavezzi's model shows that the amount of specialization and level of per capita output depends on competition, agent connectivity, agent thresholds, and initial conditions such as number of agents and their connectivity. An agent's potential to choose a specialization is limited by the amount of other agents performing the same task, as well as the stimulus level for that task. Agents of course have to know about the level of competition, or be directly aware of the changing stimulus levels. In either of these two situations, agents are required to have excess knowledge of their economic environment. While non-social, we have found that a lot of the effects claimed by Lavezzi are also evident in our model.

The existing social models have several other shortcomings, several of which we look to address. In these models, agents are only able to perform one task per unit of time. In our model, we aim to deal with situations where agents can divide their time among several tasks. Take for example something like human agents, such as those found in (Kohler et al., 2007), who have several tasks to perform in each year such as farming, hunting, getting water and getting wood.

In the social inhibition model, which is aimed at age based specialization, the social influence of other agents do not directly determine the specializations that others will choose. Tasks must first be ranked in a way related to age, then agents are ranked by fit for those tasks. After that, agents are then assigned based on the number of workers needed for that task. We think that while this may be appropriate in insect colonies, it makes the model difficult to adjust to agent populations where tasks may not have priorities. In our model, we assume no priority among tasks.

Approach

Our approach is not aimed at system optimization, whereby the system itself tries to be the most productive possible. Instead, agents should be able to emerge the specializations that they are most suited for in their given environment. We assume the existence of a set of tasks T. Each element *t* in T is a task that can be performed by an agent. Each agent has a level of skill associated with each task. The skill level may be static, or it may be determined by the agent's previous success at performing the task. This allows for skill levels that may correspond with fitness functions in evolutionary algorithms. This skill level is quantifiable, comparable and monotonic, such that $sk_a(t) > sk_b(t)$ means that agent *a* is more skilled than agent *b* at performing task *t*. All agents assume they can perform the task perfectly. The level of skill is then reflected in the amount of inhibition that agent then releases when they interact with others. Agents are thus able to determine their true relative skill level through interactions with other agents. The strength of inhibition, which we refer to as the influence rate, depends on each agent.

In our test simulations, we assume that all agents have the same level of influence. This is not required, and it is quite possible for different levels to make sense in a domain. For instance, we can create the effect of age polyethism if we were to have the influence rate grow with age. In that case, to create task prioritization, we can have the level of influence vary by task as well. In addition to skill, agents have to divide their time among tasks. They therefore need to track their allocations, which they do internally. Note that while we refer to time, that is simply one idea of a resource. This model does not require the resource to be time, but it can by money, food, or any other divisible resource. The simulation is composed of a set of interacting agents within a social network that can all perform the same tasks at varying skill levels.

Problem Description

Given agent *Ag*, the set of tasks available to Ag T_{Ag} and a resource R_{Ag}, how does an agent allocate its R_{Ag} among each task *t* in T_{Ag}? So, $\Sigma x_i = S(R_{Ag})$, where *i* is each task in T_{Ag}, $S(R_{Ag})$ refers to the amount of the resource R_{Ag} available, and $x_i$refers to a fraction of $S(R_{Ag})$. The problem also involves the following conditions: The problem is continuous over a period of iterations, $S(R_{Ag})$ changes between iterations and x_i is allowed to change over iterations.

Weight-based model for resource allocation among tasks

For each agent *Ag*, we propose a set ALLOC, where $e_i \in$ALLOC $\Rightarrow$there is a task *i* in T_{Ag} and e_i represents the weight allocated to task *i*. Task weights in ALLOC are relative, therefore for a given task *i* and a resource to be allocated R_{Ag}, the amount of R_{Ag} to be allocated to task *i* is: $\frac{e_i}{S(ALLOC)}$ x $S(R_{Ag})$, where S(ALLOC) is the sum of all elements in ALLOC. We make no assumptions about the initialization of the weights in ALLOC; they can be randomly assigned, or initialized by some other method. A task having a weight of 0 will result in the task being allocated none of R_{Ag}. For simplicity, we will assume R refers to time for the rest of this paper. We also normalize the weights in ALLOC such that S(ALLOC) is always equal to 1.

Model outline

Agents influence other agents when they interact. In some networks, such as kin network, it can be assumed that they interact with all their neighbours in each time step. The amount of influence is dependent on skill level. The higher the skill level, the higher the level of influence. When an agent interacts with another, it positively reinforces its own behaviour, while also inhibiting the other agent. The amount of self-reinforcement is the same amount that it inhibits the other agents. After all agents have interacted, the agent subtracts the level of inhibition it has received from the level of activation it has provided itself. The agent also self-activates itself, such that an agent that does not interact with any other agents will still change its behaviour. These effects result in the change of the allocation levels for the agent.

Agent Properties

Agent Attributes

Each agent has the following attributes:

- An allocation set ALLOC = { ti <- [0,1]}, for all tasks $i \in$ T, where ti is the fraction of time the agent will spend on task *i*.
- A skill set SKILL = { si <- [0,1]}, for all tasks $i \in$ T, where *si* is the skill of the agent at performing the task *i*. If an agent cannot perform a task *p*, then the value of *sp* would be 0. The skill level for a task may be dynamic and updated regularly. The skill value as a function must be monotonic though, such that if agent *Ag1* has *si* 0.5 and agent *Ag2* has *si* 0.7, then we can say that *Ag2* is better than *Ag1* at performing task *i*.
- A set PODS = { pi }, for all tasks $i \in$ T, where *pi* is a 3-tuple (A, SA, I). In this 3-tuple for task *i*, *A* represents the activator store for the agent, *SA* is the level of self-activation, and *I* is the inhibitor store for the agent. The agent will increase the weight of the associated task when *A+SA* > 0, and decrease it when *A+SA* < 0.

The idea behind self-activation is the inclination of an agent to perform more of the task at which they are best. This value should be large enough that it will allow an isolated agent to specialize over a long period of time, but it should also be small enough that it doesn't overwhelm the social pressure created by stronger competitors.

Agent Inhibition

The level of inhibition I in an agent's pod for a task i is determined by several factors:

- The skill level of the agent at performing task i.
- The size of the agent's social neighbourhood.
- The influence rate, $IR = (0, 1]$, which is a parameter that determines the strength of an agent's influence. This parameter can be universal, or variable for each agent. It is also possible that the influence rate can be different for each task. We can re-create the effect of polyethism if we were to make IR dependent upon the age of an agent.

Agent interaction

When agents $Ag1$ and $Ag2$ interact, for each task $t \in$ T, we obtain the values in $Ag1$'s pod pt for task t, and $Ag2$'s pod pt for task t. The value in $Ag1$'s A will be decreased by $Ag2$'s I and vice versa. Each agent will also increase its own A by its I. This method allows agents to influence each other only when they interact.

Since agents both exchange inhibition, and inhibition level is tied to skill and influence level, the more skillful and influencial agent would have a greater effect on a neighbour than a less skillful and influential competitor. While the influence of the "better" agent would be stronger, the weaker agent would still inhibit the stronger one. It is also possible for agents to be considered to interact on every iteration, in which case agents would inhibit all others in their neighbourhood. It should be noted that the level of self-activation plays no role when agents interact.

Agent Attribute Updates

During each time period, agents will have performed their tasks based on their allocation weights (ALLOC). If the skill set is dynamic, then it would be updated based on the results of task performance. The influence rate of each agent would also need to be updated. If agents have different influence rates for each task, then the updates would need to be applied for each task.

Agents will then update their allocations based on each task pod. Given a normalized allocation ti for a task i, and a pod (a, s, x) for the same task i, then ti will be updated as:

$ti = ti + a + s$. That means that the amount of self-activator s will be added to the activator a, and the sum of that added to the current weight. If an agent was overall more skilled at a task than the other agents it interacted with, then its actiavator level a should increase. If it is less skilled overall, then the level should decrease, resulting in a negative value for a. After all task weights are updated for an agent, the values are again normalized, resulting in the sum of all weights in the agent's ALLOC being 1.

Experiments and Results

To measure the level of specialization within a population, we use a measure developed in (Gorelick et al., 2004). The measure quantifies the degree to which agents in a population are specialized. We have each agent record their task allocation amounts. These amounts are then stored in an nxm matrix, with n being the number of agents and m the number of tasks. We then normalize this matrix such that the sum of all cells is 1. The mutual information and Shannon entropy index (Shannon, 1948) are then calculated for the distribution of individuals across tasks. Finally, dividing the mutual information score by the shannon entropy score will provide a value between 0 and 1. A score of 0 indicates a population with no specialization, while a score of 1 indicates a fully specialized population (Gorelick et al., 2004).

We test our method across several parameter types. These are: the type of network, the number of tasks, the number of agents, and the influence rate. We test with two network types, small-world networks and random networks. Small-world networks (Milgram, 1967) are networks whereby most nodes are connected by a small degree of separation, with the existence of a power-law structure among many nodes. Two famous examples of a small-world network are the '6-degrees of separation' phenomenon found within the US population (Milgram, 1967) and a similar phenomenon among many sites on the World-Wide Web (Bu and Towsley, 2002). With random networks, each node will just be randomly connected with another node. We use the same amount of total edges in both network types, dependent upon the number of agents.

We tested for 2, 4, 10 and 20 tasks. Most studies involve 2 to 5 possible tasks (Waibel et al., 2006), while some insect colonies have anywhere from 20 to 40 specializations (Beshers and Fewell, 2001). Although we could have tested for more possible tasks, we observed that 20 would be sufficient to demonstrate the process. As for the number of agents, we tested smaller groups of 10, 50 and 100 agents, as well as larger groups involving 500 and 1000 agents. Each agent acts after the previous step for all others, meaning that all agents operate in the same time step. Tasks are all assumed to take the same amount of time to perform. We tested with a variety of influence rates, these being 0.05, 0.1, 0.25 and 0.5. The influence rate was the same for all agents during each run. We used a constant self-activation rate of 0.05. All agents also have the same capacity for task performance, that is to say the same amount of time available to be allocated. We ran each combination of parameters 10 times.

Each agent would be created with random task allocations. Thus for each available task, the agent would assign a percentage of their time to be spent on that task. As the metric developed in (Gorelick et al., 2004) is dependent upon these task allocations, different populations of agents would necessarily have different initial levels of specialization. As such, it is not possible to compare the initial and ending specialization levels across runs within the same network type, even with the same parameter settings. The initial populations would be the same for different network types when all other parameters are the same. Considering these conditions, we measure the change in the level of specialization over a run. In the tables given, rows represent the number of tasks and columns represent the number of agents. Tables 1 through 4 illustrate a representative sample of our overall results. They report the average division of labour (DOL) and standard deviation with influence rates (IR) of 0.05 and 0.5 for both small-world and random networks. The DOL values are average multiples (DOL at beginning of run / DOL at end of run) of the initial level of population specialization over the 10 runs for each parameter combination. Thus a value of 3.3 indicates that there was a 230% increase in the level of specialization. For brevity, the results of other influence rates are not shown.

The level of specialization increased in all 1600 runs that we simulated. In our small-world networks, the average result was a multiple 3.2 over the initial values, with a standard deviation of 0.75. With our random networks, the average result was a multiple of 3.9, with a standard deviation of 0.97. We believe that the higher increase in our random networks is due to the higher average number of connections between agents. In small-world networks, several agents have a lot of neighbours while most have only a few. As agents are influenced by interacting with others, having more interactions result in each agent moving toward its optimal state faster. This suggests that increasing the level of connectivity between agents will result in more pronounced increases in specialization.

Our results may be depressed by the emergence of equilibrium states within our populations. This is the case when adding more iterations will not result in any increase in the population's level of specialization. This emergence of equilibrium states is not surprising though as it is predicted in (Young, 1928). As the initial level of specialization is randomly between 0 and 1, it is the case that a population with a high initial level of specialization would not have much room for improvement. We would not expect to see a state of equilibrium if we had used a dynamic society, as new births, deaths, and other state changes would keep the situation in a state of flux (Lavezzi, 2003).

We noticed that in many cases agents would not become fully specialized. This may be in spite of the fact that they may be significantly better at a particular task than all competitors. This is because they would still have some pressure to perform other secondary tasks where they may still have some advantage. This became more pronounced as the number of tasks increase. In such cases, agents may possess comparative advantages in multiple tasks, and thus the motivation to increase their allocation in both. As the allocation system is weighted, the increases in both weights offset each other.

While we did notice that in most cases increasing the level of influence would also result in a higher level of specialization, this does not occur in all cases. In our simulations, the level of specialization would decline in many cases when going from an influence rate of 0.25 to one of 0.5. Because of the different initial populations and specialization levels, we are unable to study the effect of changing agent and task amounts.

Conclusion and future work

In this paper we presented a new social inhibition model for the emergence of specialization in agent societies. We showed that this model is able to significantly increase the level of specialization in a random population. While several current models deal with domains where agents can only perform one task at a time, our model deals with having agents that have to allocate their time among several tasks. We have shown that when agents are differentiated by skill level, competition and social inhibition can be used to increase division of labour. We found that our agents will increase their allocation of time among tasks for which they possess a comparative advantage over their neighbours. This follows a well established law of economics. Surprisingly, we also found that using our weight based approach, agents will not necessarily specialize on the task they are most efficient at. This is because the change in allocations for multiple tasks may offset each other. The result seems supported by real world experience, where we have yet to see a modern nation completely specialize on one product. Our model is created in a way that makes it applicable to many domains.

We intentionally kept several parameters abstract because we would like to keep the approach general. Many of the parameters used can be changed to accomodate different domains. We also didn't state how it is that agents interact for the same reason. Interaction could be either broadcast, exchanged through the environment, or exchanged through message passing. The meaning of the social network and its connections is also left open intentionally, such that it could represent a wide range of topics, such as a topographical neighbourhood, or even a collaboration network.

We currently do not account for different levels of resource availability. We would like to investigate what changes if any the model would need to work under those conditions. In addition, we assume that demand is always equal to the amount of a resource produced. It would be a good idea to investigate different levels of demand either globally or locally for each task. We would also like to see how the model performs under dynamic environmental conditions. We would like to apply the model in concrete domains such as human society simulations, or even social insect simulations. We believe that this model can encompass

	10	50	100	500	1000
2	3.3 ± 1.24	2.48 ± 0.41	2.43 ± 0.29	2.28 ± 0.16	2.28 ± 0.10
4	3.46 ± 0.95	2.86 ± 0.28	2.54 ± 0.24	2.53 ± 0.08	2.48 ± 0.12
10	3.07 ± 0.65	2.73 ± 0.31	2.64 ± 0.11	2.69 ± 0.07	2.71 ± 0.06
20	3.36 ± 0.42	3.08 ± 0.29	2.9 ± 0.26	2.92 ± 0.12	2.88 ± 0.08

Table 1: Average DOL multiple and standard deviation with IR = 0.05 in small-world networks.

	10	50	100	500	1000
2	3.82 ± 2.11	2.76 ± 0.36	2.69 ± 0.35	2.48 ± 0.16	2.53 ± 0.11
4	4.39 ± 1.46	3.4 ± 0.43	3.15 ± 0.18	3.09 ± 0.06	3.03 ± 0.15
10	3.74 ± 0.80	3.4 ± 0.47	3.35 ± 0.16	3.43 ± 0.09	3.45 ± 0.06
20	4.54 ± 1.06	3.53 ± 0.35	3.73 ± 0.32	3.72 ± 0.14	3.72 ± 0.11

Table 2: Average DOL multiple and standard deviation with IR = 0.5 in small-world networks.

several of the currently existing social interaction models, including the social inhibition model which inspires it. We didn't think it appropriate to compare our approach to the social inhibition approach here though because they have different assumptions.

References

Arnold, J. and Munns, A. (1994). Independent or attached specialization: The organization of shell bead production in california. *Journal of Field Archaeology*, 21(4):473–489.

Beaudreau, B. (2003). On the origins of large-scale specialization and exchange: A game-theoretical approach. In *The Canadian Economic Theory Conference*, Montreal, Canada.

Beshers, S. and Fewell, J. (2001). Models of division of labor in social insects. *Annu. Rev. Entomol.*, 46:413–440.

Bu, T. and Towsley, D. (2002). On distinguishing between internet power law topology generators. In *INFOCOM*.

Childe, V. G. (1951). *Social Evolution.* Watts, London.

Evans, R. (1978). *Early Craft Specialization: An Example for the Balkan Chalcolithic. In Social Archaeology: Beyond Subsistence and Dating.* Academic Press, New York.

Gordon, D.M., G. B. and Trainor, L. (1992). A parallel distributed model of the behavior of ant colonies. *Journal of Theoretical Biology*, 156:293–307.

Gorelick, R., Bertram, S., Killeen, P., and Fewell, J. (2004). Normalized mutual entropy in biology: quantifying division of labor. *American Naturalist*, 164:678–682.

Hölldobler, B. and Wilson, E. (1990). *The Ants*. Harvard University Press, Cambridge, Massachusetts.

Huang, Z. and Robinson, G. (1992). Honeybee colony integration: worker-worker interactions mediate hormonally regulated plasticity in division of labor. In *National Academy of Sciences*, volume 89, pages 11726–11729.

Julian, G. (1999). *Genetic Variation and Task Organization in the Desert Leaf-Cutter Ant, Acromyrmex versicolor.* PhD thesis, Ariz. State Univ., Tempe.

Kohler, T., Johnson, C. D., Varien, M. D., Ortman, S., Reynolds, R., Kobti, Z., Cowan, J., Kolm, K., Smith, S., and Yap, L. (2007). *The Model-Based Archaeology of Socionatural Systems*, chapter Settlement Ecodynamics in the Prehispanic Central Mesa Verde Region., pages 61–104. SAR Press, Santa Fe, New Mexico.

Larsen, J. (2001). Specialization and division of labour in distributed autonomous agents. Master's thesis, University of Aarhus.

Lavezzi, A. M. (2003). Complex dynamics in a simple model of economic specialization. Technical report, University of Pisa.

Merkle, D. and Middendorf, M. (2004). Dynamic polyethism and competition for task in threshold reinforcement models of social insects. *Adapt. Behav.*, 12:251–262.

Milgram, S. (1967). The small world problem. *Psychology Today*, pages 60–67.

Murciano, A. Millán, J. Z. J. (1997). Specialization in multiagent systems through learning. *Biological Cybernetics*, 76(5):375–82.

	10	50	100	500	1000
2	4.42 ± 2.2	2.95 ± 0.5	2.95 ± 0.37	2.82 ± 0.11	2.83 ± 0.14
4	3.62 ± 1.02	3.06 ± 0.33	2.93 ± 0.27	2.92 ± 0.11	2.91 ± 0.09
10	2.99 ± 0.47	2.75 ± 0.24	2.88 ± 0.1	2.87 ± 0.08	2.88 ± 0.05
20	2.99 ± 0.27	2.81 ± 0.16	2.83 ± 0.11	2.83 ± 0.06	2.81 ± 0.04

Table 3: Average DOL multiple and standard deviation with IR = 0.05 in random networks.

	10	50	100	500	1000
2	5.34 ± 2.79	3.79 ± 0.74	3.65 ± 0.43	3.57 ± 0.16	3.59 ± 0.15
4	5.04 ± 1.57	4.48 ± 0.37	4.23 ± 0.29	4.34 ± 0.13	4.28 ± 0.12
10	4.98 ± 0.70	4.65 ± 0.27	4.79 ± 0.2	4.81 ± 0.12	4.83 ± 0.08
20	5.25 ± 0.35	4.84 ± 0.13	4.92 ± 0.19	4.93 ± 0.08	4.91 ± 0.05

Table 4: Average DOL multiple and standard deviation with IR = 0.5 in random networks.

Naug, D. and Gadagkar, R. (1999). Flexible division of labor mediated by social interactions in an insect colony–a simulation model. *Journal of Theoretical Biology*, 197(1):123 – 133.

O'Donnell, S. (1996). Rapd markers suggest genotypic effects on forager specialization in a eusocial wasp. *Behav. Ecol. Sociobiology*, 38:83–88.

Page, R., Erber, J., and Fondrk, M. (1998). The effect of genotype on response thresholds to sucrose and foraging behavior of honey bees (apis mellifera l.). *J. Comp. Physiol. A*, 182:489–500.

Ravary, F., Lecoutey, E., Kaminski, G., Châline, N., and Jaisson, P. (2007). Individual experience alone can generate lasting division of labor in ants. *Curr. Biol.*, 17:1308–1312.

Robinson, G., Page, R., Strambi, C., and Strambi, A. (1989). Hormonal and genetic control of behavioral integration in honey bee colonies. *Science*, 246:109–112.

Shannon, C. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27:379–423, 623–656.

Spencer, A. J., Couzin, I. D., and Franks, N. (1998). The dynamics of specialization and generalization within biological populations. *Journal of Complex Systems*, 1(1):115–127.

Traniello, J. and Rosengaus, R. (1997). Ecology, evolution and division of labour in social insects. *Animal Behaviour*, 53:209–213.

Waibel, M., Floreano, D., Magnenat, S., and Keller, L. (2006). Division of labour and colony efficiency in social insects: effects of interactions between genetic architecture, colony kin structure and rate of perturbations. In *R Soc Lond B*, number 273, pages 1815–1823.

Young, A. A. (1928). Increasing returns and economic progress. *The Economic Journal*, 38:527–542.