Optimal Organizational Size in a Stochastic Environment with Externalities

Bennett Levitan [1]

José Lobo [2]

Stuart Kauffman [3]

Richard Schuler [4,5]

March 26, 1999

[1] Bios Group LP, 317 Paseo de Peralta, Santa Fe, NM 87501, 505-992-6713,
e-mail: bennett.levitan@biosgroup.com.

[2] Graduate Field of Regional Science, Cornell University, 108 West Sibley Hall, Ithaca, NY 14853, 607-255-5385, e-mail: jl25@cornell.edu.

[3] Bios Group LP, 317 Paseo de Peralta, Santa Fe, NM 87501, 505-992-6717,
e-mail: stu@biosgroup.com.

[4] Department of Economics and School of Civil and Environmental Engineering, 422 Hollister Hall, Cornell University, Ithaca, NY 14853, e-mail: res1@cornell.edu.

[5] To whom correspondence should be addressed.

Submitted to *Organization Science*.
The authors gratefully acknowledge research support provided by the Santa Fe Institute and Bios Group, LP.

Optimal Organizational Size in a Stochastic Environment with Externalities

Abstract

In this study, we explore the relationships among group size, the extent of interactions with other groups, and group performance in a stochastic environment. We have developed a modeling framework which allows the connections among the individual members constituting a group and the connections between groups (externalities) to be tuned independently. The search for improved group configurations is modeled as a random walk on a space of possible configurations whereby agents in a group periodically have the opportunity to accept or reject random changes in their characteristics. By controlling which groups have externalities with which other groups, we can manipulate the topology of the problem – the web of interactions within and between groups. We present numerical results showing that optimal group size relates to the magnitude of externalities and the length of the search period. Our main result suggests that for short search periods, large organizations perform best, while for longer time horizons, the advantage accrues to small sized groups with a small number of (but not no) externalities. However, over these long time horizons, as the extent of externalities increases, modest increases in group size enhances performance. Under all circumstances, organizations that perform best border on a regime of chaotic behavior. The results have applications at both the micro-scale for the size and structure of production units and at the macro-level for understanding the relationships between groups and communities in a hierarchy of market networks.

1 Introduction

A salient feature of economic behavior is the tendency of individual economic agents to cluster themselves into groups, organizations, firms or similar structures. These economic groupings affect each other's performance in a myriad of ways not mediated by market mechanisms, giving rise to the phenomena of externalities. Indeed, a prominent answer to the question, "why are there firms?" is that they exist primarily to internalize externalities (Coase 1937, Williamson 1985). Another widely documented feature of economic life is the persistent variety of firm sizes, both within the same and across different industries (see, *e.g.*, Acs and Audretsch 1988, Schmalense 1989, Klepper and Graddy 1990, Hannan and Carroll 1992, Hansen 1992, Bailey, Bartelsman and Haltiwanger 1994, Davis, Haltiwanger and Schuh 1996, Sutton 1997). Although both externalities and variation in firm size have been studied separately, how they affect firm performance jointly is not understood.

In this study we specifically consider: (1) How does optimal group size change as the extent of externalities among distinct groups varies?, and (2) How does optimal group size depend upon the length of the group's operational cycle? We have devised a modeling framework which allows connections among individual agents constituting a group and connections between groups (externalities) to be tuned independently. These externalities can have positive or negative effects on group performance. By controlling which groups have externalities with which other groups, we can manipulate the web of interactions within and between groups.

We assume that groups (or organizations or firms) operate in a stochastic environment in which the payoffs resulting from the choices made by individual agents are randomly assigned. The only intelligence attributed to an agent is its ability to recognize a better or a worse choice after its payoff has been determined, but there is no foresight nor cost of search. A consequence of having a large number of agents joined together in a group or organization is that every time one member of the group considers an alternative behavior or state, that change affects the payoffs of all other cooperating agents. Since the environment is stochastic, there is no guarantee that an improvement available to one agent in the group will automatically be translated into improvements for all. Rather, the benefits of forming groups arises merely because a larger group is afforded the opportunity of evaluating more alternatives. As the group's size increases, since each new alternative considered by one member alters the payoffs of all others, the total number of options available increases. Which options are accepted will depend upon the group's acceptance rule. Throughout this analysis, the acceptance rule used is that state changes experienced by a group's members are accepted only if they improve the entire group's payoff. Thus the improvement process that determines optimal group size is trial and error in nature based upon benefits to the entire group; however, this acceptance criterion does not guarantee maximum benefits to the economy as a whole in the face of externalities.

In section 2, we define our modeling framework and assumptions. In section 3, we provide an order statistics-based qualitative basis for interpreting the results, reported in section 4. We discuss the results and future directions in section 5.

2 Modeling Framework

2.1 Agents and Groups

In our setting, economic activity is performed simultaneously by M *groups*, each composed of L individual *agents*. There are a total of N agents in the economy with $N = M \cdot L$ (see Figure 1). Groups may correspond to work teams within a firm, divisions within a multi-unit firm, entire firms or even entire economies depending on how the model parameters are chosen and interpreted. Agents correspond to components within economic organizations, such as persons within a work team, production processes within a plant or communities within a metropolitan economy.

Each agent's activity is characterized by one of S discrete states. A state represents different choices of economic behavior undertaken by an agent, such as selecting a particular production technology or investment strategy. The decision to characterize economic agents as occupying discrete states departs significantly from the continuity assumptions typically made in economic analysis. However, a little reflection on the nature of economic agents and economic decisions should convince one of the ubiquitous presence of indivisibilities (see, for example, the discussion in Reiter and Sherman (1962, 1965) and Scarf (1981a,b, 1994)). In

Level	State variable	Payoff variable
agent	$a^c_{i,m}$	$\phi^c_{i,m}$
group	$\vec{\omega}^c_m$	ϕ^c_m
economy	$\vec{\Omega}^c$	Φ^c

Table 1: State and payoff variables in the model.

fact, the actual analytic abstraction is to characterize the choices facing economic agents as a smooth continuum of tradeoffs among connected states.

In our framework a *configuration* denotes a specific assignment of states to every agent in the economy (or, equivalently, the economic behaviors adopted by all agents in the economy). Let Ω denote the space of all possible configurations. We define $a^c_{i,m}$ as the state of the i'th agent in group m in configuration c, $\vec{\omega}^c_m$ as the state of group m in configuration c, and $\vec{\Omega}^c$ as the state of the economy as a whole in configuration c (see Table 1). Notationally,

$$\begin{aligned} a^c_{i,m} &\in \{0, \ldots, S-1\}, \\ \vec{\omega}^c_m &= [a^c_{1,m},\ a^c_{2,m}, \ldots a^c_{L,m}], \\ \vec{\Omega}^c &= [\vec{\omega}^c_1,\ \vec{\omega}^c_2,\ \ldots,\ \vec{\omega}_M], \\ \vec{\Omega}^c &\in \Omega, \end{aligned} \tag{1}$$

where $i = 1 \ldots L$ and $m = 1 \ldots M$. Since all groups are composed of L agents and each agent has S possible states, there are S^{ML} possible configurations for the economy.

Performance of a group may be measured using a variety of metrics such as time, costs, profits or any other relevant measure of group welfare. We use the generic term *payoff* to refer to the property a group seeks to optimize. We define $\phi^c_{i,m}$ as the contribution of the i'th agent in group m in configuration c, ϕ^c_m as the payoff of group m in configuration c, and Φ^c as the payoff of the economy as a whole in configuration c (see Table 1). These payoffs are related by

$$\phi^c_m = \frac{1}{L} \sum_{i=1}^{L} \phi^c_{i,m} \tag{2}$$

$$\Phi^c = \frac{1}{M} \sum_{m=1}^{M} \phi^c_m. \tag{3}$$

A group's payoff is the sum of the contribution of its members, and the economy's payoff is the sum of the payoffs of its constituent groups. In this work, we restrict ourselves to additive interactions among the payoffs of individual agents inside a group (an assumption that can be regarded as a first order approximation to more general interactions). All payoffs generated range from 0 to 1 and can be thought of as a fraction of some theoretically maximum payoff.

The payoff of an agent depends on its own state and that of other agents, both within the same group and possibly within other groups when there are externalities. We say agent

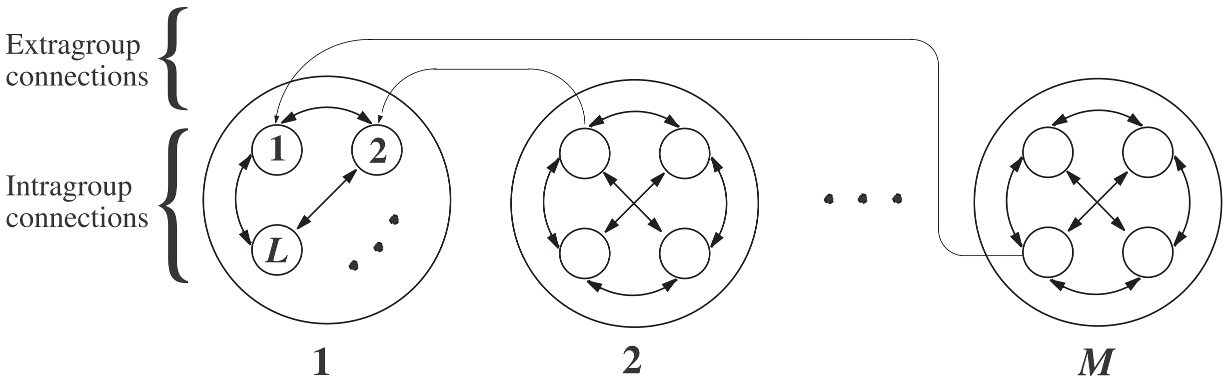

Figure 1: Depiction of connections in model for $L = 4, J = 2$ and M groups. Agents are represented by small circles. Each large circle represents a group composed of L agents. Lines within a large circle indicate intragroup connections. Lines above the groups represent extragroup connections (*i.e.*, "externalities"). Only the extragroup connections to group 1 are shown.

$a_{i,m}$ is *connected* to agent $a_{j,k}$ if $\phi_{j,k}$ depends on $a_{i,m}$. For example, actions by members within a group (*e.g.*, whistling while at work, using a wrench rather than pliers, deciding to use JAVA instead of C++ as a programming language) have direct impacts on the payoffs of other group members. Note that connections are not necessarily symmetric: one agent may be connected to a second without the second being connected to the first. By definition, every agent is connected to itself.

Many of our results below hinge on the difference between *intragroup* connections and *extragroup* connections (or *externalities*). The operational significance of the distinction between intra- and inter-group connections between agents is caused by the decision rules whereby individual agents accept or reject changes in their states. The fundamental nature of most groups or organizations is that intra-group impacts are the primary concern. An example of extragroup connections might be when a person in a company that competes with my company (an agent in another group) pollutes my water supply or develops a better way to paint their cars, hence altering my payoffs.

Let K be the number of intragroup connections per agent, excluding the connection from itself, and J be the number of extragroup connections per group. If $J = 3$, there are three possibilities for the externalities of a group: (i) One agent is affected by three external agents; (ii) one agent is affected by two external agents and a second agent is affected by one external agent; or (iii) three agents are affected by one external agent each. Thus, agent $a_{i,m}$ will be affected by itself, K connections from agents within group m, and by anywhere between 0 and J connections from agents in other groups. In our simulations the outside agents that affect a given group are chosen at random from within a randomly-chosen group. Once $J \geq L$, some agents depend upon $\lfloor J/L \rfloor$ external agents and others depend upon $\lfloor J/L \rfloor + 1$

external agents, so that the total is J [1]. In this manner, the level of external connections among groups is distributed as uniformly as possible.

We assume that each agent within a group is "maximally connected" within that group; that is, every agent is connected to every other agent within its group, therefore $L = K + 1$ (see Figure 1). This assumption corresponds to an organizational setting in which the choices of each individual affects the payoff of every other individual in the organization. An agent's payoff thus depends upon its own state and on the state of between $L - 1$ and $L - 1 + J$ other agents. By assuming that essentially arbitrary interactions are possible among agents, we can assign payoff values to the $\phi^c_{i,m}$, at random – they are drawn independently from the uniform $[0, 1]$ distribution in our simulations. While a group itself has only S^L possible configurations, it has $S^{(L+J)}$ possible payoff values due to the external connections. In this manner a "table" of payoff values is constructed for each agent. These connections and values define the entire economy in our model.

2.2 Group Search for Greater Payoffs

In our simulated world the stochastic nature of economic environments is represented by the random opportunities for agents to change their states. In every round (generation) of the simulation, a randomly-selected agent is given the opportunity to adopt a different state whose payoff consequences are selected from a fixed table of random payoffs. These payoffs are drawn from the uniform $[0, 1]$ distribution. In this manner, if any particular group size should emerge as optimal, that result will not be biased by any arbitrary, or systematic, assignment of payoffs. Because each group is maximally connected, each agent's decision to accept or reject an alternative state affects the payoffs of all other members of its group, plus all externally connected agents. However, the agent decides to accept the new state only if that change improves the combined payoff of its group, irrespective of the decision's effect on other groups or on its own individual payoff. Thus in a practical application, compensation schemes would have to be implemented whereby individual agents shared any increased group payoff.

Due to externalities, there will be many interactive effects on the payoffs of other groups. In the course of attempting to optimize its own group's payoff, an agent's decisions will constantly alter the payoff of other groups. The resulting churning of agent states (analogous to companies changing their policies in the light of other companies' decisions) is responsible for much of the interesting behavior in our model.

To summarize, the group search process is as follows:

1. The economy starts a production run at $t = 0$ with an initial configuration $\vec{\Omega}^t = \vec{\Omega}^0$ (using time periods to index configurations).

2. During each step of a production run, a randomly chosen agent in group m is given the opportunity to adopt a different state. If $S > 2$, the new state is chosen at

[1] $\lfloor x \rfloor$ is the largest integer i such that $i \leq x$.

random; if $S = 2$, the state is switched between 0 and 1. The economy's tentative new configuration is $\vec{\Omega}^{\tilde{t}}$.

3. If payoff $\phi_m^{\tilde{t}} > \phi_m^t$, the group adopts the new configuration in the next production period, so that $\vec{\Omega}^{t+1} = \vec{\Omega}^{\tilde{t}}$; otherwise, the economy keeps the same economic configuration, such that $\vec{\Omega}^{t+1} = \vec{\Omega}^t$

4. Increment t and return to step 2.

Note that the simulation is a Markovian process of order one and does not have an automatic termination point. The time for one pass through steps 2 – 4 above is referred to as a *generation.*

When deciding whether or not to accept a change in the state of an agent, a group compares a newly sampled configuration only with the currently-held economic configuration. The groups' search, propelled by the choices made by individual agents, can therefore be understood as a *random adaptive walk* in the space of possible economic configurations, Ω (Kauffman 1993). The trial and error process for finding higher group payoffs imposes no decision cost on the agents since the agent is allowed to return to the previous state if no improvement is realized. The limited memory attributed to the groups' search process carries another notable implication: a group derives no benefit from the accumulation of unsuccessful sampling.

Through time, the economy might move away from a configuration with higher aggregate payoff towards one with lower aggregate payoff if such a move benefits the group that experienced a change in the state of one of its members. Thus this group-based selection criteria does not necessarily lead to a search process that is Pareto improving in terms of the entire economy. If the implemented search procedure were guided by maximizing the economy's aggregate payoff, the random walk would eventually settle on a local optima for the economy as a whole. In contrast, a walk guided by the imperative to maximize individual groups' payoffs can result in the walk not settling on a local optima for each group. Indeed, as discussed in Section 4.4, the walk can verge on the chaotic as different groups make conflicting decisions to accept changes in the state of agents. Figure 2 depicts the contrast between an economy-wide choice rule and a group-based choice rule.

3 Order Statistics and Group Payoffs

A statistical analog to the search process faced by each group is selecting the largest order statistic from a sequence of samples. Where $J = 0$, the group payoffs associated with each group's state can be thought of as random samples drawn independently from identical distributions. As an illustration, consider the simple limiting case when $L = 1$, $J = 0$ and $S = 2$ for M groups. Each agent acts independently and ultimately will try both of its allowed states and choose the better of the two (the second order statistic[2]). Starting

[2] The n'th order statistic is the largest value in a sample of size n.

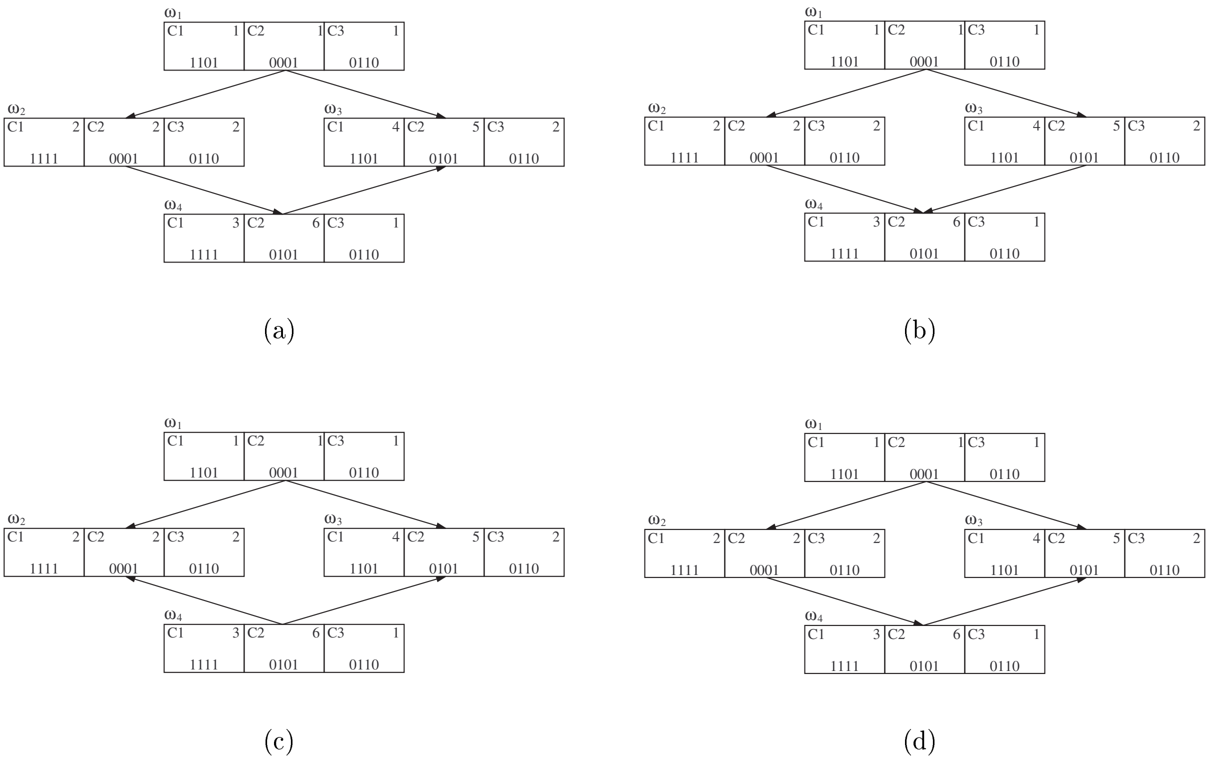

Figure 2: Depiction of simulation dynamics for economy-wide and group-based search rules. The example shows an economy with $M = 3$ groups and $L = 4$ agents per group. Each large rectangle represents the economy, and the small rectangles labeled C1 – C3 represent the individual groups. Within each group, the four binary digits indicate the state of that group and the number in the upper right indicates the payoff to that group in that state. Four of the 2^{12} states of the economy are show, $\omega_1 \ldots \omega_4$. The states differ by the flipping of one agent's bit. The arrows indicate the flow between states of the economy dictated according to decisions by individual groups or the economy as a whole. Dynamics dictated by changes in: (a) group 1's payoff; (b) group 2's payoff; (c) group 3's payoff; and (d) the economy's payoff.

from any initial configuration, the economy will eventually arrive at its globally optimal configuration, since each agent will ultimately find its best state when there is no correlation between groups ($J = 0$). The economy's expected average payoff at this optimum is 0.67, the expected value of the larger of two independent draws from a uniform [0,1] distribution[3].

Consider now the case when $L = 2$ with $J = 0$ and $S = 2$: The outcomes for two agents within the same group are linked. There are four possible states for each group [$\vec{\omega}_g \in \{00, 01, 10, 11\}$]. The best result is to pick the highest of the four payoffs; however, since the payoffs are the average of two random selections drawn from a uniform distribution, the average payoff for two connected agents can be shown to follow a piecewise truncated Beta probability distribution (triangular in the case where $L = 2$). While selecting a fourth order statistic (rather than the largest of two, as in the $L = 1$ case) suggests a better expected payoff, choosing that statistic from a symmetric piecewise truncated Beta distribution, instead of from the uniform density, has the offsetting effect of forcing the probability mass away from the extremes and toward the mean.

As group size increases further, the upper bound to the group's average payoff will rise because an ever higher order statistic is searched for. However, as L increases, the distribution underlying that order statistic is the computed average of an ever larger group of independent choices. Since the "sample" whose average is being computed is drawn from the uniform probability distribution function with finite mean and variance, the Central Limit Theorem applies, so that the density of the sample means approaches a normal distribution with an increasingly declining variance. Thus as L increases, the possible greater payoff resulting from the opportunity to select an ever higher order statistic is counterbalanced by the fact that underlying distribution is converging towards the population mean of 0.5.

Furthermore, as L increases, the probability that the adaptive walk will become stuck at a local optimum increases. Consider a simple example with $L = 2$ (Figure 3). If the group's search begins at (0,0), moves to either (0,1) or (1,0) will not be taken since (0,0) has a higher average payoff. By comparison, if moves were dictated solely by the payoff to the first agent, the walk would move from (0,0) to (1,0) to (1,1). With moves accepted or rejected based upon group payoff, the only way that agents can locate at the (1,1) configuration, the global optimum, is if (0,0) is not the initial configuration. As L increases, these local optima become more and more prevalent in the search space. Thus the possibility of getting stuck at a local optimum also exerts a strong negative force on the system-wide average outcome as L increases.

4 Numerical Results

4.1 Characteristics of Simulated Searches

Searching for optimal configurations is an extremely difficult combinatorial problem for which analytic solutions are not known, except in several limiting cases (*e.g.*, $L = 1, J = 0$).

[3] The expected value of y_n, the largest of n random variables drawn independently from a uniform [0,1] distribution, equals $n/(n+1)$. For a treatment of order statistics see Mood, Graybill and Boes (1974).

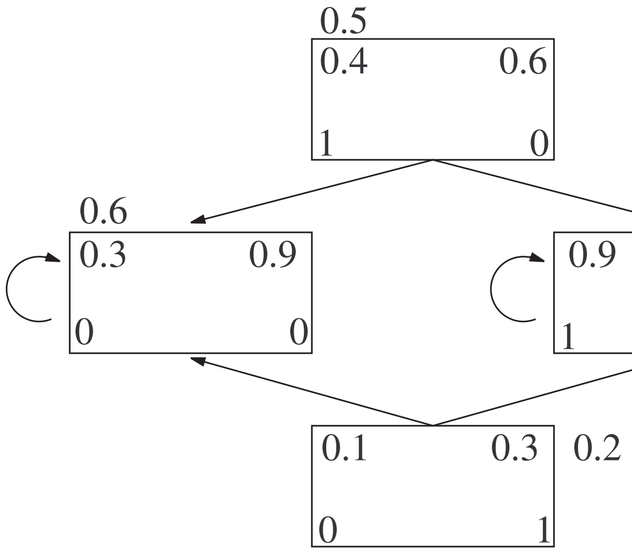

Figure 3: Possible payoffs for a group of size $L = 2$ with $S = 2$ and $J = 0$. The decimal numbers inside each box denote the payoffs to agents 1 and 2 under the given configuration; the decimal number outside each box denotes the group's average payoff. The arrows indicate state transitions that increase group average payoff.

Therefore, we determine results numerically, an approach used extensively in related work (Kauffman and Levin (1987) on the "NK" model, which allows tunable degrees of interagent connections but no concept of groups within the population.) Previous analyses predicated on the NK paradigm include: Lobo and Schuler (1997) on optimal organizational size, Auerswald, Kauffman, Lobo and Shell (1999) on learning-by-doing, and Kauffman, Macready and Dickinson (1985) on a variant with distinct patches; but none of these analyses provide a distinction or independent tuning between intragroup and extragroup connections.

We have simulated the group search process for economies with $N = 100$ agents, $S = 2$ possible states per agent and group sizes L varying from 1 to 11. Since the agents within groups are maximally connected, the number of groups $M = N/L$ varies from 9 to 100. We consider a large range of external connections, from 0 to 50. Simulations were run for 2000 generations, which generally proved sufficient to show steady state behavior of the dynamic economy. Ensemble averages were computed by compiling cumulative statistics for 500 different economies (sets of connections and payoff tables) with 2 simulations per economy. (The simulation code is written in ANSI C and is available upon request from the authors.) In all results below, "average group payoff," $\langle \Phi^c \rangle$, refers to the payoff of a group averaged over all groups and all simulations at the corresponding generation.

4.2 Optimal Group Size

Figure 4 plots the average group payoff per generation for selected values of L and J. Several features of the results stand out. First, in these "trial and error" economies, the familiar S-shaped learning curves emerge for nearly all configurations. Second, when there are no external connections ($J = 0$), or when the magnitude of external connections is low ($J = 4$), it is advantageous, in the long run, for groups to form, but to be small in size, since the payoffs of larger sized groups increase at slower rates toward the end of the search process. Once the level of external connections is high ($J = 12$), a group size of $L = 5$ is associated with higher average payoff at the end of the search; as the level of external connections is increased even further ($J = 18$), greater payoff accrues to the larger organizations. At a level of externalities of $J = 50$ (not shown in Figure 4), the optimal organizational size, in the long run, is $L = 11$. For larger group size, average group payoffs are larger when $J > 0$, and the larger the group size, the larger the level of external connections associated with the highest group payoff.

These results suggest that larger organizational size provides a buffer against the many stochastic shocks brought about by external connections. An alternative interpretation is that in large organizations where operational changes are decided upon incrementally, the probability of getting stuck at a locally optimal configuration, or globally sub-optimal configuration, is high. In these cases, an increased number of external interactions enhances the likelihood of being dislodged from a sub-optimal position and of being able to continue the search for a global optimum.

Figure 5 shows the average group payoff achieved as a function of L after adaptive walks of varying duration ($g = generation$); the highest payoff achieved per generation is indicated with an "$\times$." Across a broad range of values for external connections ($J = 0,\ 4,\ 12,\ 18$), larger-sized groups achieve higher payoffs when adaptive walks are of brief duration. As the duration of the groups' search increases, however, smaller organizational size proves to be more beneficial. (When the level of external connections is very high ($J = 50$), however, larger group-size is optimal regardless of the duration of the search.)

These results suggest that for organizations engaged in activities characterized by short turn-around periods or brief production cycles, large organizational sizes may be preferred to small ones [Summarized in Fig. 6]. The reason is that in searches (or production runs) of short duration, the number of alternative configurations that can be sampled from any arbitrary starting position by each group is $(S-1)L+1$ and for the economy as a whole is $LM(S-1)+1$. So with J small, the order statistic analogy suggests that larger groups get more chances to find a better outcome, initially. As the duration of the walk increases, however,the chances of getting stuck at a sub-optimal configuration increases with group size; whereas with $L = 1$ and J small, eventually nearly all agents find the larger of their two order statistics.

Figure 7 plots average group payoff as a function of group size (L) and magnitude of external connections (J) after 50, 100, 700 and 2001 generations. The results described above can all be seen here, but certain temporal patterns are more evident. For short production cycle business applications (small number of generations), there is little to distinguish over

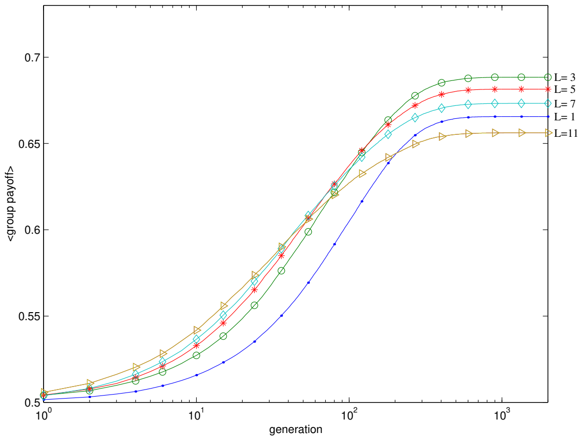

(a) $J = 0$

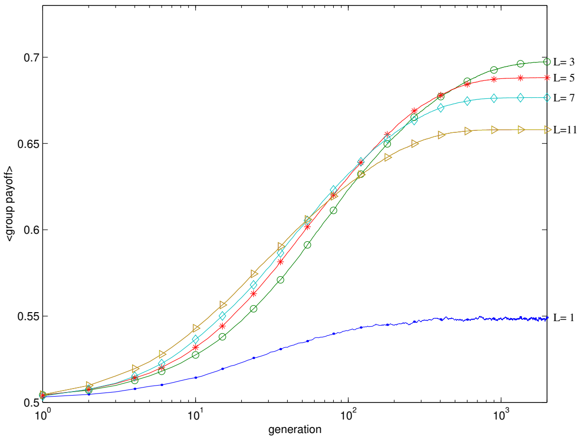

(b) $J = 4$

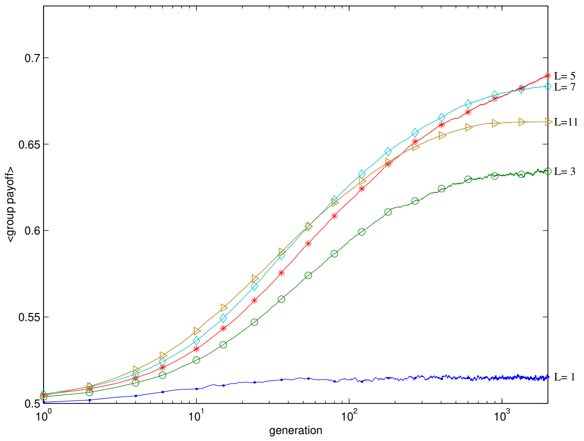

(c) $J = 12$

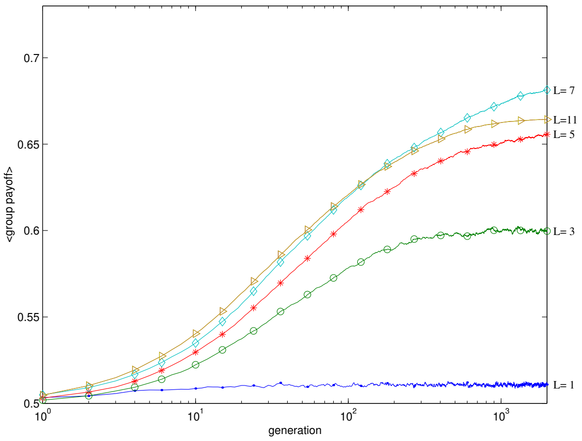

(d) $J = 18$

Figure 4: Average group payoff over time (generations) for different group sizes, L, and four different levels of externalities $J = 0, 4, 12$ and 18.

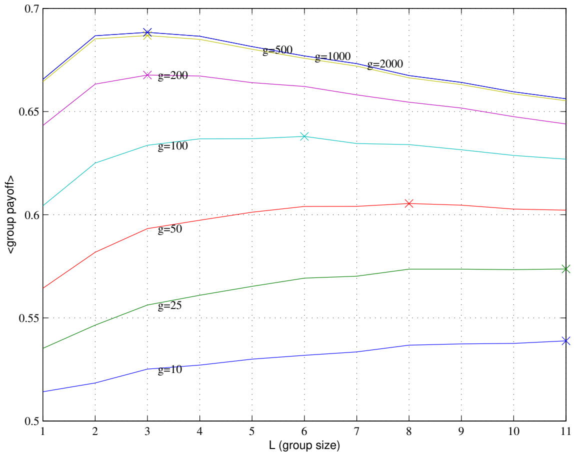

(a) $J = 0$

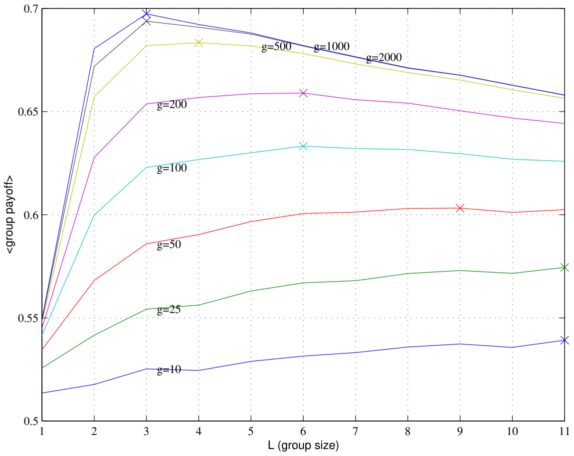

(b) $J = 4$

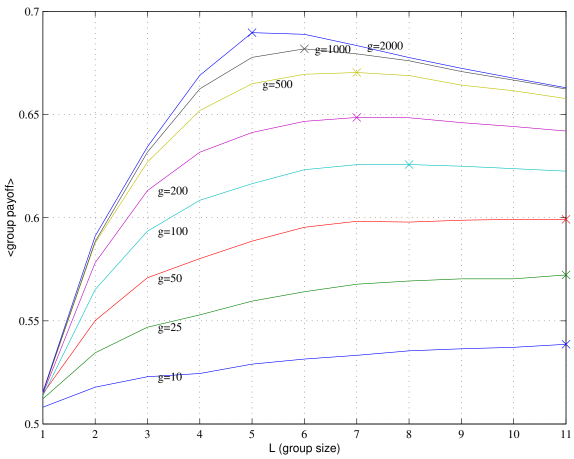

(c) $J = 12$

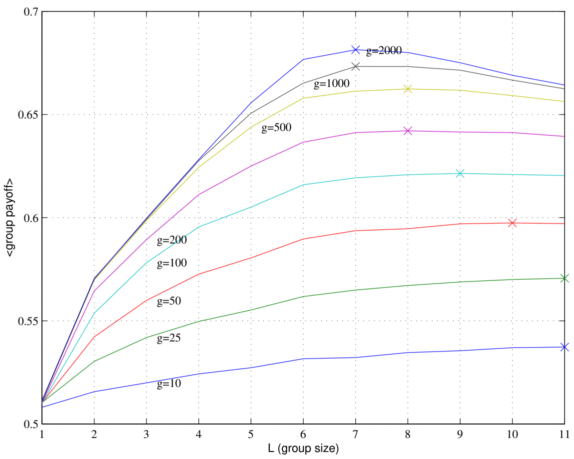

(d) $J = 18$

Figure 5: Average group payoff per group size (L) for varying generations (g) and four different levels of externalities $J = 0, 4, 12$ and 18.

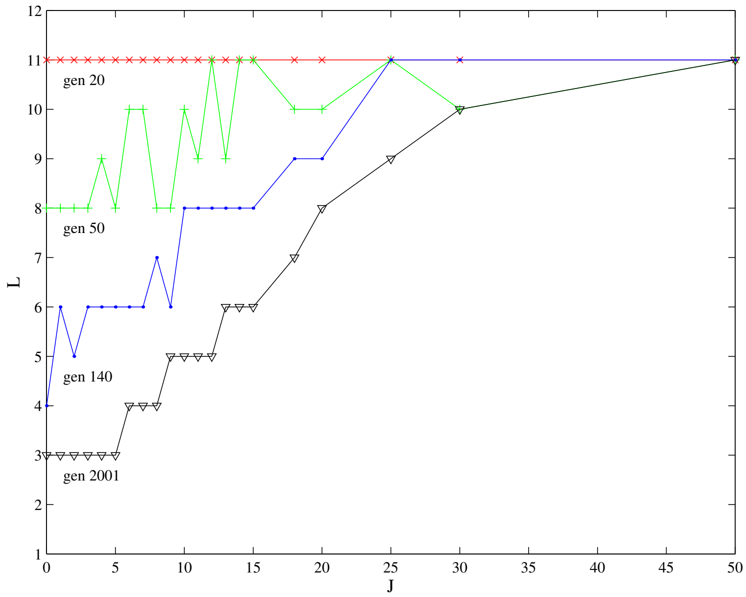

Figure 6: Optimal group size (L) corresponding to varying levels of externalities (J) after 20, 50, 140 and 2001 generations.

a large range of L and J. There is a modest monotonic improvement as L increases, though the interactions with J are subtle for high L. If one chooses a poor organization size for the "J of his business," the drop in payoff is minor – though the payoff is small to begin with. As the production cycle lengthens, a strong ridge reflecting values of L and J for reasonable to good performance becomes more and more evident. In these cases, choosing the wrong group size can be catastrophic.

It is of interest to consider what is the best (L, J) combination for different time horizons. Figure 8(a) plots the best combination (L, J) for different generations. The trajectory of optimal (L, J) combinations moves from high L, low J, at short walks, to low L and higher J for larger generations. The move towards larger group sizes at the longest time horizons considered (2001 generations) is curious, and merits further investigation. We compare these results to those for a wider range of production choices per agent ($S = 5$ in Fig. 8(b)) below.

4.3 Transition Between Ordered and Chaotic Regimes

Figure 9 illustrates the average rate at which agents, given the opportunity to choose new states, continue to do so, as a function of L and J for four different durations of activity ($g = 51, g = 100, g = 700$ and $g = 2001$). This *flip rate* per site measures the extent to which the trial and error search process continues to induce changes in the state of agents over time, and is calculated as a running average over the previous 50 generations of all generations larger than $g = 50$. For longer production cycles, as in Figs. 9(c) and 9(d), the economy exhibits two broad regimes in the LJ plane. An "ordered" regime emerges where agents stop flipping for a range of low J, high L, and a "chaotic" regime where agents keep

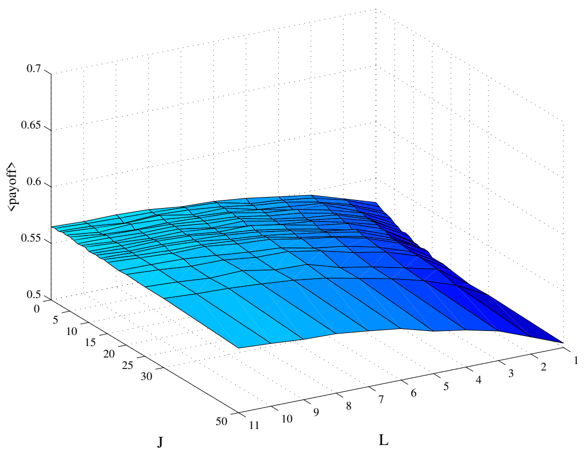

(a) $g = 20$

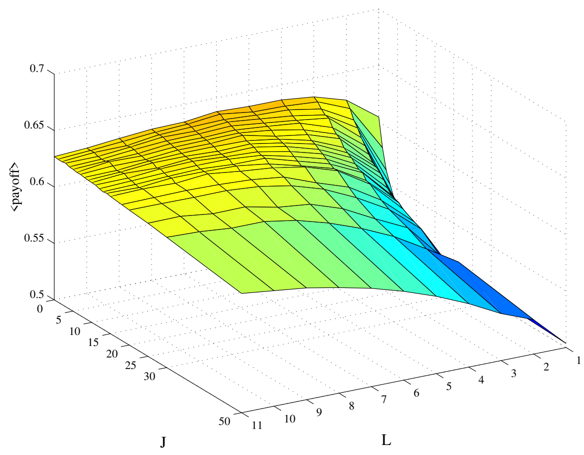

(b) $g = 100$

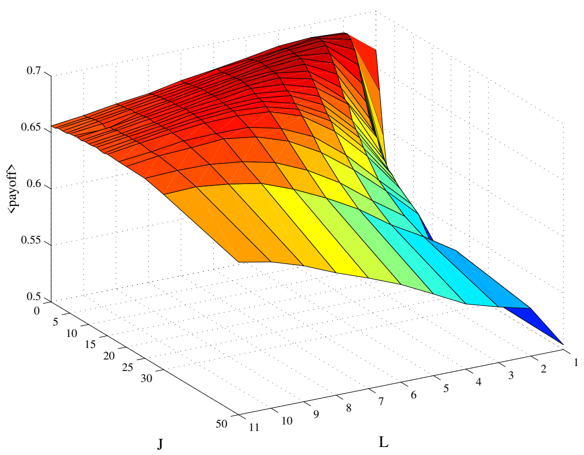

(c) $g = 700$

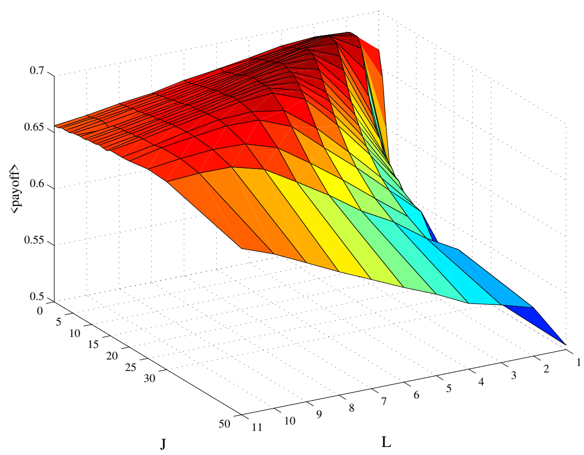

(d) $g = 2001$

Figure 7: Average group payoff as a function of group size, L, and level of externalities, J, after 20, 100, 700 and 2001 generations.

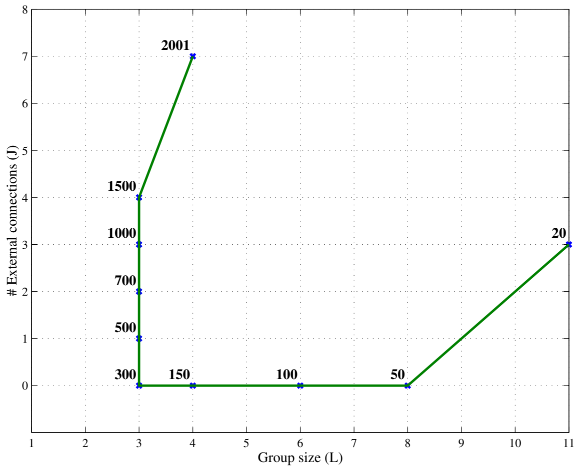

(a) $S = 2$

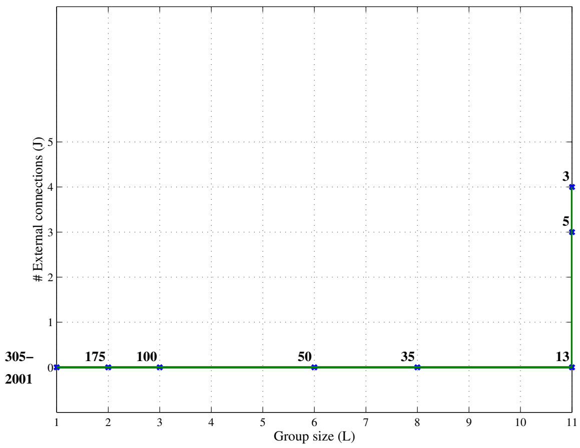

(b) $S = 5$

Figure 8: Trajectory of the combination of group size (L) and level of externalities (J) associated with the highest average group payoff (the number associated with a point denotes the roughly the earliest generation at which that (L, J) combination became optimal). (a) $S = 2$; (b) $S = 5$.

flipping rapidly for high J, low L.

The payoff of any given group depends on the choices of its constituent agents to accept or reject alternative states, but in the presence of externalities adaptive moves by other groups may impinge on its payoff. In the ordered regime each group finds a local or global optimum that is mutually consistent with the optima attained by the other groups that affect its payoff and vice versa. Once these mutual optima are attained, further changes in the states of agents in all groups stops. By contrast, in the chaotic regime, adaptive moves by one group continue to alter the payoffs of other groups faster than those groups can find their own optimal payoffs. The economy remains in persistent chaotic motion as groups chase ever moving optima, and agents keep changing their states. Given the restriction that each group must move incrementally (can change only one agent's state at a time), the ordered regime can be thought of as sets of local Nash equilibria among all the groups in the sense that no individual group has an incentive to accept further state changes in response to changes by another group or to external shocks to itself.

It is particularly interesting to compare figure 9, showing the ordered and chaotic regimes, with figure 7, showing payoff in the LJ plane. Note that in figure 7(c) and 7(d) a curved ridge of high payoff runs diagonally from low L, low J values to high L and high J values. This ridge represents a type of watershed or threshold where payoff decreases down the two slopes flanking the ridge. Comparing this curved ridge of maximal payoff with the corresponding flip rate curves in figures 9(c) and 9(d) suggests that the ridge of maximum payoff in the LJ plane falls in the stable regime parallel and just adjacent to the regime where flip rates rise.

The observation that optimal payoff is associated roughly with a boundary of a region with constant low flipping rate is profoundly interesting. The model economy faces an extremely hard combinatorial optimization problem with a huge configuration space. Achieving optimal payoffs over time seems to require a balance between "exploration" and "exploitation" (Herriott, Levinthal and March 1985, Levinthal and March 1985, Macready and Wolpert 1998 and March 1991). This balance appears to be optimally struck somewhere near the phase transition between the ordered regime and the chaotic regime, where persistent exploration, i.e. flipping, continues. In much of the ordered regime (generally larger group size, L, together with few externalities, J) the economy easily becomes trapped on locally sub-optimal configurations and eventually exhibits stable behavior. These numerical results suggest that if the number of external connections were to increase for an economy whose groups had achieved their optimal sizes, this increase in the number of external shocks might drive the economy toward a chaotic regime. Deep in the chaotic regime, (generally low L, high J), few or no groups find or sustain locations of high payoff.

The comparison between Figs. 7 and 9 also sheds light on optimal group size as a function of time. Where $J > 0$, the moves of other groups can serve the useful role of helping each group avoid becoming trapped on poor local configurations. Figure 9 shows that with increasing time horizons, the flip rates decline dramatically for large group sizes (L) with large numbers of external connections (J); whereas regimes with small L and large J remain chaotic even after 2,000 generations. In the face of a highly inter-connected economy, large sized groups maintain a reasonably high level of payoff and rapidly achieve stable configurations over a wide range of J.

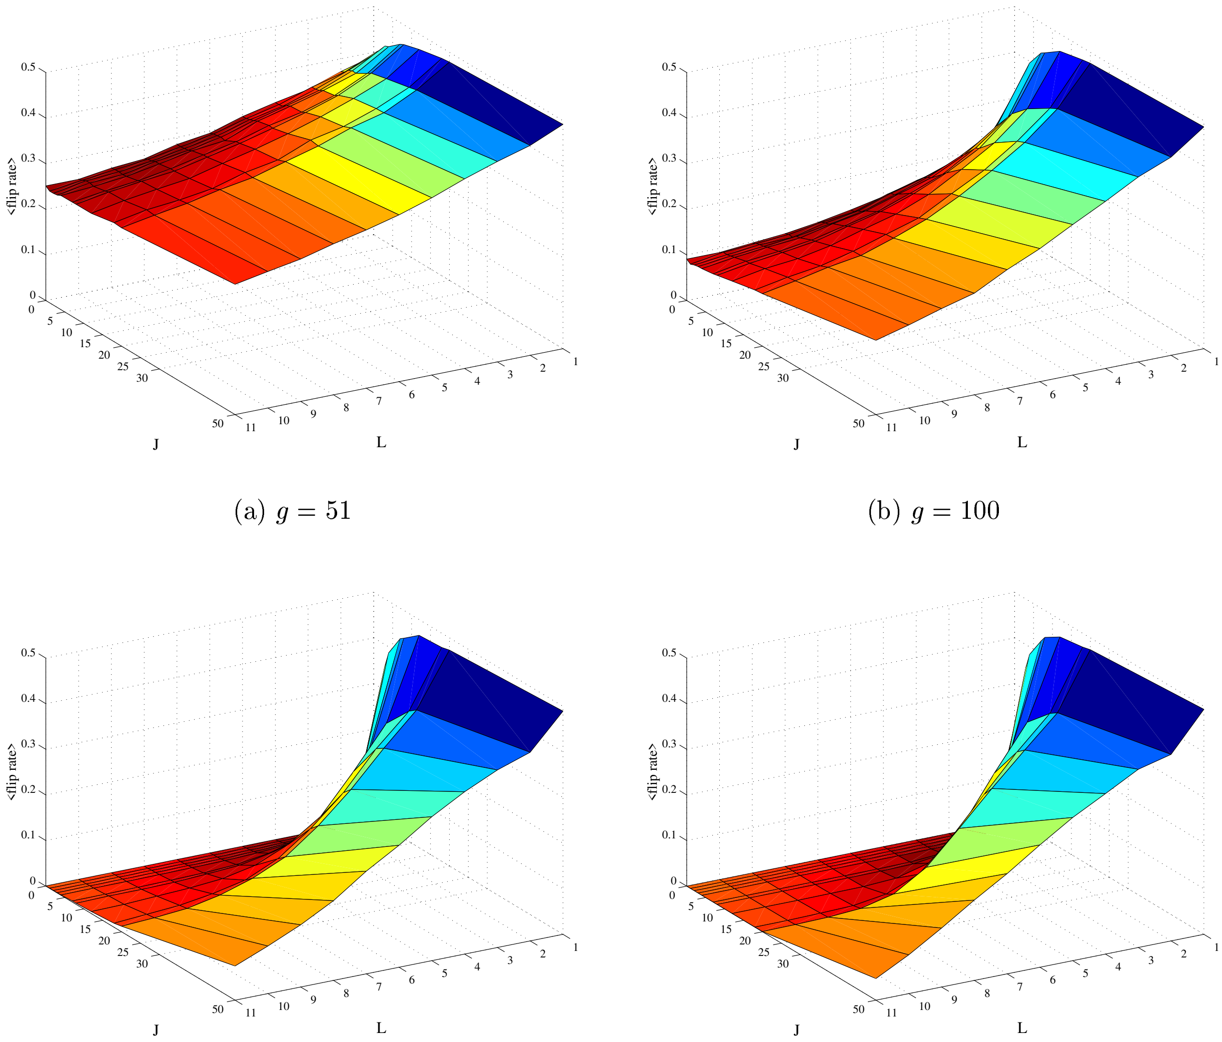

(a) $g = 51$

(b) $g = 100$

(c) $g = 700$

(d) $g = 2001$

Figure 9: Average flip rate as a function of group size, L, and magnitude of externalities, J, for generations 51, 100, 700 and 2001. (See main text for the definition of "flip rate").

5 Discussion

Through the "trial and error" type of search process modeled here we have studied how group size influences average group payoff. We have presented a modeling framework which can be applied usefully to the study of organizations at a variety of scales, ranging from individuals cooperating in groups of varying size to large multi-divisional corporations. The framework is appropriate for organizations where each activity, agent, group or community are performing similar tasks (examples might be a collection of grain farmers, artisans, computer programmers, auditors, or, at a hierarchical level, automotive divisions under General Motors). Intriguing results emerge from our numerical explorations. We show that the relationship between group size and average group payoff hinges on both the time horizon (how long the product cycle is, for example) and the extent of external connections among groups, as might be represented by trade-group exchanges, other information networks, geographic proximity and external competitive or environmental effects.

We show that for short search periods, as represented by few opportunities for each agent to try alternative states, larger organizational sizes fare better because in combination, larger interconnected groups can experiment with a larger number of initial alternatives. As the search period lengthens, however, the advantage of large group size does not keep pace with the continual ability of smaller groups to experiment. This consequence results from the nature of incremental, trial-and-error experiments, which allows large groups a greater variety of initial experiments, but which in the long-run furnishes a much greater chance of being stuck at a locally optimal configuration. By contrast, small groups, having a limited number of configurations, may exploit each random opportunity more fully. The order-statistic analogy is that large groups offer a higher potential order statistic, but with a smaller chance of reaching that pinnacle; whereas, smaller groups offer a potentially smaller top prize, but a greater chance of attaining it.

We also show that high levels of external group connections enhance the payoff of a large group far more than they do for smaller groups, and in fact these externalities actually diminish the payoffs for very small organizations since those groups are constantly being knocked away from their optimal configurations. For large groups that may be stuck at some configuration well below their global optimum, those external shocks stand a much greater chance of improving overall group performance. The magnitude of externalities is an exogenous variable in our model, but suppose that the groups in the model represent individual firms; then the different level of external connections among firms might represent different industries or different economic sectors. Whether or not an economic sector is dominated by small or large firms could result from which firm size is optimal, given the extent to which firms in that sector affect each other's performance. The number of interactions might also be influenced by the geographic location of groups (*e.g.*: rural vs. urban settings).

The search for improved configurations, particularly in an environment that has large numbers of inter-group connections, borders on the "edge of chaos." With complexity comes the risk of uncertainty, and as larger group sizes are formed and operated for long periods of time with trial and error attempts at improvement, the chance of slipping into a chaotic regime diminishes somewhat. Small groups do better in an environment where there are a small number of interconnections among groups. If, however, the complexity of the environment increases (larger J), groups may be thrown into a chaotic regime — unless the groups respond by increasing their size.

A topic warranting further analysis is the effect that an increased number of production states for each agent ($S > 2$) has on optimal group size. A comparison of the very different trajectories of optimal (L, J) combinations for $S = 2$ and $S = 5$ (Figs. 8(a) and 8(b)) makes it tempting to conclude that when agents are *not* restricted to two production choices, banding themselves into large groups may be less desirable. However, additional work is required for $S > 2$, particularly where there are large numbers of externalities, before drawing firm conclusions. Furthermore, with $S > 2$, selecting alternative production states randomly may be unrealistic – using a variety of metrics, it is easy to imagine some of the S states being closer to one another than others. As an example, consider a sequence of production techniques of increasing capital intensity. Under these conditions, technologies might more realistically change incrementally, with state i changing to states $i - 1$ or $i + 1$ only.

Another topic requiring extensive exploration is the modeling of sequential activities (such as assembly line production processes). The essential difference between serial processes and the parallel processes modeled in this paper is that in an assembly line, an increase in group size implies increased specialization (fewer choices) by each agent. In exchange, each agent may gain far greater experience (more repetitions per unit time). The other unique attribute of sequential processes is that the failure of one agent halts the entire process. This aspect suggests that assembly line type processes might be modeled by multiplicative payoffs (geometric means) of the individual agents' contributions, but our approach does not incorporate the advantages of assembly line procedures over parallel processing in any natural way.

Computational and interpretative considerations led us to make the strong simplifying assumption that all groups in the economy are of the same size. As a consequence, the search process is the same for all groups, regardless of their size. It may be more realistic to have smaller groups more agile in their search; whereas larger organizations might experience fewer changes in the states of their agents per unit time. We can hypothesize that under a size-dependent group search rule, smaller groups could find their optimal configuration faster than large ones, while remaining more vulnerable to the stochastic shocks provided by externalities.

Thus, a desirable extension of our framework would be to develop a truly dynamic model describing how individual agents join or leave groups. While the present analysis is dynamic in the sense of describing how the payoffs of a predetermined size group with a predetermined number of external connections evolves over time, at any snapshot of time the analysis represents a static comparison of payoffs among the different conceivable group configurations. What is needed is an explicit dynamic mechanism describing why and how agents and groups try, select and abandon new configurations. This search mechanism could also be used to make non-random explorations of alternate production techniques where $S > 2$ and some metric amongst the states is available. It would be particularly intriguing to use this formal dynamic extension to explore behavior near the ridge of best performance that borders the chaotic region (Figs. 7(d) and 9(d)). With a truly dynamic system of organizational evolution, we could begin to explore those circumstances under which stable optimal organizations are formed, when they are never attained and when the organizations venture past their optimal configurations and enter a regime of perpetual vacillation.

References

[1] Auerswald, P., S. Kauffman, J. Lobo and K. Shell (1998) "The Production Recipes Approach to Modeling Technological Innovation: An Application to Learning by Doing." Center for Analytic Economics Working Paper 98-10, Ithaca: Cornell University. (Forthcoming in *Journal of Economic Dynamics and Control*.)

[2] Acs, Z. and D. Audretsch (1998) "Innovation in Large and Small Firms: An Empirical Analysis," *American Economic Review*, 78, 678-690.

[3] Bailey, M.N., E.J. Bartelsman and J. Haltiwanger (1994) "Downsizing and Productivity Growth: Myth or Reality?" Center for Economic Studies Discussion paper 94-4. Washington, D.C.: U.S. Bureau of the Census.

[4] Coase, R.H. (1937) "The Nature of the Firm," *Economica*, 4, 386-405.

[5] Davis, S., J. Haltiwanger and S. Schuh (1996) *Job Creation and Destruction.* Cambridge, MA: The MIT Press.

[6] Hannan, M. and G. Carroll (1992) *Dynamics of Organizational Populations.* New York: Oxford University Press.

[7] Hansen, J. (1992) "Innovation, Firm Size, and Firm Age," *Small Business Economics,* 1, 37-44.

[8] Herriott, S.R., D.A. Levinthal and J.G. March (1985) "Learning from Experience in Organizations," *American Economic Review,* 75, 298-302.

[9] Kauffman, S. and S. Levin (1987) "Towards a General Theory of Adaptive Walks on Rugged Landscapes," *Journal of Theoretical Biology*, 128, 11-45.

[10] Klepper, S. and E. Graddy (1990) "The Evolution of New Industries and the Determinants of Market Structure," *Rand Journal of Economics*, 21, 24-44.

[11] Levinthal, D.A. and J.G. March (1981) "A Model of Adaptive Organizational Search," *Journal of Economic Behavior and Organization,* 2, 307-333.

[12] Lobo, J. and R.E. Schuler (1997) "Efficient Organization, Urban Hierarchies and Landscape Criteria," *Self-Organization of Complex Structures,* (Schweitzer, F. ed.), 547-558. Amsterdam: Gordon and Breach

[13] Macready, W.G. and D.H. Wolpert (1998) "Bandit Problems and the Exploration/Exploitation Tradeoff," *IEEE Transactions on Evolutionary Computation*, 2, 2-22.

[14] Kauffman, S.A., Macready, W.G. and Dickinson, E. (1994) "Divide to Coordinate: Coevolutionary Problem Solving" *Santa Fe Institute Technical Report*, SFI-94-06-031.

[15] March, J.G. (1991) “Exploration and Exploitation in Organizational Learning,” *Organization Science*, 2, 71-87.

[16] Mood, A.M., F.A. Graybill and D.C. Boes (1974) *Introduction to the Theory of Statistics*. New York: McGraw-Hill.

[17] Reiter, S. and G.R. Sherman (1962) “Allocating Indivisible Resources Affording External Economies of Diseconomies,” *International Economic Review*, 3, 108-135.

[18] Reiter, S. and G.R. Sherman (1965) “Discrete Optimizing,” *SIAM Journal*, 13, 864-889.

[19] Scarf, H.E. (1981a) “Production Sets with Indivisibilities Part I: Generalities,” *Econometrica*, 49, 1-32.

[20] Scarf, H.E. (1981b) “Production Sets with Indivisibilities Part II: The case with Two Activities,” *Econometrica*, 49, 395-423.

[21] Scarf, H.E. (1994) “The Allocation of Resources in the Presence of Indivisibilities,” *Journal of Economic Perspectives*, 8, 111-128.

[22] Schamalensee, R. (1989) “Inter-Industry Studies of Structure and Economic Performance,” *Handbook of Industrial Organization* (Schamalensee, R. and R. Willig eds.), 951-1009. New York: North Holland.

[23] Sutton, J. (1997) “Gilbrat’s Legacy,” *Journal of Economic Literature,* 35, 40-59.

[24] Williamson, O.E. (1985) *The Economic Institutions of Capitalism.* New York: Free Press.