Organization of Innovation in a Multi-Unit Firm: Coordinating Adaptive Search on Multiple Rugged Landscapes*

Myong-Hun Chang
Department of Economics
Cleveland State University
Cleveland, OH 44115
216-687-4523, -9206 (Fax)
m.chang@csuohio.edu
www.csuohio.edu/changm

Joseph E. Harrington, Jr.
Department of Economics
The Johns Hopkins University
Baltimore, MD 21218
410-516-7615, -7600 (Fax)
joe.harrington@jhu.edu
www.econ.jhu.edu/people/harrington

November 2000

Abstract

In Chang and Harrington (2000a), a computational model of a multi-unit firm is developed in which unit managers continually search for better practices. Search takes place over a rugged landscape defined over the space of unit practices. There it is shown that a more centralized organization is optimal when markets are not too different and the horizon is not too long. The robustness of those results are explored here with respect to the shape of the landscape. In particular, we find that centralization does better when the search space is larger and there is a stronger correlation in a consumer's preferences across different dimensions. A richer description of comparative dynamics is also provided.

*This paper was presented at the conference on Complex Behavior in Economics (Aix-en-Provence, France - May 2000), the 2000 Taipei International Conference on Industrial Economics (Taipei, Taiwan - June 2000), the International Conference on Computing in Economics and Finance (Barcelona, Spain - June 2000), and World Congress of the Econometric Society (Seattle - August 2000). We thank the participants for their comments. This research is supported by the National Science Foundation through Grant SES-0078752.

1 Introduction

A distinguishing feature of a multi-unit organization - such as a retail chain and a multi-unit manufacturer - is that it is comprised of units that are trying to solve similar but not identical problems. That these problems are distinct suggests it is desirable to give these units considerable discretion in how they try to solve these problems. That these problems are similar suggests that there may be ways in which units can pool their knowledge to jointly perform better. A difficulty is that these units are run by agents that tend to be motivated more by the prospect of enhancing the performance of their individual unit than the performance of the organization. As a result, giving them complete discretion may not result in the desired outcome. A major task for a multi-unit organization is then to balance giving units the freedom to handle the idiosyncratic features of their environments while, at the same time, coordinating various units' decision-making so as to take advantage of what is common to them. Chang and Harrington (2000a) took an initial step towards understanding the nature of this task by constructing a computational model of a retail chain and exploring the use of organizational structure - specifically, the allocation of authority - to solve this dilemma. The purpose was to identify the relevant features of the environment that affect the relationship between organizational structure and the long-term dynamics of adaptive learning by a multi-unit organization.

In the current paper, we extend that model by considering the impact of stores' environments on the optimal organizational structure. First, we assess the role of the difficulty of the problem faced by stores by changing the dimensionality of the space of stores practices. A higher dimensional search space is shown to enhance the performance of the centralized organization. Second, we consider the effect of consumer preferences in terms of the degree of correlation across the various dimensions. Decentralization performs better when there is less correlation. Third, we provide a more in-depth examination of the comparative dynamics of different organizational structures. Though tentative, our results suggest that there are three phases to learning by stores and they strongly determine how an organizational form performs over time.

2 Model

A retail chain is assumed to be made up of a corporate headquarters (HQ) and $M \geq 2$ stores. Each store is in a distinct market and has a set of N practices such that store i's operation in any given period is fully described by a vector, $z^i \equiv (z_1^i, z_2^i, ..., z_N^i)$, where z_j^i is store i's practice for the jth dimension of its operation and $z_j^i \in \{1, ..., R\}$, $j \in \{1, ..., N\}$ and $i \in \{1, ..., M\}$. Thus, there are R feasible practices for each dimension and, at any point in time, a store is represented by a point in $\{1, ..., R\}^N$.

2.1 Consumers

In each of these M markets, there are B consumers. Let $\{1, 2, \ldots, B\}$ denote the set of consumers in market i. All consumers in market i shop at store i. Each consumer has an ideal set of store practices which lies in $\{1, ..., R\}^N$ and is referred to as a consumer's type. Letting $w^{k,i} \equiv (w_1^{k,i}, w_2^{k,i}, ..., w_N^{k,i})$ be the type of consumer k in market i, it is assumed that such a consumer receives utility from buying x units at a price of p from store i equal to:

$$u(x; p, w^{k,i}, z^i) = \left[\overline{L} - \sqrt{\sum_{j=1}^{N}(z_j^i - w_j^{k,i})^2}\right]^{\gamma} \cdot x^{\beta} - p \cdot x \tag{1}$$

where $\beta \in (0,1), \gamma \geq 1$, and $\overline{L}$ is chosen so that $\left[\overline{L} - \sqrt{\sum_{j=1}^{N}(z_j^i - w_j^{k,i})^2}\right] > 1$ for all $(w^{k,i}, z^i) \in \{1, ..., R\}^{2N}$.[1] Given the above utility function, a consumer is assumed to optimize with respect to her purchase quantity, x.[2]

The distribution of consumer types is generated as follows. Each market has a seed which is denoted S_0^i for market i. For each consumer in a market, we randomly draw the consumer's seed from $\{S_0^i - G, ..., S_0^i + G\}$ according to a triangular density function. Letting s^k denote consumer k's seed, this consumer's type is drawn from $\left\{s^k - E, ..., s^k + E\right\}^N$ according to a uniform density function. E is a key parameter in the analysis as it controls the degree of correlation in consumers' preferences and will influence the shape of the landscape faced by stores. When $E = 0$, a consumer desires the same value for all dimensions (equal to its seed). Hence, preferences along one dimension are perfectly correlated with preferences along another dimension. As E increases, this correlation weakens. As a motivation for this structure, suppose preferences are determined by a consumer's income and some idiosyncratic elements (which may be determined by many factors, perhaps genetic or personal history, that are weakly related to income). Higher income may imply a whole host of preferences - a desire for a high quality-high price array of goods, a high level of attentive service (at the cost of sales people being pushier), and a low product variety-high inventory mix (which reduces the risk of having to return to the store and incur additional search costs). Higher income implies a desire for higher values along all of these dimensions. If income was the primary driver of consumer preferences then the correlation should be high and this is captured by a low value for E. As idiosyncratic elements enter and become more important, preferences are driven by a variety of factors and are apt to be more of a hodge-podge. This is captured by a high value for E.

[1] Letting $d \equiv \sqrt{\sum_{j=1}^{N}(z_j^i - w_j^{k,i})^2}$, it can be shown that $\overline{L} - d > 1$ is sufficient for $\frac{\partial^2 u}{\partial d \partial \gamma} < 0$ so that an increase in γ reflects a consumer's higher marginal dissatisfaction from actual store practices deviating from most preferred practices.

[2] It can be shown that the optimal purchase quantity for consumer k in market i is $x^*(p, w^{k,i}, z^i)$ $= \left(\frac{\beta}{p}\right)^{\frac{1}{1-\beta}} \left[\overline{L} - \sqrt{\sum_{j=1}^{N}(z_j^i - w_j^{k,i})^2}\right]^{\frac{\gamma}{1-\beta}}$.

The structure of our model is identical to that in Chang and Harrington (2000a) with the exception of the specification of consumer preferences. In that paper, $E = 0$, so that a consumer's preferences (or type) are represented by a scalar, and the distribution of consumers in a market is specified to be an exact (discrete) triangular density function. There is then no randomness in the distribution of consumer types. In the current model, there is randomness in the determination of the market seed and, given that market seed, in the distribution of consumer types.

2.2 Stores and the Chain

Given practices $z^i \equiv (z_1^i, ..., z_N^i)$ and the type distribution of consumers in market i, store i's profit is

$$(p-c)\left(\frac{\beta}{p}\right)^{\frac{1}{1-\beta}} \sum_{k=1}^{B} \left[\overline{L} - \sqrt{\sum_{j=1}^{N} (z_j^i - w_j^{k,i})^2} \right]^{\frac{\gamma}{1-\beta}} \tag{2}$$

Each store is assumed to set its price so as to maximize store profit. Given this profit function, the optimal price is fixed and independent of store practices: $p^* = \frac{c}{\beta}$. Taking into account the optimal price, a store's profit is proportional to:

$$\sum_{k=1}^{B} \left[\overline{L} - \sqrt{\sum_{j=1}^{N} (z_j^i - w_j^{k,i})^2} \right]^{\sigma} \tag{3}$$

where $\sigma \equiv \frac{\gamma}{1-\beta} > 1$. An important feature of this profit function is that it is decreasing in the distance between the store's practices and those desired by its customers.

The profit for the chain is a simple sum of stores' profits:

$$\sum_{i=1}^{M} \sum_{k=1}^{B} \left[\overline{L} - \sqrt{\sum_{j=1}^{N} (z_j^i - w_j^{k,i})^2} \right]^{\sigma} \tag{4}$$

2.3 Structure of the Landscape

Innovation involves searching over a landscape where a store's landscape is defined over the N-dimensional practice space with the height of the surface of the landscape being given by (3). This is analogously true for the chain though with the height of the surface of its landscape being given by (4). In this section, we characterize how a store's landscape depends on the correlation in consumers' preferences, as controlled by E, and the dimensionality of the space of practices, as controlled by N. This is to provide some background towards understanding later results. These simulations assume $R = 100, G = 25, B = 992$, and $\sigma = 3$.

Table 1 provides results for $N = 2$ and $N = 3$ when the market seed is 50. For $N = 2$ the number of local optima ranges from one to three on the basis of 100 replications of a store's landscape - each replication using a fresh set of consumer types drawn from $\{50 - E, ..., 50 + E\}^N$. Note that the number of local optima declines as E rises. The landscape was also examined for $N = 3$. On the basis of 20 replications, we found that the number of local optima increases in N in that the maximum number of local optima rose from 3 to 5 when N was raised from 2 to 3. As a caveat, our later results are based on much higher values for N and we do not know how general are the properties in Table 1.

2.4 Innovation

In each period, a store manager receives a new idea regarding store practices. Generally, an *idea* is a randomly selected element from $\{1, ..., R, *\}^N$ where $*$ denotes that an idea leaves a dimension unaffected. To be concrete, consider an example with $N = 5$. Idea $(76, *, 55, *, *)$ involves changing dimension 1 to 76 and dimension 3 to 55. If current store practices are $(14, 58, 34, 6, 99)$ then the new store practices are $(76, 58, 55, 6, 99)$ if this idea is adopted. In the simulations reported in this paper, we consider only one-dimensional ideas.[3]

Organizational structure impacts innovation by virtue of how it determines the allocation of authority within the chain regarding which new ideas are adopted. Generally defined, the degree of centralization is measured by the number of dimensions that HQ controls where "control of a given dimension" means having the authority to determine the adoption of those ideas that involve changes in that dimension. To save on computing time, we focus on two extreme forms of organization: *full centralization* where HQ controls all dimensions and *full decentralization* where store managers control all dimensions.[4]

In a given period, the process of innovation consists of the evaluation and potential adoption of internally generated ideas (stage 1) and externally generated ideas (stage 2). In a decentralized organization, a store manager comes up with an idea, evaluates it, and decides whether to adopt it in stage 1 (the criterion for adoption is described below). HQ observes all newly adopted ideas and automatically passes them along to the other stores.[5] In stage 2, these ideas are sequentially evaluated by each store manager. Using the same criterion for adoption, they decide whether to implement them in the current period. In a centralized organization, a store manager evaluates

[3] For $E = 0$, we have run the model when ideas can be high as N-dimensional and all qualitative results are robust.

[4] When ideas are multi-dimensional and both stores and HQ control some dimensions, some ideas may then be controlled by *both* stores and HQ. Such a possibility was considered in an earlier version of Chang and Harrington (2000a). It was assumed that the probability that HQ controls an idea is increasing in the proportion of dimensions that it controls which that idea impacts.

[5] In Chang and Harrington (2000a), we consider a class of decision rules that determine which newly implemented practices HQ decides to communicate throughout the chain.

his internally generated idea and decides whether or not to pass it to HQ in stage 1. In stage 2, HQ sequentially evaluates those ideas that were sent up and decides whether or not to mandate them across the chain. A key implicit assumption is that HQ has less information than store managers about the local market environment and this prevents them from selectively implementing new ideas in particular markets. Rather, they either mandate an idea across the chain or discard it.[6]

In deciding what to do with an idea, a store uses the following evaluation criterion. A store manager adopts an idea (under decentralization) or communicates it to HQ (under centralization) if and only if (iff) it raises store profit. Otherwise, the idea is discarded; not being worth the store manager's time to consider further. Thus, if store i's current practices are $(z_1', ..., z_N')$ and if its practices from adopting a new idea are $(z_1'', ..., z_N'')$, this idea is adopted iff:

$$\sum_{k=1}^{B} \left[\overline{L} - \sqrt{\sum_{j=1}^{N} (z_j'' - w_j^{k,i})^2} \right]^{\sigma} > \sum_{k=1}^{B} \left[\overline{L} - \sqrt{\sum_{j=1}^{N} (z_j' - w_j^{k,i})^2} \right]^{\sigma} \quad (5)$$

Under centralization, HQ mandates an idea across the chain iff doing so raises chain profit and discards it otherwise. To avoid an additional source of randomness in the model, the evaluations by store managers and HQ are assumed to be perfect; that is, each gets a perfect signal of what profit would be if an idea was adopted.

3 Results

In our model, innovation is modelled as random search carried out in a finite fixed space of ideas. This particular approach is rooted in the concept of a fitness landscape, defined in a multidimensional space in which each attribute of a store is represented by a dimension of the space and a final dimension indicating the performance (profit) of the store. An adaptation of a store is then represented by movement on the landscape toward a location reflecting higher fitness value. In the context of population genetics, Kauffman (1993) demonstrated that the topography of the fitness landscape is determined by the degree of interdependence of the fitness contribution of the various attributes of an organism. Taking the Darwinian perspective from organizational ecology, Levinthal (1997) uses this connection in the context of organizational attributes to examine the effectiveness of organizational adaptation at the population level.[7] As shown in Section 2.3, the stores in our model also face rugged landscapes with their ruggedness influenced by the degree of correlation in consumer preferences,

[6] In a sense, we are limiting the centralized chain to a very coarse set of actions. In spite of that, centralization can outperform. Presumably, it could do even better with a more refined set of actions.

[7] For elaborations of a similar idea in other contexts, see Carley and Svoboda (1996), Carley and Lee (1998), Kauffman, Lobo, and Macready (1998), Page and Ryall (1998), Auerswald, Kauffman, Lobo, and Shell (1999), and Kollman, Miller, and Page (2000).

as controlled by E, as well as the difficulty in adaptive search, as controlled by N. Given that the ruggedness of the landscape is likely to influence the effectiveness of adaptive search, we proceed to examine how parameters that determine the shape of the landscape, E and N, influence the properties identified in Chang and Harrington (2000a).

3.1 Simulation Design[8]

For each set of parameter values, the computational experiment consists of 400 replications. Each replication involves a randomly drawn vector of initial store practices (which are assumed to be identical for all stores), a vector of ideal practices for each consumer in each market, and a sequence of TM ideas, one for each of the M stores in each of T periods. All reported profit numbers are averaged across stores in the chain and across the 400 replications.

For the simulations reported below, we assume $R = 100$, $G = 25$, $B = 992$, $\sigma = 10$, $M = 3$, and $T \in \{500, 1000\}$. In order to examine the impact of inter-market heterogeneity, market seeds are specified to be symmetric around 50: $(S_0^1, S_0^2, S_0^3) = (50 - \alpha, 50, 50 + \alpha), \alpha \in \{0, 1, \ldots, 5\}$. Thus, higher values for α correspond to greater inter-market heterogeneity. Given the seed for market i, we assign a seed to consumer k, denoted s^k, for each of 992 consumers, where s^k is drawn from $\{S_0^i - 25, ..., S_0^i + 25\}$ according to a triangular density function. Given s^k, the type of consumer k is completely described by her ideal practices in each of the N dimensions of a store's operations. Given the correlation parameter E, each consumer k's type is drawn from $\left\{s^k - E, \ldots, s^k + E\right\}^N$ according to a uniform density function. Other than the difference noted in Section 2.1, the model and simulations in Chang and Harrington (2000a) is the special case of $E = 0$ and $N = 10$.

3.2 Benchmark Result

We begin with the benchmark case of $E = 0$ and $N = 10$ which serves to replicate the results in Chang and Harrington (2000a). Define "$v_C - v_D$ in period t" to be the average profit under centralization in period t minus the average profit under decentralization in period t.[9] Figure 1 shows the time series on this profit differential over 1000 periods and for various degrees of inter-market heterogeneity (recall that markets are, in expectation, identical when $\alpha = 0$ and the difference between markets is increasing in α). This differential tends to be positive in the earlier periods when α is low (though not always in the very earliest periods which is a point to which we

[8] The simulation codes were written in C^{++} and are available upon request from Myong-Hun Chang (m.chang@csuohio.edu). The pseudo-code is in Appendix A.

[9] The averaging is done across replications. Also note that in calculating the profit of both organizational forms for a given replication, we used the same initial practices and the same sequence of ideas and thereby controlled for two sources of randomness.

will return in Section 3.4) which reveals that centralization is yielding higher profit in those periods. Eventually the differential declines and, when α is sufficiently high, becomes negative; indicating that decentralization is out-performing in the long-run.

For various values of E and N, Table 2 reports the organizational form that generates higher average profit and shows its dependence on α and T.[10] Generally, centralization tends to outperform more when the horizon is shorter and when inter-market heterogeneity is lower. This result is robust to the values of E and N though is not as clean for higher values of N. For example, when $N = 30$, the decentralized organization often outperforms when markets are identical ($\alpha = 0$) though this is never statistically significant. The associated time series for the profit differential when $E > 0$ and $N > 10$ will be examined later but they also reflect the same qualitative properties as those in Figure 1. These findings establish the robustness of the main results in Chang and Harrington (2000a) with respect to the dimensionality of the search space and the structure of consumer preferences. This is re-stated here as Property 1.

Property 1: Centralization outperforms when markets are sufficiently similar and the horizon is sufficiently short.

Given that markets are heterogeneous, the benefit of decentralization is clear - it allows each store manager to tailor practices to the local market. How can then a centralized structure outperform? It was argued in Chang and Harrington (2000a) that there is an implicit cost to decentralization. As stores tailor their practices to their markets in a decentralized chain, these practices drift farther apart. As a result, a new practice adopted by one store is increasingly likely to be incompatible with the current practices of other stores. In essence, stores come to target distinct consumer types (that is, local optima) and what works for one type of consumer is unlikely to work for another type of consumer. Inter-store learning is then less under decentralization and this is detrimental to the rate of improvement in store practices. The virtue of a centralized structure is that it enhances inter-store learning by keeping stores close in store practice space so that they are targeting similar consumers. With these two countervailing effects, a centralized structure outperforms as long as markets are not too different.

Returning to Figure 1, note that centralization is superior earlier in the horizon when learning is most active. During this time, a centralized chain more effectively progresses toward some optimum or set of optima. In Chang and Harrington (2000a), it is shown that mutual learning - by which is meant that one store finds another store's adopted idea profitable - is indeed higher under centralization. That superiority, however, dissipates over time as stores in the decentralized chain individually identify desirable local optima and begin to independently converge to them. While

[10] A description of how the test statistic is constructed is provided in Appendix B.

mutual learning is less under decentralization, the ultimate superiority of its global optimum tends to result in the decentralized form outperforming in the long run.[11]

3.3 Altering the Store's Landscape

For when the market seeds are $(S_0^1, S_0^2, S_0^3) = (49, 50, 51)$, we simulated 400 replications of the innovation process for $N \in \{10, 20, 30, 40, 50\}$ and $E \in \{0, 2, 4, 6, 8, 10\}$. For $N = 10$, Figure 2 shows the time series on the profit differential for various values for E. As E is raised, the series shifts downward and the time over which it is positive becomes shorter. For example, moving from $E = 2$ to $E = 6$, the peak of the curve falls from around 3.8 ($\times 10^{28}$) to around 3.0 and the period at which it goes from being positive to being negative falls from around 500 to around 250. The general tendency (which is confirmed by Table 2) is that less correlation in a consumer's preferences enhances the relative performance of the decentralized structure.

Property 2: An increase in E favors decentralization.

As reported in Table 1, an increase in E results in fewer local optima and thereby a smoother landscape. With fewer optima, it becomes more likely that decentralized stores will target the same optimum and, if that is the case, it is more likely that the new practices implemented by one store are of value to other stores. Thus, a smoother landscape enhances mutual learning under decentralization and thus lessens the advantage of centralization.

For $E = 0$, Figure 3 shows the time series on the profit differential for various levels of dimensionality of the space of store practices. As N is raised, the series shifts downward in the earlier periods and upward in the later periods. For example, when $N = 10$, the profit differential peaks at a value around 4.5 ($\times 10^{28}$) around period 50. It is not too far below that value as late as period 200 though it has fallen by about 75% by period 600. When instead $N = 20$, the series peaks at a lower value, around 3.0, but much later, around period 700. Though increasing the dimensionality may then lower the profit differential for some periods, the relative performance of centralization over 1000 periods is clearly improved and this is most easily seen in Table 2.

Property 3: An increase in N favors centralization.

As N increases, the operations of a store become more complex in the sense of involving a greater number of dimensions. This makes search more difficult as the space is larger and that results in the average time to reach an optimum rising. By Table 1, it is also true that there are more local optima. Combined with the

[11] In Chang and Harrington (2000a), we also explored the stationary profit path when the distribution of consumer preferences in a market follows a random walk. In that case, the centralized form can outperform in the long-run as well.

argument provided above, this suggests that centralization becomes more attractive as the dimensionality of the space rises because, under decentralization, it is less likely that stores will target the same optimum and thus there is less mutual learning. That Property 3 may be due to the longer time it takes before getting close to an optimum is consistent with the result that centralization does better in shorter horizons because, in both cases, we are observing that centralization does better when stores are at more sub-optimal practices. All this suggests that mutual learning, which is greater under centralization, is more important when practices are farther away from a local optimum.

The logic just described is that a higher dimensional search space slows down the learning process and thus lengthens the time over which centralization outperforms. But there may be more going on for consider Figure 4 which plots the differential profit path for various values of N. Generally, when N is higher, the path is lower in the earlier periods and higher in the later periods, as we've already noted. Also note, however, that for $\alpha \geq 2$, the maximum value for the profit differential (where the maximum is taken over time) tends to be higher when there are more dimensions to the search space. For example, it is higher for $N = 50$ than for all lower values of N. Secondly, notice that the differential is positive for some time periods when N is sufficiently high while, for lower values of N, it is never positive. For example, when $\alpha = 3$, centralization outperforms more than 60% of the time (that is, for more than 600 of 1000 periods) when N is 40 or 50 but only does so for less than 10% of the time when N is 10 or 20. While this could just be a complex implication of a higher dimensional search space slowing down the search process, it is possible that something else is driving these properties.

3.4 Comparative Dynamics

An examination of the time series for the profit differential reveals three phases. These phases tend to be more magnified for higher values of N though they appear to generally be present. In Figure 5, the profit differential initially falls and goes negative which indicates that the decentralized structure is outperforming and is increasingly doing better. This is the first phase. The second phase is when the differential begins to rise and, depending on the value of α, becomes positive. But, regardless of how diverse markets are, centralization is performing increasingly better during this phase. The differential then peaks and enters the third phase during which it steadily falls and either becomes negative or appears to is trending to be negative.

In trying to understand this pattern, we have a working hypothesis as to how learning is occurring over time and how it varies between the type of organizational form. This is rather speculative at the moment and we are still trying to develop measures which could assess its validity. After offering this explanation, we will present what statistics we do have.

- Phase 1: In that stores' initial practices are randomly selected, it is safe to presume that they are highly sub-optimal. It is then quite possible that a store can effectively learn on its own because it is not difficult to find new ideas which are an improvement on existing practices. Thus, inter-store learning may not be that important in the early periods. Given that a higher rate of inter-store learning is the advantage of centralization, this suggests that the benefit of centralization is relatively weak. On the other hand, the cost of centralization - imposing unprofitable practices on some stores - is still present. These forces could explain why the profit differential is initially decreasing and negative.

- Phase 2: As stores get out of having highly sub-optimal practices, they are likely to move into the basin of attraction for some optimum or a small set of optima. At this point, finding useful ideas becomes more difficult under either organizational form. Perhaps it is here that inter-store learning starts becoming important. It may now be incrementally more valuable to be able to learn from other stores which is easier to occur under centralization since, by fiat, stores must be in the basin of the same optimum. If so, this could explain why the profit differential becomes increasing and positive.

- Phase 3: Stores now begin to hit the limit of centralization because it constrains them to have identical practices even though their markets are different. In contrast, under decentralization, stores can approach their global optimum. As a result, the profit differential becomes decreasing and may become negative.

Towards trying to assess the validity of this description, we have measured the relative learning rates across organizational structures. A store's learning rate is defined to be the proportion of ideas available to a store that are actually adopted. Let $n(O, i, t, h)$ denote the number of ideas adopted by store i in period t in replication h when the organizational structure is $O \in \{C, D\}$. Since there are M stores and each store generates one idea per period, the store's learning rate is measured by $n(O, i, t, h)/M$. We can then construct the per-store learning rate at t averaged over the X replications:

$$k(O, t) \equiv \frac{1}{X} \sum_{h=1}^{X} \left[\frac{1}{M} \sum_{i=1}^{M} \left(\frac{n(O, i, t, h)}{M} \right) \right].$$

The relative learning rate is defined to be $k(C, t)/k(D, t)$.

Figure 6 reports the 20-period moving average for the relative learning rate.[12] Initially, it is above one and rising which indicates that more ideas are being adopted under centralization and this advantage is increasing. Eventually, however, it begins to decline and ends up below one (at least for $\alpha = 5$); revealing that more ideas are

[12] In that the relative learning rate is found to be very noisy, it has proven to be more informative to examine a moving average (which is still quite noisy).

being adopted by the decentralized chain. This fits in with our description of phases 2 and 3. During phase 2, the higher rate of mutual learning when stores have identical locations results in more ideas being adopted. During phase 3, the constraint imposed under centralization becomes increasingly restrictive while, under decentralization, stores are unconstrained to adopt any idea that is profitable to them. This statistic does not provide support for our explanation of phase 1, however. Indeed, evidence shedding light on it remains elusive.

Also reported is the average return to learning which, in any period, is defined as the change in chain profit from the previous period divided by the number of adopted ideas by all stores. If no ideas are adopted then it is specified to be zero. The reported measures are averaged across the replications. The relative average return to learning is the average return to learning for the centralized structure divided by the same measured for the decentralized structure. Note that with the decentralized chain, an idea is adopted iff it raises a store's profit. Under centralization, an idea may be adopted even if it lowers a store's profit though it must be the case that it raises the profit of some other stores (so that total chain profit rises). We then expect the average return to learning to be higher for the decentralized chain.

Not surprisingly, we find that the relative average return to learning is typically below one which indicates that, on average, adopted ideas are more profitable under decentralization. What is more informative is that the relative rate is declining over time. The ideas that are adopted under decentralization are increasingly more profitable relative to centralization. Offsetting this is that the relative learning rate is higher under centralization due to the greater similarity in stores' practices. Though we do not have a tight connection yet, we believe it is the interaction of these two measures that is determining the relative performance of different organizational forms.

4 Concluding Remarks

This paper has explored how basic parameters in the market environment - the structure of consumer preferences and the dimensionality of the space of practices - alter the shape of the landscape and thereby impact the relative performance of centralized and decentralized multi-unit organizations. Consistent with our earlier work, decentralization is not a panacea. Imposing uniform practices from on high can lead to superior performance even though individual agents face heterogeneous environments. The relevant question is then: when does decentralization outperform? The results of this paper show that decentralization tends to be superior in markets for which the practice space is simple (that is, there are few dimensions to what a store does) and there are weak correlations in consumer preferences.

As part of this line of research, we are currently adapting this model to allow for competing chains and consumer search and thereby be able to investigate the interaction of market structure and organizational structure (see Chang and Harrington, 2000b). The setting is now one in which each market is served by several chains. Consumers engage in search to find the store that best fits their needs while chains compete by discovering and adopting better practices. The spillover of ideas is not just within a chain but also across chains; in particular, we will allow a store to learn about the new practices of competing stores in their market. Initial results suggest that the interaction of stores learning and consumers searching generate rich dynamics and new insight into the role of organizational structure. In particular, centralization does even better in the presence of competition and consumer sorting.

5 Appendix A

```
// Pseudo Code
main ( )
{
Assign consumer types in all markets;
Initialize store practices init_prac(i) in t = 0 ∀i ∈ {1, 2, ..., M};
Create innovation list I(i, t) ∀i ∈ {1, 2, ..., M}, ∀t ∈ {1, 2, ..., T};
Set organizational form, O_i ∈ {D, C};
Set t = 0;
Initialize store profits for all i ∈ {1, 2, ..., M};
Initialize chain profit;
Set t = t + 1;
While (t < T) {
    Switch (O_i) {
    Case D: PROC I: Adoption Process for a Decentralized Chain
    Case C: PROC II: Adoption Process for a Centralized Chain
    }
    Report chain profit for period t;
    t++;
}
Return 0;
} // End Main
```

PROC I: Adoption Process for a Decentralized Chain

```
For each store i ∈ {1, 2, ..., M} {
    Evaluate the profitability of I(i, t);
    If the idea is profitable for store i then do {
```

```
            Store i adopts the idea;
            Store i's profit is revised;
            I(i,t) enters the HQ InBox;
        }
}
// Consideration of the ideas in the HQ InBox
For each idea l in the HQ InBox {
        Identify the source store of the idea, s[l] ∈ {1,2,...,M};
        Evaluate the profitability of the idea for store k for each k ∈ {1,2,...,M}, k ≠
            s[l] {
            If the idea is profitable for store k, then do {
                  Store k adopts the idea;
                  Store k's profit is revised;
            }
        }
}
Revise chain profit;
// End PROC I
```

PROC II: Adoption Process for a Centralized Chain

```
For each store i ∈ {1,2,...,M} {
        Evaluate the profitability of I(i,t) for store i;
        If the idea is profitable for store i, then do {
            If the chain profit improves as the result of mandating I(i,t) on all stores,
                then do {
                  Intra-chain mandate: All stores adopt I(i,t);
                  Revise chain i's profit;
            }
        }
}
// End PROC II
```

6 Appendix B

The test statistic used in Table 2 is constructed as follows. Let $v_C^{t,h}$ denote the profit of a centralized chain in period t for replication h. Similarly define $v_D^{t,h}$ for when the chain is instead decentralized. Figure 1 reports $\left(\frac{1}{X}\right)\sum_{i=1}^{X}\left[v_C^{t,h} - v_D^{t,h}\right]$ where X is the number of replications. Next define $V_C^h(T) \equiv \sum_{t=1}^{T}\left(\frac{1}{T}\right)v_C^{t,h}$ and $V_D^h(T) \equiv \sum_{t=1}^{T}\left(\frac{1}{T}\right)v_D^{t,h}$ as average chain profit over the first T periods for a centralized and decentralized chain, respectively. Defining $\delta^h(T) \equiv V_C^h(O) - V_D^h(O)$, we can construct the following test statistic:

$$Z = \frac{\overline{\delta}(T)}{\frac{\sqrt{\left(\frac{1}{X}\right)\sum_{h=1}^{X}\left(\delta^h(T)\right)^2 - \left(\overline{\delta}(T)\right)^2}}{\sqrt{X}}}$$

where

$$\overline{\delta}(T) \equiv \left(\frac{1}{X}\right)\sum_{h=1}^{X} V_C^h(T) - \left(\frac{1}{X}\right)\sum_{h=1}^{X} V_D^h(T).$$

References

[1] Auerswald, P., S. Kauffman, J. Lobo, and K. Shell, "The Production Recipes Approach to Modeling Technological Innovation: An Application to Learning By Doing," *Journal of Economic Dynamics and Control*, 24 (1999), 389-450.

[2] Carley, K. M. and D. M. Svoboda (1996), "Modeling Organizational Adaptation as a Simulated Annealing Process," *Sociological Methods & Research*, 25 (1996), 138-168.

[3] Carley, K. M. and J.-S. Lee, "Dynamic Organizations: Organizational Adaptation in a Changing Environment," *Advances in Strategic Management*, 15 (1998), 269-297.

[4] Chang, M.-H. and J. E. Harrington, Jr., "Centralization vs. Decentralization in a Multi-Unit Organization: A Computational Model of a Retail Chain as a Multi-Agent Adaptive System," Santa Fe Institute, Working Paper 00-02-10, 2000a (*Management Science*, forthcoming). [www.santafe.edu/sfi/publications/00wplist.html]

[5] Chang, M.-H. and J. E. Harrington, Jr., "Competition and the Organizational Structure of Multi-Unit Firms," June 2000, Revised September 2000b. [www.econ.jhu.edu/People/Harrington/cyborg9-00.pdf]

[6] Kauffman, S., *The Origins of Order*, New York: Oxford University Press, 1993.

[7] Kauffman, S., J. Lobo, and W. G. Macready, "Optimal Search on a Technology Landscape," Santa Fe Institute, Working Paper 98-10-091, 1998.

[8] Kollman, K., J. H. Miller, and S. E. Page, "Decentralization and the Search for Policy Solutions," *Journal of Law, Economics, and Organization*, 16 (2000), 102-128.

[9] Levinthal, D. A., "Adaptation on Rugged Landscapes," *Management Science*, 43 (1997), 934-950.

[10] Page, S. E. and M. Ryall, "Does Strategy Need Computer Experimentation?," *Advances in Strategic Management*, 15 (1998), 299-326.

Frequency Distribution with $N = 2$ (100 replications)

No. Local Optima	$E = 0$	$E = 1$	$E = 2$	$E = 4$	$E = 6$	$E = 10$	$E = 20$
1	0	4	31	79	90	96	100
2	0	18	53	21	10	4	0
3	100	78	16	0	0	0	0
No. Replications	100	100	100	100	100	100	100

Frequency Distribution with $N = 3$ (20 replications)

No. Local Optima	$E = 0$	$E = 1$	$E = 2$	$E = 4$	$E = 6$	$E = 8$
1	0	0	5	16	18	19
2	0	1	6	4	2	1
3	0	19	9	0	0	0
4	0	0	0	0	0	0
5	20	0	0	0	0	0
No. Replications	20	20	20	20	20	20

Table 1: Landscape Ruggedness As a Function of E

	$E=0$		$E=2$		$E=4$		$E=6$	
T	500	1000	500	1000	500	1000	500	1000
$N=10$								
$\alpha=0$	$\mathbf{C}$	$\mathbf{C}^*$	$\mathbf{C}^*$	$\mathbf{C}^*$	$\mathbf{C}^*$	$\mathbf{C}^*$	$\mathbf{C}^*$	$\mathbf{C}^*$
$\alpha=1$	$\mathbf{C}^*$	$\mathbf{C}^*$	$\mathbf{C}^*$	$\mathbf{C}$	$\mathbf{C}^*$	$\mathbb{D}$	$\mathbf{C}$	$\mathbb{D}^*$
$\alpha=2$	$\mathbb{D}$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$
$\alpha=3$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$
$\alpha=4$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$
$\alpha=5$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$
$N=30$								
$\alpha=0$	$\mathbb{D}$	$\mathbf{C}$	$\mathbb{D}$	$\mathbb{D}$	$\mathbb{D}$	$\mathbf{C}$	$\mathbb{D}$	$\mathbf{C}$
$\alpha=1$	$\mathbf{C}^*$	$\mathbf{C}^*$	$\mathbf{C}^*$	$\mathbf{C}^*$	$\mathbf{C}$	$\mathbf{C}$	$\mathbf{C}$	$\mathbb{D}$
$\alpha=2$	$\mathbf{C}^*$	$\mathbf{C}^*$	$\mathbf{C}^*$	$\mathbb{D}$	$\mathbb{D}$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$
$\alpha=3$	$\mathbb{D}$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$
$\alpha=4$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$
$\alpha=5$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$
$N=50$								
$\alpha=0$	$\mathbf{C}$	$\mathbf{C}^*$	$\mathbf{C}$	$\mathbf{C}$	$\mathbb{D}$	$\mathbb{D}$	$\mathbf{C}$	$\mathbb{D}$
$\alpha=1$	$\mathbf{C}$	$\mathbf{C}^*$	$\mathbf{C}$	$\mathbf{C}^*$	$\mathbb{D}$	$\mathbf{C}$	$\mathbf{C}$	$\mathbf{C}^*$
$\alpha=2$	$\mathbf{C}^*$	$\mathbf{C}^*$	$\mathbf{C}^*$	$\mathbf{C}^*$	$\mathbf{C}$	$\mathbf{C}$	$\mathbb{D}$	$\mathbb{D}^*$
$\alpha=3$	$\mathbf{C}^*$	$\mathbf{C}^*$	$\mathbf{C}^*$	$\mathbf{C}^*$	$\mathbb{D}$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$
$\alpha=4$	$\mathbf{C}$	$\mathbb{D}$	$\mathbb{D}$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$
$\alpha=5$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$	$\mathbb{D}^*$

Table 2: ***Ex Ante*** **Optimal Organizational Form**

[Note: Those results with "*" are statistically significant at the 5% level.]

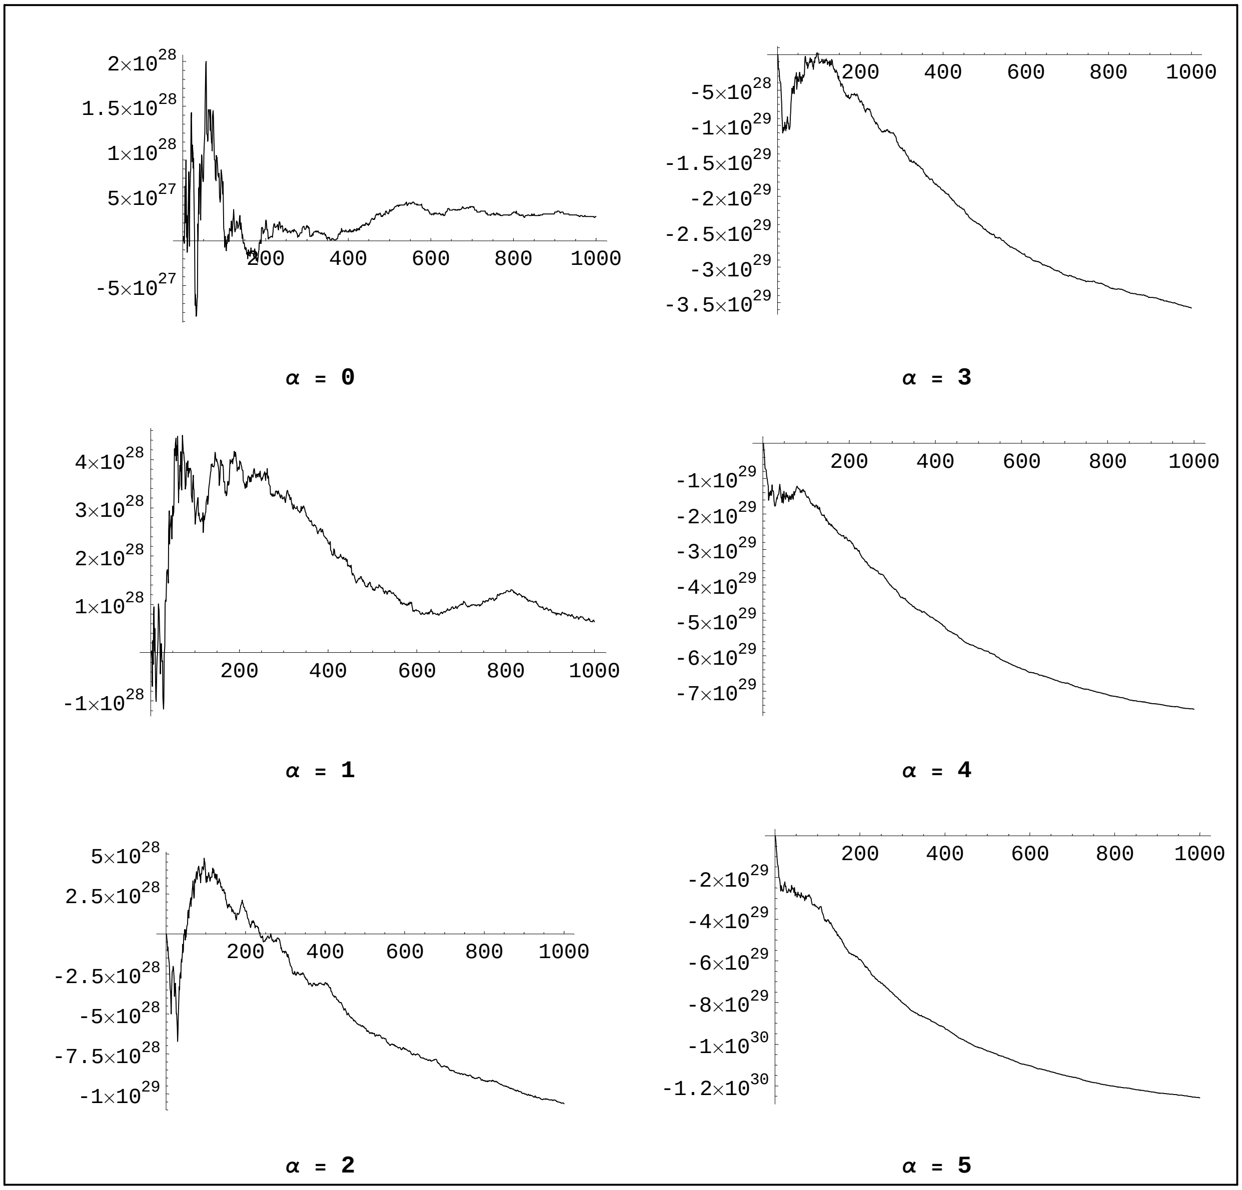

Figure 1: $(V_C - V_D)$ time series for $E = 0$ and $N = 10$

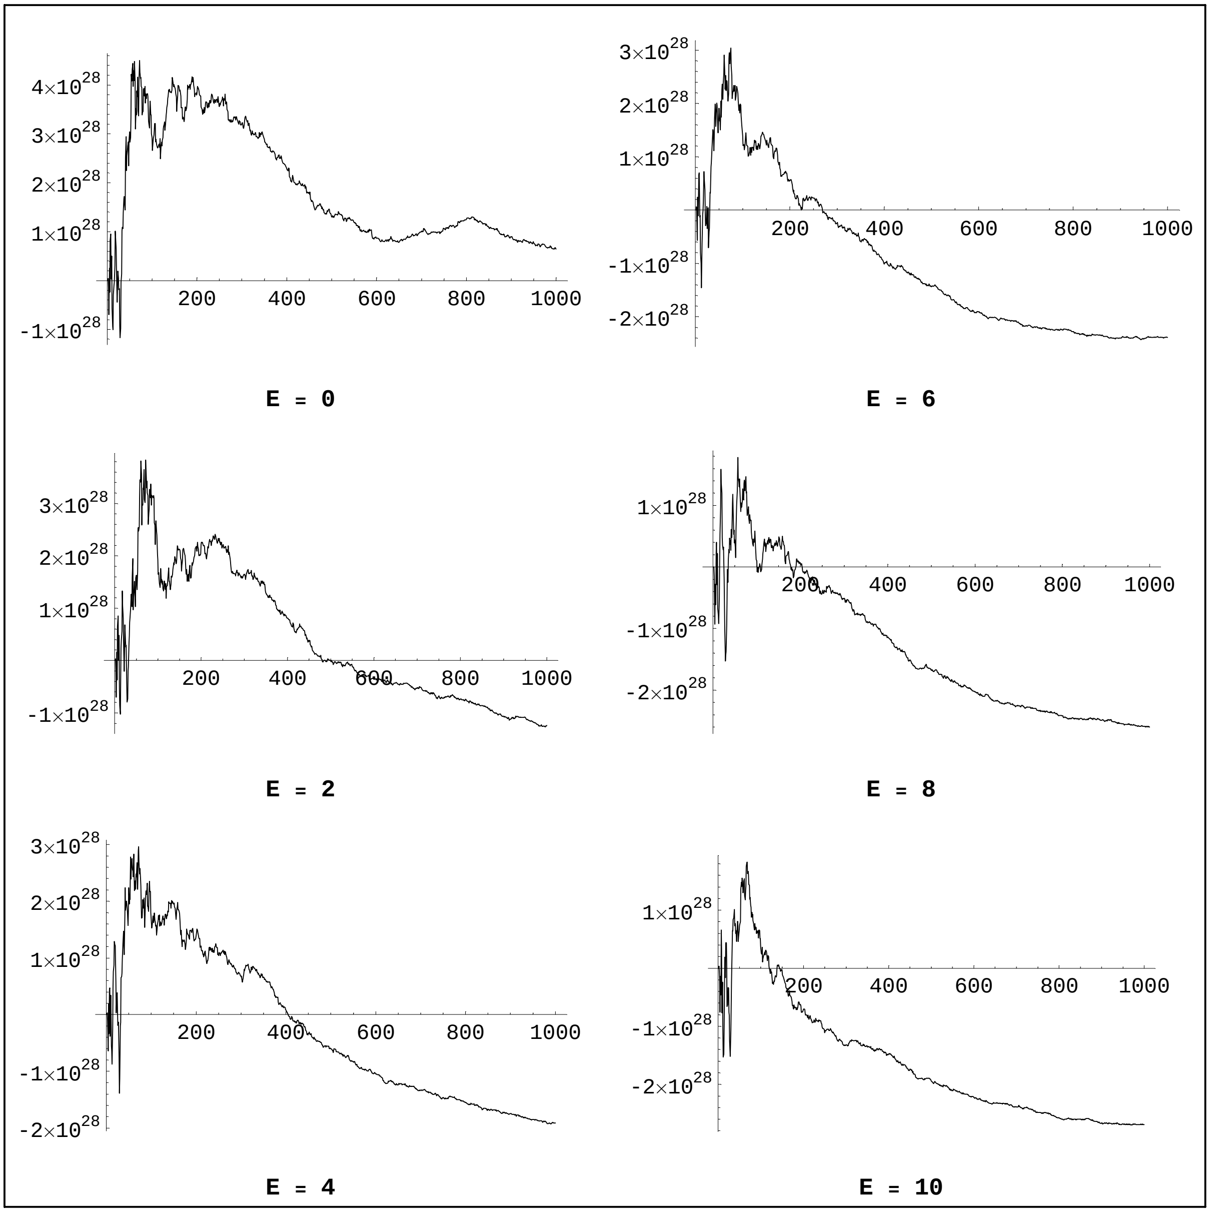

Figure 2: $(V_C - V_D)$ time series for α = 1 and N = 10

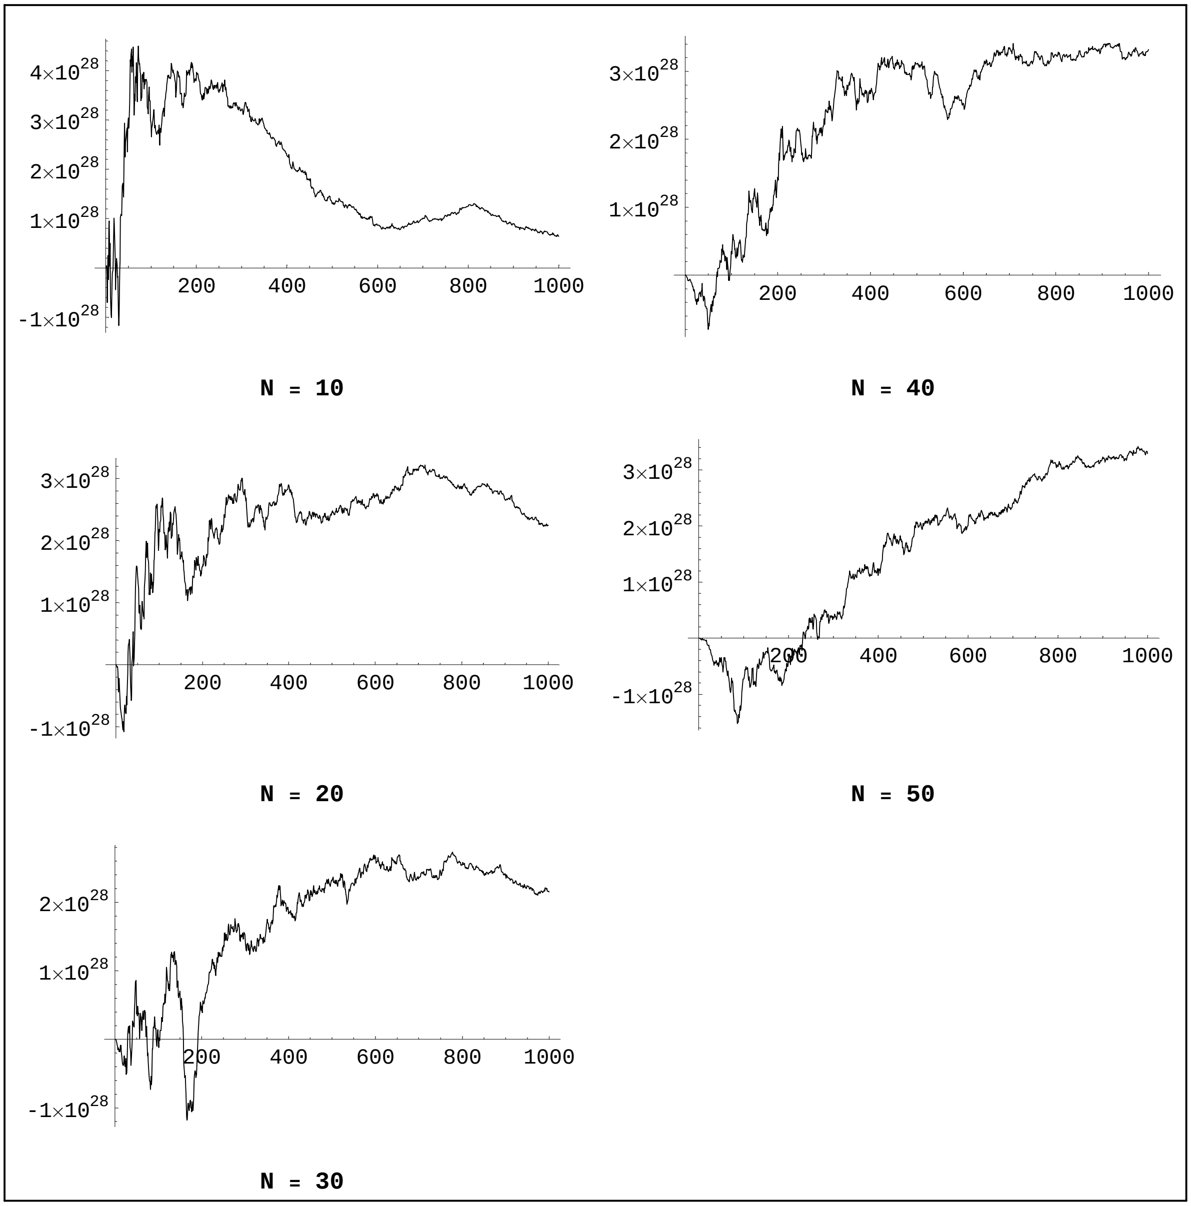

Figure 3: ($V_C - V_D$) time series for α = 1 and E = 0

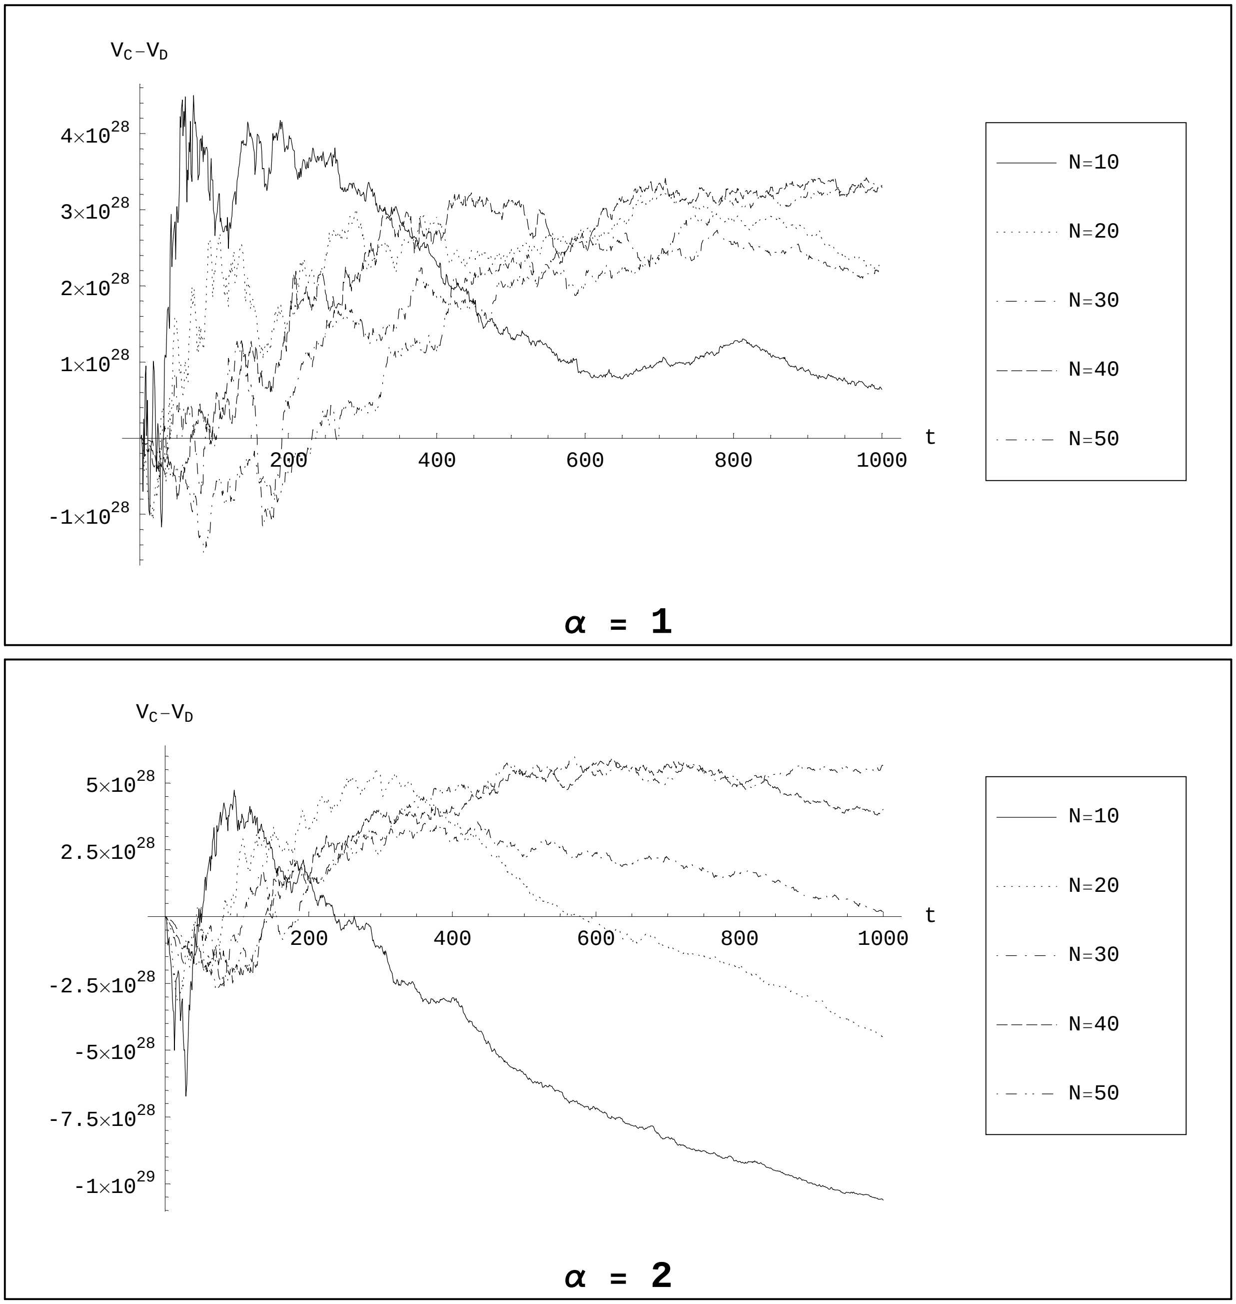

Figure 4-A: $V_C - V_D$ for various N $[E = 0]$

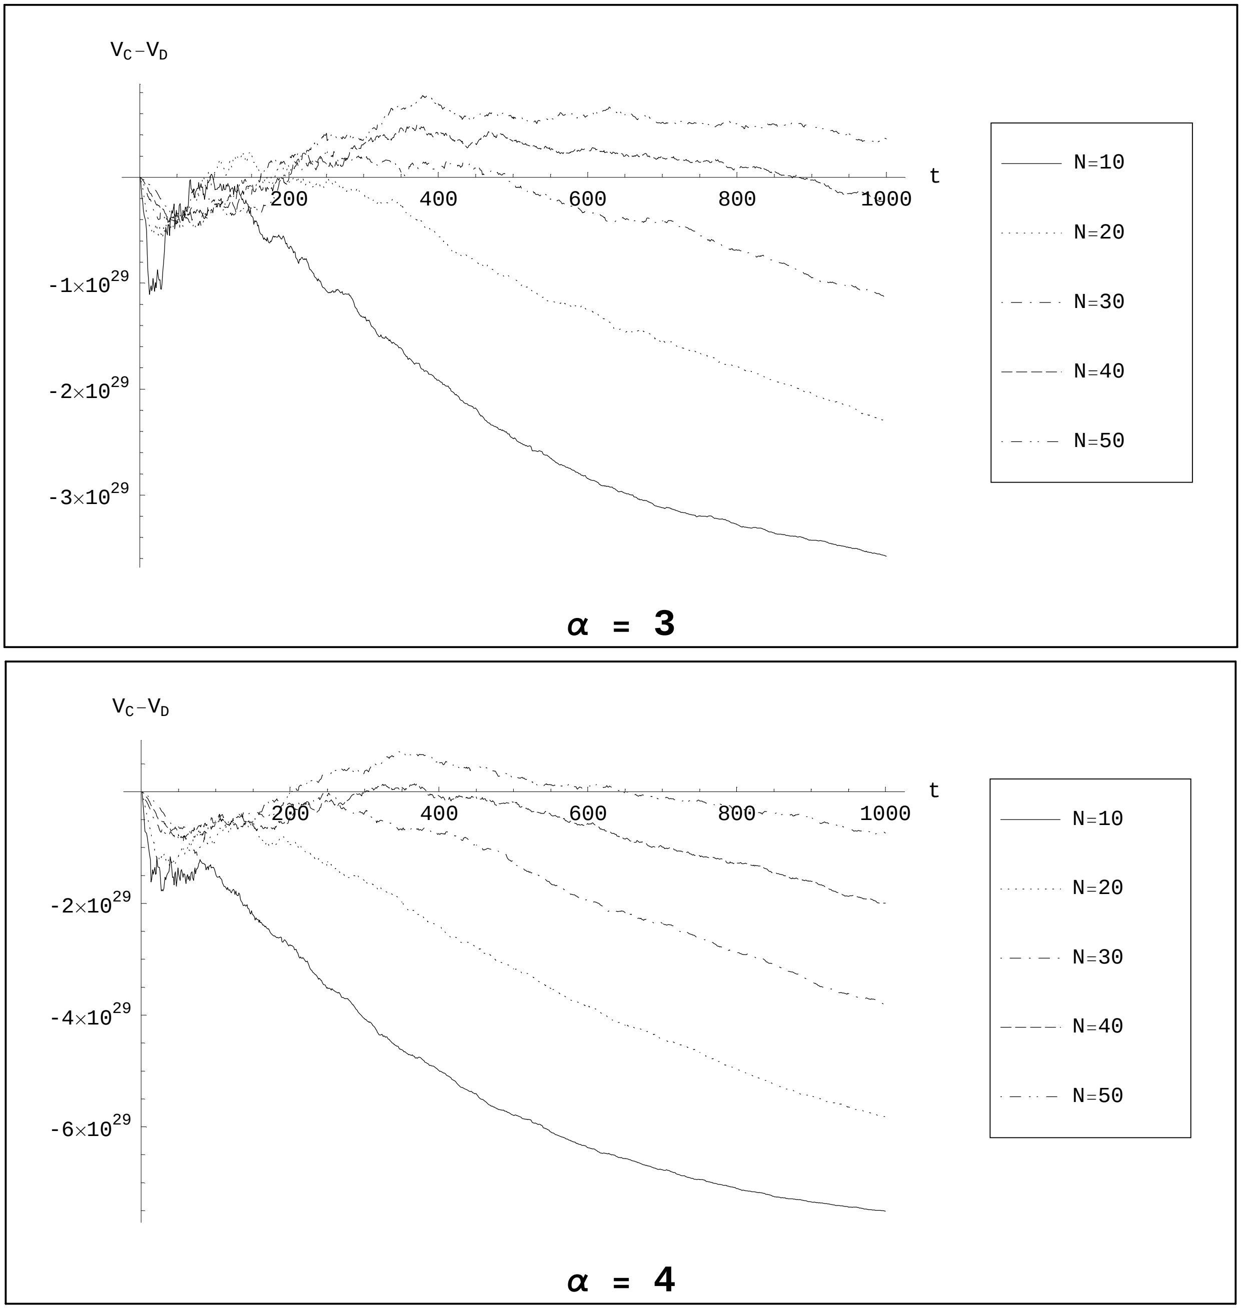

Figure 4-B: $V_C - V_D$ for various N [$E = 0$]

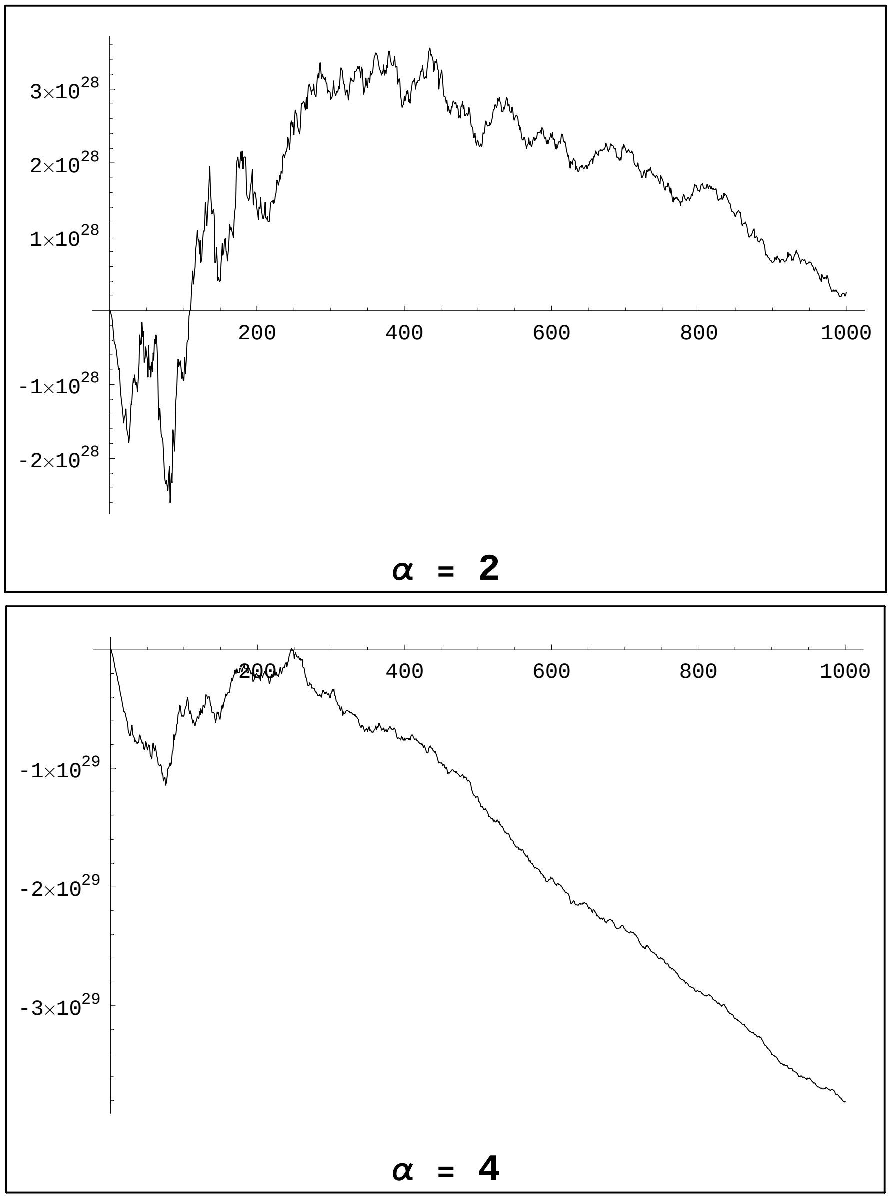

Figure 5-A: $V_C - V_D$ for $N = 30$ $[E = 0]$

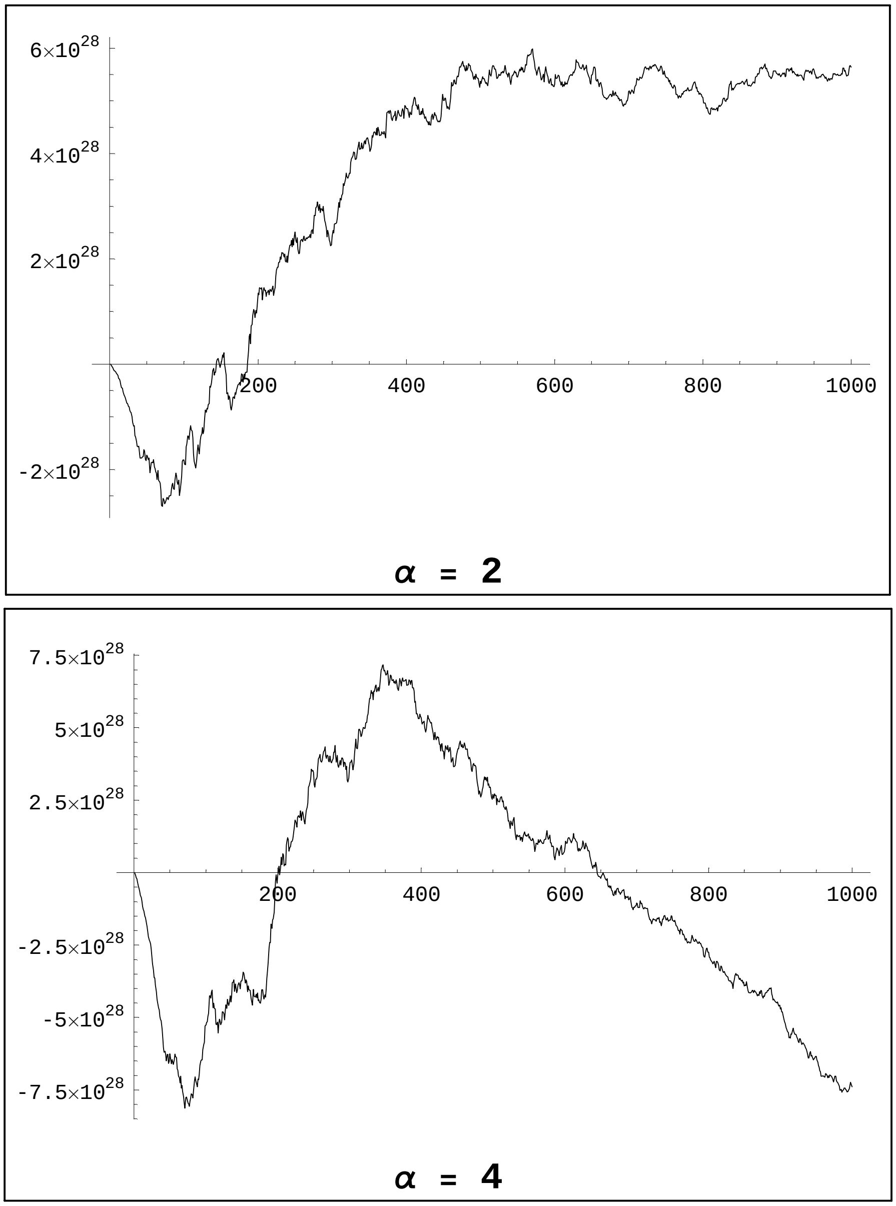

Figure 5-B: $V_C - V_D$ for $N = 50\,[E = 0]$

20 – Term Moving Average for α = 2

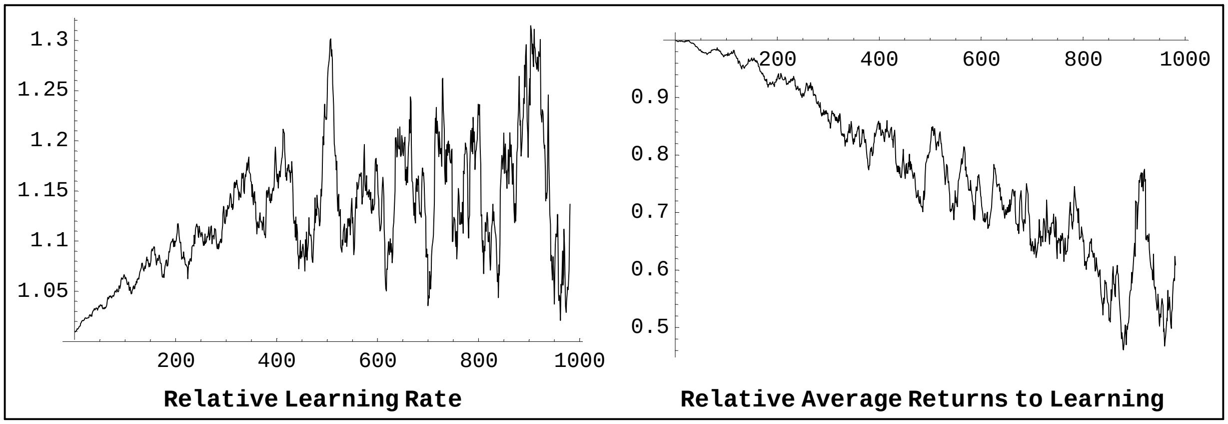

20 – Term Moving Average for α = 5

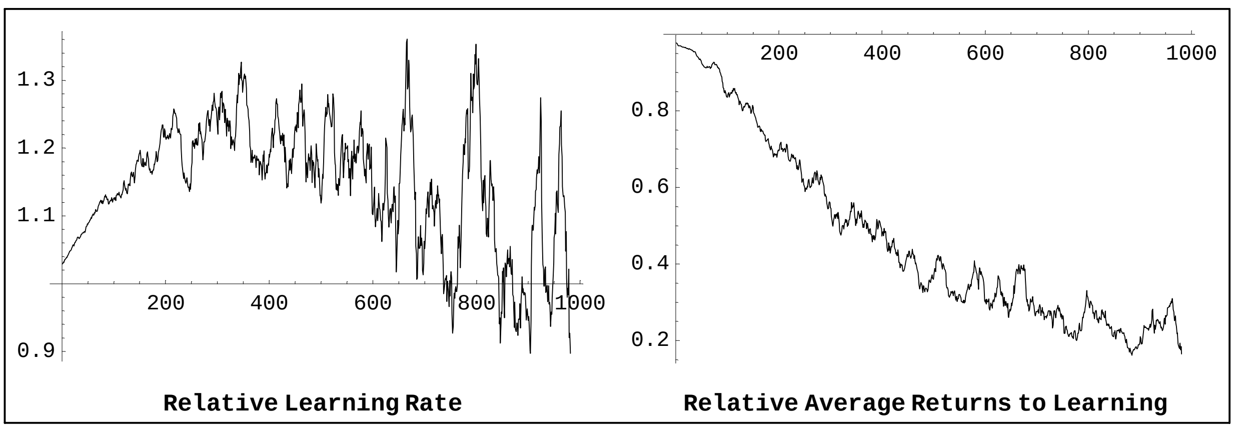

Figure 6: Relative Learning for (E = 0, N = 30)