Why is Five a Crowd in the Market Share Attraction Model: The Dynamic Stability of Competition

September 6, 2001[1]

By

Paul Farris*,
Phil Pfeifer*,
David Reibstein**,
And
Erjen van Nierop***

Send comments to: farrisp@virginia.edu

The authors wish to thank Thomas Gruca, Praveen Kopalle, Prasad Naik, Scott Neslin, and Peter Leeflang, and others for their valuable comments during the development and revision of this manuscript.

* Darden Graduate School of Business Administration, University of Virginia
** The Wharton School, University of Pennsylvania
*** Tinbergen Institute, Erasmus University Rotterdam

[1] This paper is a substantial extension of the earlier paper, "Number and Competitors and Dynamic Stability of Competition: Five is a Crowd in the Market Attraction Model." This paper replaces and extends the previous one.

Abstract

As economists have increasingly accepted the notion that non-equilibrium behavior may be a reflection of the "real world," interest has grown in modeling the emergence of turbulence. Using simulation to conduct a non-equilibrium analysis of a market characterized by the market share attraction model shows that the relatively simple competitive behavior of optimizing one's own spending could lead to surprising instability. Optimizing one's own spending for the next period requires that firms either (1) use their competitors' last-period budgets and the MSA model to determine a profit-maximizing budget or (2) use current customer response (dollar marketing input-dollar profit output) to set profit-maximizing budgets. These two approaches are equivalent in that they produce identical spending decisions. These decisions are myopic in the sense that they are profit maximizing under current competitive conditions. The surprising finding is that this behavior leads to dynamic instability if the number of competitors exceeds a threshold level of four. Even though an equilibrium exists, myopic profit maximizing firms will not find it if the number of competitors exceeds four. In this paper, we provide an analytical explanation for the emergence of this instability and demonstrate the close relationship between this nonlinear system of competing firms and the standard logistic map. We go on to show that this threshold for instability persists even if one firm has perfect knowledge of future competitor actions. In this work, we also identify conditions that affect the existence and level of the instability threshold.

Not only in (biological) research, but also in the everyday world of politics and economics, we would all be better off if more people realized that simple nonlinear systems do not necessarily possess simple dynamical properties. May (1986)

1. Introduction and Background on the Market Share Attraction Model

The purpose of this paper is to explore the dynamical properties of competition in a market where shares are determined based on the Market Share Attraction (MSA) model. This model, prevalent in the marketing literature, is the simplest representation of customer reaction to competition among firms that we know. It posits that a firm's share of market is equal to its share of "attraction."

$$s_i = \frac{A_i}{\sum_{j=1}^{m} A_j} \qquad [1]$$

The basic "us/(us + them)" form of this model captures two essential elements of competition. Competitors compete (1) for a limited objective resource (market share) and (2) with limited decision variables resources (brand attraction bought with investments in marketing). Limiting (or assigning a cost to) both objective and decision variables is an essential part of the notion of competition.

If the attraction of a brand is proportional to the firm's marketing effort, we can substitute marketing effort for attraction in [1] and have what Kotler (1984, page 231) refers to as "…the fundamental theorem of marketing." To apply this model to marketing mix decisions, it is necessary to specify the antecedents of "attraction." It is usual practice to assume that marketing effort or attraction is a function of the firm's marketing spending (on all mix elements both past and present). See Cooper and Nakanishi (1988, page 25) for a discussion of the subtle differences between the concepts of "marketing effort" and "attraction."

Several features of the MSA model help explain its popularity. As pointed out by Naert and Bultez (1973), the MSA model is theoretically appealing because it is logically consistent—shares are between zero and one and necessarily sum to one across all the firms in the industry. (Simple linear and multiplicative share models do not have this property.) Bell, Keeney, and Little (1975) show how this model derives from four simple axioms.[2] Nakanishi and Cooper (1974, 1982) point out that the model's parameters can be estimated using standard statistical packages. Finally, in an empirical study of two markets, Naert and Weverbergh (1981) tentatively conclude that the MSA model is a better predictor of market shares than either the linear or multiplicative model.

[2] (1) Attractions are non-negative and their sum is positive; (2) Zero attraction implies zero market share; (3) Equal attraction implies equal market share; and (4) A change in attraction has a symmetrically distributed effect on competitive market share.

Consequently, many researchers have used the MSA model in equilibrium analyses of marketing competition (Gruca, Kumar, and Sudharshan 1992, Karnani 1984 and 1985, Monahan 1987, and Steckel 1984). While the authors differ with respect to the exact form of the relationship between a firm's marketing spending and the resulting "attraction" of the brand, they all use [1] as the relationship between share and attraction.

While not discounting the importance of equilibrium analysis, we argue that the path toward equilibrium also deserves attention. In addition to exploring what an industry looks like when it reaches its equilibrium destination, it should also be useful to find out something about the journey.

The purpose of this paper is to do exactly that—explore the journey toward equilibrium of profit maximizing firms competing in a very simple MSA market. As the opening quote suggests, the simple *nonlinear* system of profit-maximizing firms competing in a MSA market exhibits some surprising and complicated dynamical properties. As first reported in Farris et al. (1998), depending on the number of firms in an industry, profit-maximizing firms can either: (1) quickly reach equilibrium; (2) slowly and predictably move to equilibrium; (3) move chaotically in a long path that eventually reaches equilibrium; (4) move chaotically on a path that never reaches equilibrium; or (5) after a sequence of generally increasing budgets that renders the industry unprofitable, suddenly abandon the industry.

Since reaching equilibrium is by no means a sure thing, these results suggest that the path toward equilibrium deserves its share of attention. Equilibrium analysis of markets is certainly less useful if the market cannot reach equilibrium. These results might also help explain the behavior of real markets. In particular, we believe that these results help explain why industry shakeouts and period of turbulence occur. And at a minimum, the results are a surprising implication of the simple profit maximizing behavior of firms competing in an MSA market.

The paper is organized as follows: Section 2 presents the simple MSA model used in our initial simulation study. Section 3 presents the results of that study. Section 4 is our explanation of why these results occur. Section 5 examines the generalizability of the findings, and Section 6 concludes with a review of the findings, some directions for further research, and some potential implications.

2. Formal Statement of the Base-Case Market Share Attraction Model

Because our purpose is to explore the basic nature of marketing competition, we keep things as simple as possible. We assume that attraction equals marketing effort, and marketing effort equals marketing spending. Thus we ignore the question of how a firm should allocate its spending across mix elements. Notice too that we have implicitly assumed that all firms are equally effective in spending their marketing dollars and that attraction achieved is proportional to dollars spent. In Section 5, we will relax both of these assumptions and explore firm difference in marketing effectiveness (Section 5.3) and nonlinear attraction/spending relationships (Section 5.7).

Letting X be marketing spending of firm i, the "base case" Market Share Attraction model is:

$$s_i = \frac{X_i}{\sum_{j=1}^{m} X_j} \qquad [2]$$

Keep in mind that X_i represents total marketing spending—not just advertising. Notice also that all firms are identical at this point. No firm has an advantage over any other. Shares and profitability will depend solely on how the firms select their spending levels relative to the competition. It is in this sense that the base case MSA model represents the essence of marketing competition. It is also worth noting that this base case model (even with no parameters) yields diminishing returns (decreasing marginal effectiveness) for marketing dollars for any firm facing a fixed competitive spending level.

We assume that firms compete for a share of industry contribution.[3] Letting *IC* be the total industry dollar contribution (assume for now that IC is unaffected by marketing spending), the net profit (*NP*) to firm *i* from spending X_i can be written as:

$$NP_i = IC\left(\frac{X_i}{\sum_{j=1}^{m} X_j}\right) - X_i$$

Taking the first derivative of NP_i with respect to X_i and solving for zero gives the following equation for X_i*, the spending level that maximizes firm *i*'s Net Profit.

$$X_i^* = \sqrt{IC \sum_{j \neq i} X_j} - \sum_{j \neq i} X_j \qquad [3]$$

It may come as something of a surprise that the equation for a firm's optimal spending is not a function of *m*, the number of firms in the industry. This is true because the base case MSA model [2] predicts shares will be assigned based on share of total spending–regardless of how many firms contribute to that spending. The number of competitors in the industry does not affect customer reactions. Under an assumption that firms are profit maximizing, we can use equation [3] to predict the spending level of firm *i* in any period.

We will use equation [3] to determine a firm's optimal spending in any period. We do not assume the firm actually uses this equation to set spending. In particular, we do not

[3] This could be true because firms compete for a unit market share and have identical dollar contribution margins, or because they compete for a dollar market share and have identical percentage contribution margins, or because they compete directly for a share of industry dollar contribution margin.

assume that firm *i* knows the spending of its competitors. Instead, we assume only that the firm has enough knowledge of the customers to be able to determine its optimal spending. Our firms simply know the customers well enough to be able to optimize spending at any point in time. While this may be unrealistic, it is consistent with the profit maximizing assumptions underlying all equilibrium analyses. We relax this assumption in Section 5.

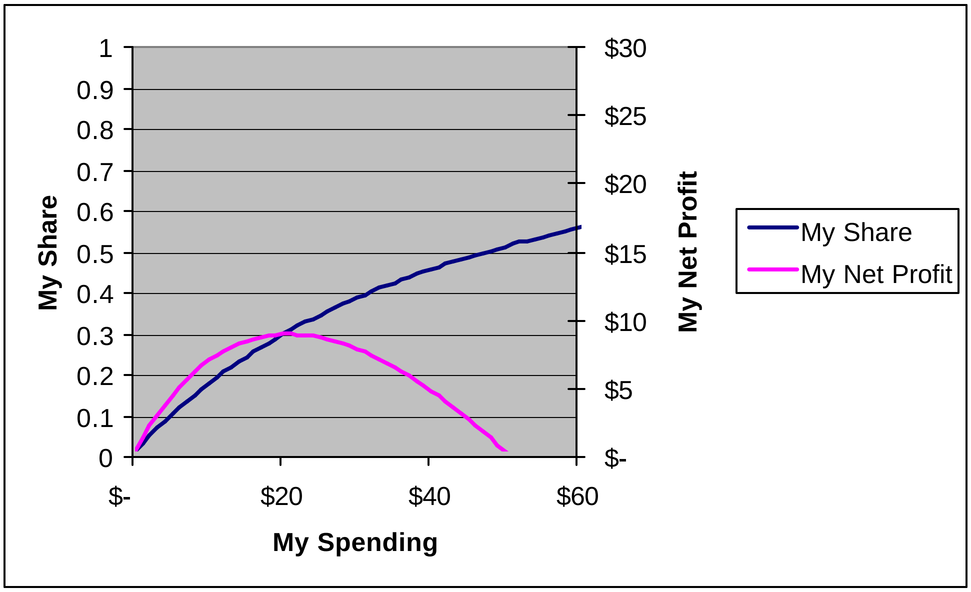

Figure 1a. Economic Analysis of the Focal Firm's Spending Decision

Figure 1a depicts the economics of the spending decision of the focal firm. We assume the firm knows the customers well enough to know how spending affects share. For this market situation, the focal firm knows it gets no sales if it spends $0 and about a 55 percent share if it spends $60. The relationship between share and spending, while strictly increasing, increases at a decreasing rate. It is nonlinear. The cost of spending, however, is strictly linear. Thus net profit first increases and later decreases as the constant cost of spending overtakes the decreasing returns from that spending. For the customers pictured here, optimal spending is about $21. At that spending, the firm expects a 29 percent share and a net profit of $9.

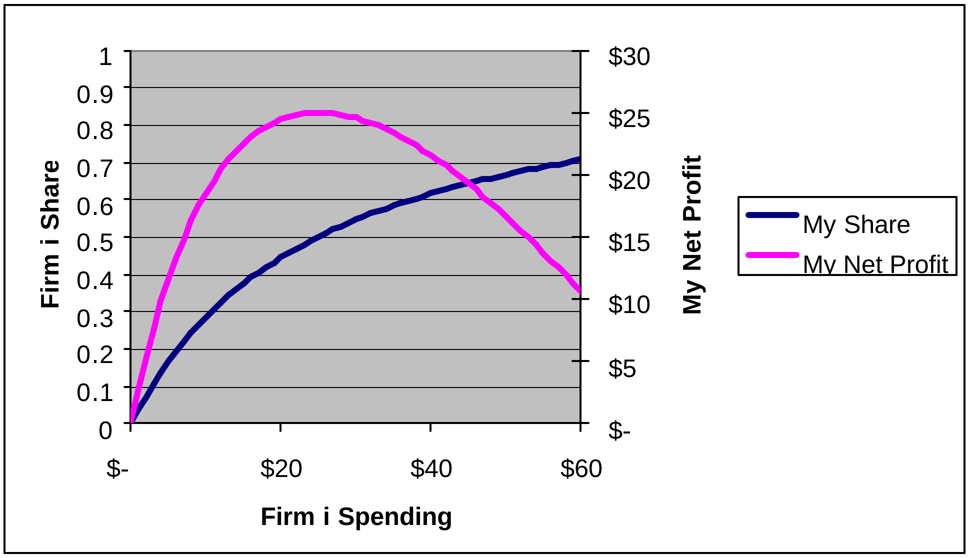

Figure 1b. Economic Analysis of the Focal Firm's Spending Decision in a New Situation

Figure 1b depicts a different situation, one in which market conditions have changed for the better. The customers are now much more responsive to the focal firm's spending—reflected in the higher share curve. In this new situation, optimal spending is about $25, expected share is 50 percent, and anticipated net profit is $25.

Figures 1a and 1b represent the assumed knowledge and behaviors of the firms in the industry. They know the customers well enough to be able to choose spending accordingly in order to maximize net profit. So, while the firms pay very close attention to the customers, they pay no direct attention to the competition.

But of course we know that our firms do react *indirectly* to the competition. For in the model we have constructed, the only reason the firm's customer response changes is that competitors changed their spending. (Competitors' spending is $50 in Figure 1a and $25 in 1b) Thus, if competitors change their spending, the firm's customer response changes, and the firm subsequently changes its spending. So spending is a direct function of customer response and an indirect function of competitive spending. Our firms do react to competitive spending by focusing on the customer response that changes with each competitive action.

Ultimately, budgeting decisions must be converted to dollars invested and dollars returned if the firm wishes to understand the profit implications of its decisions. It is clear, of course, that firm profits depend on both its spending and the total spending of competitors. Consequently, profit-maximizing firms' budgeting decisions must implicitly or explicitly account for competitive spending.

Although it may be tempting to think about customer response as if it were somehow separate from competitors' actions, we do not believe this distinction is useful or possible for the MSA. Customer response in the MSA model (see equation [2]) is a function of own spending and competitive spending and nothing else. Thus, customer response is totally and completely determined by *own and competitive spending*. Forecasting customer response to one's own budget *requires* some assumption about competitive spending. If the firm experiments to determine the current customer response to *dollar* spending and then sets its future budgets accordingly, it operates under an implicit assumption that customer response to spending will not change—which is true only if competitive spending remains the same as it was during the experimental period.

Equation [3] (our equation for the firms' profits) captures the behavior of the firms in this system. The focal firm's optimal spending turns out to be a function of only two things: the industry contribution *IC* and total competitive spending. Figure 2 graphs this relationship. If the rest of the industry spends \$50, firm *i* will spend \$21 (the example in Figure 1a). If the rest of the industry spends \$25, firm *i* will spend \$25 (the example in Figure 1b). More generally, if the rest of the industry spends less than *IC*/4, firm *i* will outspend the rest of the industry. If the rest of the industry spends more than *IC*/4, firm *i* will under-spend the rest of the industry. If the rest of the industry spends *IC* or more, firm *i* will spend nothing. Under no circumstances will firm *i* spend more than *IC*/4—and spending *IC*/4 is optimal only if the rest of the industry spends exactly *IC*/4.

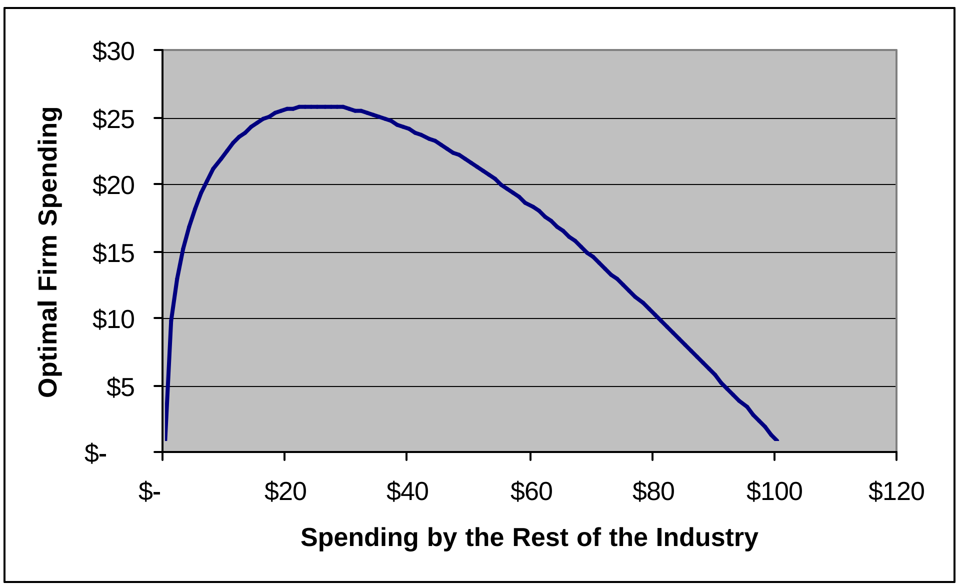

Figure 2 - Optimal firm spending as a function of spending by the rest of the industry for IC = \$100.

Before examining the non-equilibrium behavior of profit maximizing firms (the journey), we present the equilibrium solution (the destination). The Nash equilibrium is the point where no competitor has an incentive to change its decision variable unilaterally (in this case, marketing spending). An equation for the Nash equilibrium follows almost immediately from [3]. At the Nash equilibrium, all *m* firms will spend equally and optimally. If X_{NE} represents the Nash equilibrium spending per firm, then:

$$X_{NE} = \sqrt{IC \sum_{j \neq i} X_{NE}} - \sum_{j \neq i} X_{NE}$$

resulting in:

$$X_{NE} = \frac{IC(m-1)}{m^2}$$

If the m firms spend at the Nash equilibrium, the net profit to each firm is $NP_{NE} = IC/m^2$, and the total net profits for all m firms is IC/m. Mills (1961) first presented these equilibrium results.

It is here where many researchers end their journey. For some, finding and studying the equilibrium is an end. For us, it is just a beginning. As we have seen from our earlier research (Farris, et al., 1998), it will not be a trivial matter for firms to end up at the equilibrium. With a small number of firms it appears easy. For $m = 5$ firms it appears the firms can end up searching forever, and if six or more firms participate the search could lead to the abandonment of the industry.

In the next section, we will review and extend this earlier research and try to shed some light on the rather bizarre findings.

3. Simulation Analysis

Assuming that the base case MSA is a good representation of customer reactions to competitive spending, we describe a simulation designed to explore marketing spending patterns. Simulation is quick, flexible, and well-suited to the task. At some point, more advanced techniques such as those illustrated by Biscani, Gardini, and Kopel (2000) could be used.

We simulated the behavior of m firms for selected values of m starting all m firms at identical spending levels in period 0. In subsequent analyses, we will test to see if the initial spending levels affect the overall behavior. The total industry dollar contribution *IC* was set at $100, an arbitrary amount. At the beginning of each period of the discrete-time simulation, the m firms optimized their spending based on the market conditions at the end of the previous period.

Figures 3A through 3D illustrate the main results from this simulation study. Each of these figures graphs the time series of total industry spending of the *m* profit-maximizing firms. The surprise (to us) was that total industry spending showed stable patterns for two-, three-, and four-firm industries, but chaos reigned when a fifth firm was added. How and why does this happen? Notice also that with fixed industry size and costs, industry total profits are simply $100 minus the industry spending amounts shown in these graphs. Thus, two- though four-firm industries showed stable profit patterns while the fifth firm led to instability.

Let us comment in more detail. With two competitors, the industry moves quickly and steadily to equilibrium. With three and four competitors, industry spending oscillates around the eventual equilibrium level. (Notice, however, that the equilibrium spending level depends on *m* as illustrated earlier.) The major difference among the patterns for *m* = 3 and 4 is the rate at which the firms converge to equilibrium—very quickly for *m* = 3 and quite slowly for *m* = 4. These patterns appear to be unaffected by the starting value.

The story changes most dramatically when we move to five firms. Industry spending appears to be chaotic, showing almost no pattern or predictability. While we have shown earlier that an equilibrium spending level exists, it appears as if profit-maximizing firms may never get there unless they start there or stumble there through blind luck. In period 21, this industry does stumble very close to equilibrium—but doesn't stay there for very long. Indeed, this model appears balanced at the very edge of chaos.

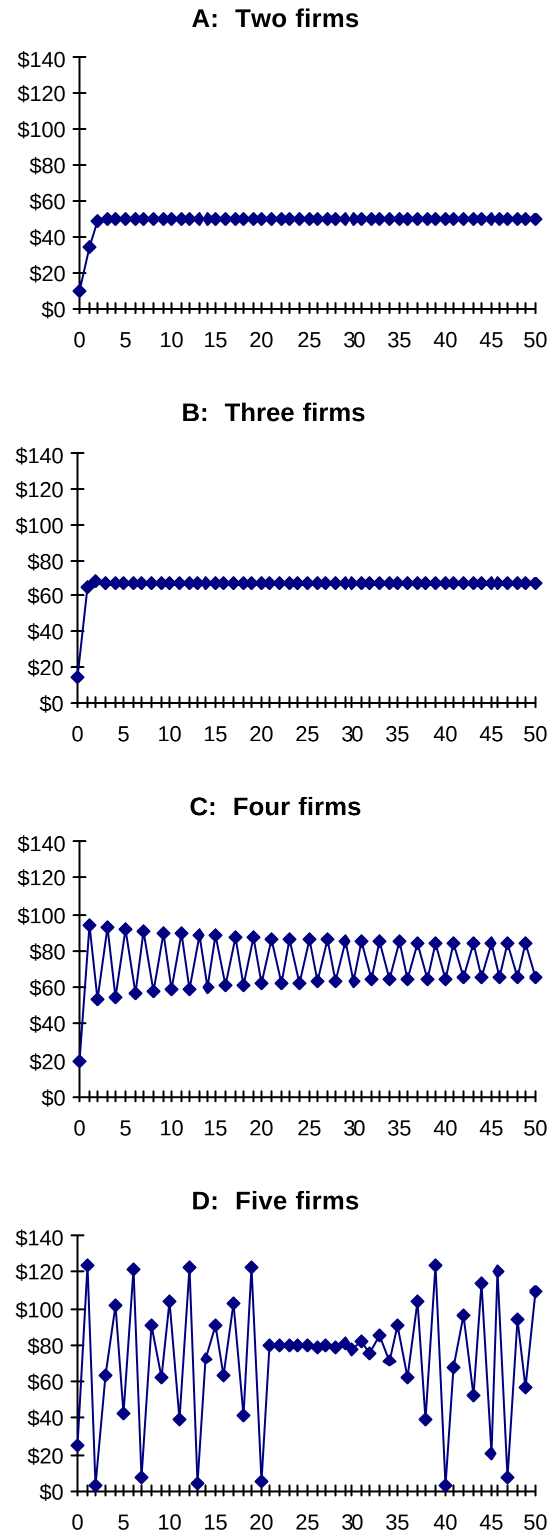

Figure 3 – Industry Spending for different numbers of competitors in the market

The pattern exhibited by the five-firm simulation is interesting because it may indicate a behavior of *collective* competition that emerges from the more basic (in our opinion) notion of competition. In other words, it is difficult to see how and why simple profit-maximizing firm behavior leads to equilibrium for two to four firms and then chaos for five firms.

Another reason for exploring this phenomenon further is that it may partially explain industry "shakeouts." Henderson (1976) claimed that "a stable competitive market never has more than three significant competitors, the largest of which has no more than four times the share of the smallest." Lambkin (1992) argued that a period of industry turbulence preceded shakeouts. Clearly, the chaotic pattern shown by the five-firm industry looks "turbulent" and could lead to shakeouts in an industry model allowing firm exit.

Before we examine the generalizability of these results, let us share our mathematical explanation for these initial results. Why does this happen? And what happens when $m > 5$?

4. An Explanation for the Emergence of Chaos at Five

We return to the graph in Figure 2 to help explain our simulation results. Recall that Figure 2 graphs the firm's optimal spending as a function of spending by the rest of the industry. In our multi-period model, we can reinterpret Figure 2 as graphing the optimal spending for firm i at period t as a function of spending by the rest of the industry in period t-1. Since all firms are identical, we can then use the appropriate straight line to map the spending of firm i at period t (a value on the y-axis) to the spending by the rest of the industry at period t (a value on the x-axis). We then repeat the process—using the spending by the rest of the industry at t to determine the spending of firm i in period t+1. Using lines and dots to keep track of all these moves results in a spider-web diagram that corresponds to the spending graphs in Figures 3A through 3D.

Consider first the spider-web diagram for m = 2 firms. Starting at t = 0 with each firm in the industry spending \$5, firm i will spend \$17.36 in period 1. If firm i spends \$17.36, the rest of the industry (in this case one other firm) spends \$17.36. The change in the "rest of the industry" generates a new optimal, and firm i spends \$24.31 in period 2, as does the rest of the industry. Again a new optimal budget is generated for firm i, and it spends \$25.00 in period 3. Now, however, the optimal budget for period 4 is the same as for period 3. Firm i stays at \$25.00 because that is the NE. It is clear, that for any starting value, spending will quickly converge to the NE.

The spider-web diagram for m = 3 (with initial spending of \$5 per firm), shows that spending now oscillates around the NE of (2/9)***IC*** as it converges rather quickly to the NE. The diagram for m = 4 is similar except that the rate of convergence to the NE is much slower. After 10 periods, spending is still quite a distance from $\boldsymbol{X_{NE}}$ = (3/16)***IC***.

The spider-web diagram for m = 5 is most interesting. (What was the spider

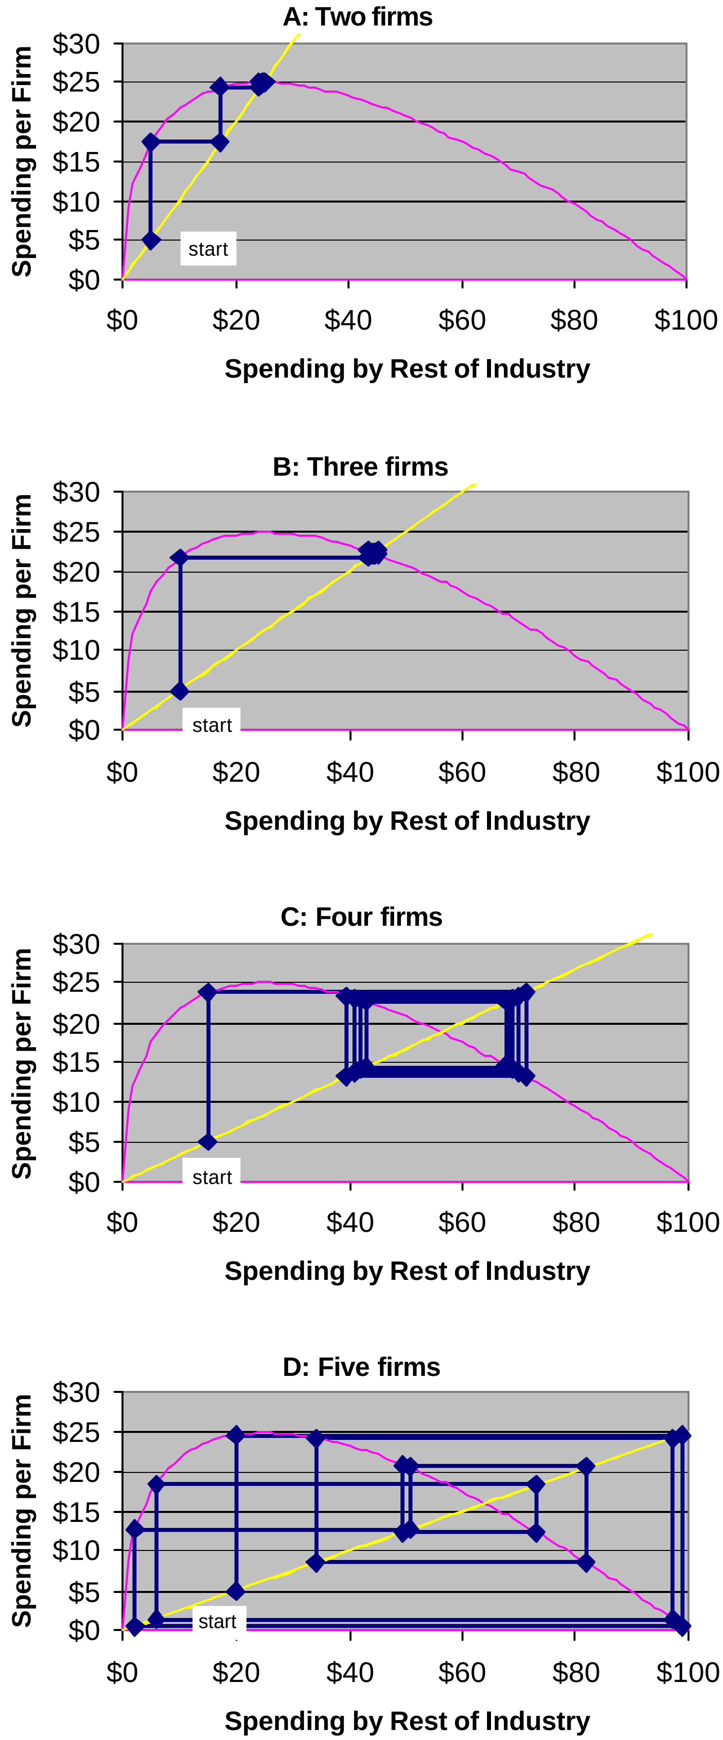

thinking?) The flatter straight line (because there are more firms in the industry) causes the spending pattern to be chaotic. While the pattern is perfectly predictable (there is no error in this model), it appears to be without pattern. Notice too that the straight line happens to go through the point ($100, $25) because for five firms, spending $25 by one firm is equivalent to spending of $100 by the rest of the industry. This means the spider is guaranteed to stay between zero and IC. It appears that spending may never hit the NE—or do so only by blind "luck." (It is difficult to even refer to luck here because this is a totally deterministic model.) If it ever did hit the NE, spending would stay there.

It turns out to be easy to show that a five-firm industry can find the equilibrium after any number of periods. By starting at the equilibrium and working the process backwards, we can generate a sequence of spending levels that lead to the NE. So, for example, if we want to show an industry that hits equilibrium in period 100, all we need to do is start the industry at the 100th spending amount in this backward sequence.

These spider-web diagrams have shown us that the flatness of the straight lines (the number of firms in the industry) determines the spending patterns we will see. For m=2, 3, and 4, the lines are steep enough to guarantee convergence to the NE. For m=5, it appears the line is just flat enough so that convergence is not guaranteed and spending behaves chaotically as the rest of the industry's 4X reaction is enough to keep the industry from finding the NE.

4.1 What Happens with m = 6 Firms?

Let us turn now to a six-firm industry. The straight line in the spider web diagram is now just a little flatter. But this extra flatness makes a huge difference in that the spider is no longer guaranteed to stay in bounds. For m = 6, almost any starting point will, within a few periods, lead to a situation where firm spending exceeds $20. When this happens, the industry is unprofitable (industry spending of $120 exceeds the industry contribution of $100), and our spider is left hanging.

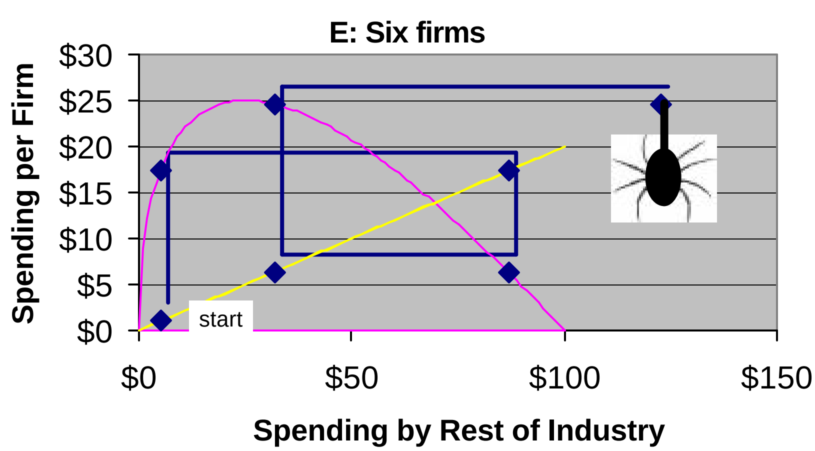

When this happens, what is a spider to do? Mathematically, equation [3] produces a negative optimal spending amount. Since negative spending is infeasible, it makes more sense to suggest the firm will spend zero. And if all m firms spend zero, shares for the subsequent period are indeterminate. We will label this outcome as an industry diverging into unprofitability. And because the subsequent all-zero situation is indeterminate, there is no easy way for our profit-maximizing firms to recover except to start over. In short, industries with more than five firms are simply not sustainable. In the base case MSA model, too many cooks definitely spoil the broth.

Before moving on, we must qualify our statement above that industries with more than five firms are unsustainable. As was the case with the five-firm industry, it is possible for a six-firm industry to stumble onto the equilibrium. We can also work the six-firm process backward from equilibrium to find a spending sequence that both stays in bounds and leads to equilibrium. Because we can construct such a sequence, it is possible (but very unlikely) for six-firm industries to stumble upon the equilibrium. The analysis techniques used by Biscani, Gardini, and Kopel (2000) would be useful to identify the exact conditions under which our simulations diverge to unprofitability or stumble onto the equilibrium.

4.2 The MSA Model, Competition, and the Logistics Equation

Readers familiar with the theory of dynamic systems may recognize our initial simulation analysis (with all firms spending the same amount) to be an exploration of the following recursive equation[4]:

$$X_t^* = \sqrt{IC(m-1)X_{t-1}^*} - (m-1)X_{t-1}^*$$

If we let Z_t^* represent firm spending as a fraction of IC ($Z_t^* = X_t^*/IC$), the relationship simplifies even further to one involving a single parameter:

$$Z_t^* = \sqrt{(m-1)Z_{t-1}^*} - (m-1)Z_{t-1}^*$$

With one more transformation letting $Y_t^* = [(m-1)Z_t^*]^{\wedge}.5$, we can rewrite the relationship as:

$$\left(Y_t^*\right)^2 = (m-1)Y_{t-1}^*(1-Y_{t-1}^*)$$

Note its striking similarity to the standard logistic map (see for example, Mira 1987; Devaney 1989).

$$Y_t = mY_{t-1}\left(1-Y_{t-1}\right)$$

Much is known about the behavior of this standard logistic map, and now we see that this knowledge sheds some light on our findings. Just as $m = 4$ is a bifurcation point separating stability from chaos in the standard logistic map, so too is $(m-1) = 4$ (or $m = 5$) the boundary between stability and chaos in the dynamic system of firms maximizing profit in a base case MSA market.

[4] Since we are modeling identical firms, we can replace spending for the "rest of the industry" with $(m-1)X_{t-1}$.

There is a further similarity worth noting. In 1845, P.F. Verhulst first used the simple logistic map to model the development of populations in limited environments (May, 1976). Just as organisms multiply quickly when there are very few of them and die off when their numbers approach the carrying capacity of their environment, so too does marketing spending expand when not many dollars are spent and contract when spending approaches *IC*, the limit the industry can support. The *m* parameter in the standard logistic map represents the unencumbered "growth" rate of the biological population, and (m-1) represents the rate with which the industry reacts to current spending.

4.3 Dynamical Properties of Profit Maximizing Firms in a MSA Market

Before moving on, let us summarize the findings of our initial simulation study. In an industry where share is determined by the base case MSA model, the profit-maximizing behavior of firms does not always lead to equilibrium. The following outcomes are possible:

(a) the industry can move quickly to equilibrium;
(b) the industry can move slowly but predictably to equilibrium;
(c) the industry can move chaotically and stumble upon the equilibrium;
(d) the industry can move chaotically on a path that never reaches equilibrium; or
(e) the industry can become unprofitable and abandoned.

We have also shown that industry outcome is dependent on the number of firms in the industry. With two and three firms, outcome (a) always results. With four firms, outcome (b) results. With five firms, depending on initial conditions outcomes (c) or (d) may result. And with more than five firms, outcomes (c) or (e) results. Adding even small amounts of error (as we show below) will make (c) extremely unlikely.

If we combine (a) and (b) and ignore (c) as too unlikely to worry about, our fundamental result can be stated even more succinctly.

> The spending of profit-maximizing firms in the base case MSA market converges to equilibrium if there are fewer than five firms, behaves chaotically if there are exactly five firms, and diverges into unprofitability if there are more than five firms.

5. Generalizing the Results

It is almost impossible to consider these results without reflecting on whether they might be artifacts of a highly abstract model of competition, or, whether such turbulence has parallels in the world of competition between real firms. Two ways of pursuing this question suggest themselves. One is to look for actual competitive spending patterns and number of competitors that reflect the outcomes seen in our simulation study. Clearly, this needs to be done and requires considerable effort and time.

A second method that may guide selection of the most appropriate industries to examine is to modify the model parameters and restrictions to learn more about the circumstances under which this kind of turbulence is most apt to manifest itself. In other words, can we make some changes to the model to prevent/eliminate this behavior? Or, in still other words, how generalizable is the fundamental result articulated in the text box above?

Please keep in mind that what follows are the findings from extending our simulation analyses and must be considered tentative. Also, keep in mind also that we ignore outcome (c) listed above. Certainly the potential to stumble onto equilibrium always exists in any of the extended models we tested. Each of these extensions was analyzed separately and independently.

5.1 Firm Differences in Initial Spending Amounts

Starting the firms at different spending levels does not affect the fundamental result. For less than five firms, firms converge to the equilibrium regardless of where each firm starts. Five firms converge surprisingly quickly to a common spending value, which then follows the chaotic pattern illustrated in Figure 3D. (After only five periods, the range in the spending of five firms dropped to about 1 percent of total industry spending regardless of the range in initial spending.) So, whereas total spending is chaotic for a five-firm industry, there is strict order within that chaos. Firms quickly learn to move in lock step—even though those steps are chaotic. Our five firms march to the same drummer despite the fact that the drumbeat is unpredictable. Industries with more than five firms diverge into unprofitability regardless of where each firm starts.

5.2 Imperfect Estimation/Implementation of Optimal Spending

Does including error in the firms' spending affect the results? Certainly real firms do not have perfect knowledge of the customers nor do they have perfect control over much they spend in any period.

Young (1998) observed that a "remarkable feature of stochastic dynamical systems is that their long-term (asymptotic) behavior can differ radically from the corresponding deterministic process no matter how small the noise term is. The presence of noise, however minute, can qualitatively alter the behavior of the dynamics. But there is also an unexpected dividend. Since such processes are often ergodic, their long-run average behavior can be predicted much more sharply than that of the corresponding deterministic dynamics, whose motion usually depends on the initial state."

To investigate this possibility, we extended the model to allow actual spending for each firm in each period to be equal to optimal spending times a multiplicative error:

$$X_t = X_t^* e^e$$

where ε is Normal(0,σ^2).

Error in implementing spending does not appear to change the fundamental findings with a couple of exceptions. For five firms without error, the chaotic spending stays "in bounds." Firm spending is always greater than \$0 and less than $IC/(m\text{-}1)$. Because the chaotic behavior sends the industry perilously close to $IC/(m\text{-}1)$ spending, however, including error means that eventually (m-1) of the firms will collectively spend more than IC. When that happens, the remaining firm will see an unprofitable industry. The profit maximizing action for that firm is to exit the industry. Consequently including error means the turbulence inherent in the five-firm industry will eventually cause a firm to spend zero.

The other exception concerns the small chance that an industry with five or more firms stumbles onto the equilibrium. If this happens in a system without error, the industry stays at the equilibrium forever. With error, however, firms always spend slightly different amounts and consequently stumble onto equilibrium with zero probability.

In our base case MSA, introducing error into the budgeting decision not only did not achieve stability for $m>4$, it interrupted the smooth movement to equilibrium for $m = 4$ and made the overall patterns less recognizable between simulation runs of the same parameters. It is possible that another type or amount of "noise" might increase the tendency toward a noisy equilibrium. For example, if we introduce variations in assumptions about competitive spending, this might interfere with the "lock-step" patterns of budget changes that soon emerge under the present structure.

5.3 Firm Differences in Marketing Effectiveness

Consider now a more meaningful form of firm difference. Suppose some firms are simply better than others—either because they are more effective in their marketing spending or they are simply more preferred by the marketplace. To handle this, assume that attraction is proportional to spending and allow firms to have different constants of proportionality:

$$A_i = a_i X_i$$

Our base case analysis had $a_i = 1$ for all i.

Note that with different levels of effectiveness, firms will spend different amounts at equilibrium. Firms with more effective marketing spend at higher levels and achieve higher profits for $m>2$. For $m = 2$, the firms spend the same, but achieve different shares and profit levels.

As concerns our main variable of interest, industry spending, three distinct patterns emerge, depending on the differences in marketing effectiveness.

1. Below some minimum differences in effectiveness, the firms still spend in approximate lock step, and for $m>5$ they abandon the industry simultaneously (all spending zero).

2. The first threshold of differences causes the firms with the least effective marketing to temporarily spend zero. The subsequent effect of this zero spending is to allow the remaining firms to reach a lower level of spending corresponding to the number of firms with some spending (two, three, or four). In subsequent periods, this lower spending entices the less effective firms to reenter the industry. This re-entry results in a higher level of spending causing the less effective firms to again spend zero. The result is a regular cycle of higher and lower industry spending caused by the periodic exit of the least effective firms.

3. The second threshold of differences in marketing effectiveness causes the least effective firms to permanently exit the industry—effectively reducing the number of firms. As a consequence, even industries with a large initial *m* can eventually behave like smaller industries. For example, an 8-firm industry may transform to a 4-firm industry that converges to equilibrium.

5.4 Carryover Effects

Up to this point, we have only looked at markets where expenditures have all their impact in the current period. We now consider carryover effects. Suppose that a fraction λ of this period's attraction carries over to the subsequent period.

$$A_{i,t} = X_{i,t} + l A_{i,t-1}$$

With carryover, we see lover levels of equilibrium spending. The Nash Equilibrium with carryover is simply (1-λ) times the Nash Equilibrium without carryover.

With carryover, industries with three or fewer firms converge quickly to equilibrium. Industries with $m = 4$ firms appear to converge very, very slowly to equilibrium. And industries with five or more firms do not reach equilibrium.

The big change is in the volatility and pattern of industry spending for five or more firms. When all *m* firms simultaneously spend a high amount the previous period, they spend zero in the next period but do carryover some of last period's attraction. This means that in the subsequent period, the industry looks very attractive and all firms return with healthy spending. The result is that carryover eliminates abandonment, leading to oscillation for $m>4$.

5.5 Decision Timing

Until now we have modeled all firms making spending decisions simultaneously at the beginning of each period. Suppose instead that each firm updates its spending with probability *p*. The base case had $p = 1$. Since the periods are of arbitrary length, this model represents a situation where firms randomly examine and adjust their spending.

As you might expect, lowering *p* has the same effect as reducing *m*. On average, the behavior of these simulations depends on *pm*. For *pm* > 5, we observe extended periods of chaotic spending and/or abandonment.

5.6 One Smart Firm

Suppose one "smart" firm reads this paper and the other (*m*-1) do not. How should the one smart firm behave in order to take best advantage of the competition? We ask this question not so much because we believe one smart firm is a viable situation, but rather to continue to better understand the dynamics of competition in a MSA market. Consider the following strategies for the one smart firm:

(a) Do nothing. Play dumb and behave like everyone else.
(b) Spend to maximize net profit after anticipating one competitive move.
(c) Spend at the equilibrium level.
(d) Spend at some other level chosen to maximize average net profit.
(e) Select the optimal spending pattern that maximizes average net profit.

Clearly, strategy (e) dominates, but what does it look like? And how well do the others do? Table 1 summarizes our findings.

Table 1. Net Profits to the One Smart Firm

m	Play Dumb	Anticipate One Move	Spend the Equilibrium	Spend the Optimal Amount	Spend Optimally	All at Equilibrium
2	*$ 25.00*	*$ 25.00*	*$ 25.00*	*$ 25.00*	**$ 50.00**	*$ 25.00*
3	*$ 11.11*	*$ 11.11*	*$ 11.11*	$ 12.50	**$ 50.00**	*$ 11.11*
4	*$ 6.25*	*$ 6.25*	*$ 6.25*	$ 8.33	**$ 50.00**	*$ 6.25*
5	$ ***6.09***	*$ 4.00*	$ 13.81	$ 35.00	**$ 50.00**	*$ 4.00*
6	abandoned	*$ 2.78*	$ 41.42	$ 43.43	**$ 50.00**	*$ 2.78*
7	abandoned	abandoned	$ 41.87	$ 46.38	**$ 50.00**	*$ 2.04*

The industry reaches equilibrium (yellow)
Chaos
The competition oscillates (green)
The smart firm and the competition oscillate (magenta)

The play-dumb strategy is equivalent to our base case. Notice that the average profit for the firms in the chaotic *m* = 5 industry is higher, on average, than the net profit at equilibrium. On average, the firms benefit from the chaos. (The table reports net profits at equilibrium for the cases that reach equilibrium. The $6.09 amount for the *m* = 5 industry is the average over 100 periods starting at spending of $5.)

If the one smart firm anticipates one competitive move, its spending has a calming effect on the industry. The one smart firm spends a small amount when the industry spends a large amount, and vice versa. Consequently, the industry reaches equilibrium for six or fewer firms. Although the one smart firm outperforms its competitors in the initial periods, that advantage goes away when the firms hit equilibrium. The table reports equilibrium net profits. The one smart firm gains little (relative to equilibrium) by anticipating one competitive move.

If the one smart firm jumps right to the equilibrium spending level, the industry quickly joins it if there are fewer than five firms. For five or more firms, the industry oscillates between high and low spending values. The one smart firm benefits a great deal from this oscillation as the net profit it makes in periods when competitors spend the small amount more than makes up for the loss it makes when the competition spends the large amount. This overreaction by the industry means that the profits to the one smart firm increase with the size (and overreaction) of the industry.

The one smart firm can take additional advantage of competitive oscillation by carefully selecting its spending level. The optimal single spending level for the one smart firm turns out to be the amount X, which causes competitors to oscillate between $(IC\text{-}X)/(m\text{-}2)$ and zero. The one smart firm intentionally induces pulsing behavior in the industry and gets 100 percent share every other period.

Inducing competitive oscillation between zero and some positive amount is also a characteristic of the one smart firm's optimal strategy. In the optimal strategy, the one smart firm careful pulses its own spending to induce the desired competitive spending. In one period, the one smart firm spends IC knowing that in the subsequent period competitors will spend zero. When the competitors spend zero, the one smart firm needs only spend an epsilon amount to capture 100 percent share and profits of IC minus epsilon. In the subsequent period, the competition will spend some delta greater than epsilon against the one smart firm's spending of IC. In the limit, the one smart firm will always get 100 percent share (epsilon against competitive zeros or IC against competitive deltas) and profits that oscillate between IC (when it spends epsilon) and zero (when it spends IC). In the limit, the one smart firm's pulsing behavior will achieve average net profits of $IC/2$ regardless of the number of firms in the industry.

5.7 Nonlinear Attraction/Spending Relationships

The base case assumes attraction is a linear function of spending ($A_i = X_i$). Here we consider two nonlinear attraction/spending relationships—one decelerating (decreasing returns) and the other accelerating (increasing returns).

The decelerating relationship considered is one in which attraction is proportional to spending raised to a power ($A_i = X_i^b$, with $b < 1$). This functional form is quite common and is used in what Cooper and Nakanishi (1988) call multiplicative competitive interaction (MCI) models. Notice that our base case MSA model is equivalent to the MCI model one spending element with $b = 1$. The accelerating relationship considered is one in which attraction is an exponential function of spending ($A_i = \exp[bX_i]$, with $b < 1$). This functional form is also quite common, and the resulting model would usually be called a multinomial logit (MNL) model. In keeping with Cooper and Nakanishi (1988), we refer to both the MCI model and the MNL model as attraction models since both incorporate [1].

Our analysis of the MCI model was limited by the fact that a closed form expression for optimal spending was not available. Numerical analyses for beta equal to 0.4 revealed that industry spending converges to equilibrium even for very large industries. Thus, the fundamental results change if attraction is a decelerating function of spending. Our initial hunch is that system performance is a function of m^b.

For the MNL model with all firms in lock step, optimal firm spending turns out to be a linear function of competitor spending:

$$X_i^* = X_j + \ln[(m-1)/2]/b + \ln\left(IC - 2 + \sqrt{(IC-2)^2 - 4}\right)/b\,.$$

With the increasing returns (attraction to spending) inherent in the MNL model, profit-maximizing firms will outspend competing firms by a constant amount each period. This results in industry spending increasing linearly and consistently over time. Eventually, the industry reaches unprofitability. This happens regardless of the number of firms. The number of firms affects only the speed at which the industry marches to unprofitability. As shown by Gruca and Sudharshan (1991), this simple MNL model has no equilibrium.[5]

Once the industry reaches unprofitability (by spending more than IC/m), the above equation represents a local optimum spending level for each firm. At this local optimum, the firm expects to lose money, but would lose even more money if it spent slightly more or less than the "optimal" amount. Only the most aware of firms would realize that at this point spending zero is preferred to spending "optimally." Firms focused on using customer response to maximize profits around their current spending level might easily get stuck on the local optimum—incurring several periods of increasing loses before realizing the situation is hopeless. Sound familiar?

In summary, nonlinear attraction/spending relationships do change the fundamental results. When attraction is a decelerating function of spending, we suggest that system performance is a function of m^b. When m^b is less than five, spending will converge to equilibrium. When greater than five, spending will diverge into unprofitability. When attraction is an exponential function of spending (increasing returns to spending), profit maximizing firm behavior leads to industry spending increasing linearly and inevitably into unprofitability.

[5] As already noted by Gruca and Sudharshan (1991), this simple MNL model has no equilibrium when $A=\exp(ß*X)$. Basuroy and Nguyen (1998) show that the more general $A=\exp(g(x))$ can have an equilibrium if $g(x)$ is sufficiently concave. In this case, there is a range of decreasing returns, and equilibrium exists. It is, however, still unclear, whether it will be reached with profit optimizing budget rules like we use in this paper.

5.8 Variable Market Size

Consider a market whose size is a nonlinear, decelerating function of total industry spending of the following form:

$$IC = B(\sum_i X_i)^q$$

where q is a parameter less than 1 and B is a base market size. Our analysis of an industry with $q = 0.4$ found convergence for industries with two or three firms, oscillation for a four-firm industry, and simultaneous exit for industries with five or more firms.

6. Discussion

This paper illustrates the complex dynamical properties of a nonlinear system of firms maximizing profit in a simple MSA market. While the system easily reaches equilibrium when the number of firms is small, it becomes chaotic with five firms and diverges into unprofitability with more than five firms.

By way of explaining these results, we showed that the recursive equation describing firms' spending behavior is very similar to the standard logistic map—an equation used to describe the boom and bust behavior of biological populations. In hindsight, it comes as no surprise that the behavior of firms competing for share of market is similar to that of living populations competing for survival in a world with limited carrying capacity.

6.1 Managerial Explanation

This mathematical explanation and the results of our sensitivity analyses suggest a managerial explanation. The failure to reach equilibrium is caused by firms paying too much attention to the customers compared to understanding competitive action (which, of course, might be viewed as reactions to our previous decision). Each firm's reaction to customer response, while optimal from its own point of view, becomes an industry overreaction when all m firms react simultaneously. The firm's "mistake" is in not recognizing that m-1 other firms are looking at the same customers and potentially reacting in the same way. The end result is m individually optimal reactions adding together to from an industry overreaction. When m is small enough, this overreaction does not keep the industry from reaching equilibrium. But at some point (m = 5 in the base case), the overreaction prevents an equilibrium outcome.

It is somewhat perverse that when the number of competitors is small (three or fewer), and one might think that it would be relatively easy to consider their moves, the naïve optimizations can "stumble into equilibrium." As the number of competitors increase, forecasting their collective behavior seems to become more important, but potentially more difficult.

The concept of industry overreaction also helps explain the results of many of our sensitivity analyses. Industry outcome (whether or not an industry reaches equilibrium) will depend on the "reactivity" of the industry. The more reactive the industry, the less likely is it to reach equilibrium. In the base case, reactivity depends only on m. The higher the m, the more reactive the industry, and the less likely is an equilibrium outcome. The sensitivity results in Section 5.5 refine this concept by suggesting that it is the number of reacting firms, pm, that affect industry outcome.

Sections 5.7 and 5.8 show that customer responsiveness to spending also contributes to the reactivity of the industry. If the relationship between attraction and spending is decelerating (X^{β}), customers are less responsive to advertising than in the base case. Consequently, the industry is less reactive, and industries with a large number of firms can reach equilibrium. In contrast, if the relationship between attraction and spending is accelerating ($e^{\beta X}$), customers are more responsive to advertising. Therefore the industry is more reactive, and equilibrium is harder (impossible) to achieve. Finally, if total industry contribution grows in response to spending (Section 5.8), the customers and industry are more reactive and equilibrium is more difficult to obtain.

The sensitivity results for carryover effects (Section 5.4) are more challenging to interpret. It turns out that carryover effects do not dampen the responsiveness of customers to spending. Current attraction is still a linear function of current spending, and the carried-over attraction is a (sunk) constant with respect to current spending decision of the firm. One dollar of spending buys one additional unit of attraction, just as in the base case. However, carryover effects do limit the firm's ability buy the optimal amount of attraction whenever it carries over more attraction than is optimal for the current period. In those cases, the firm spends zero but would like to sell back some attraction. This boundary effect means firm spending is slightly less responsive than in the base case.

6.2 Implications for Researchers and Managers

The findings reported in this paper are relevant to both researchers and managers. For those who study markets, we argue that while equilibrium analyses are important so too should be the analysis of non-equilibrium behaviors. The complicated dynamical properties illustrated in this paper are neither easily dismissed mathematical curiosities nor model defects to be fixed. Just as researchers in the life sciences have come to accept these properties as important descriptors of the actual behaviors of living populations, so too should we consider reaching the same conclusion about markets. This work suggests that turbulence in real markets may be caused not only by exogenous random factors, but also by the simple behavior of firms acting in a nonlinear system.

A potentially important implication for managers is that failing to project competitive actions and attending exclusively to customer response can lead to overreaction as an industry. This is somewhat at odds with the more prevalent notion that competitive "wars" are typically caused by paying too much attention to competitors and not enough to customers. This reinforces the idea that an important managerial function is attending to competition and anticipating their actions and reactions. Unfortunately, this paper also indicates that predicting competitor behavior will be more difficult when the number of competitors increases. If you think you are smarter than the rest, Section 5.6 suggests ways to take advantage of the overreaction of your competitors. Short of that, the paper offers an explanation for what may be happening in your industry. Of course this raises the question of whether actual markets exhibit the kind of instability to the simulated industry. Our belief is, "yes, sometimes, they do." We believe there is evidence that markets and marketers in the real world do often over-react. Airline price wars, dot.com advertising budgets, and hotel/office space capacity booms and busts are examples. Often these are attributed to lags in the feedback effects. Possibly, some of these are due to inherently unstable numbers of competitors targeting the same market. Many firms who were convinced of the need to move at Internet speed are now licking their wounds and developing respect for what they don't know about their competition.

As an industry, is it easy to "create" spending rules that lead to equilibrium? "Putting the brakes" on spending changes ensures that almost always, eventually, the firms reach the equilibrium. In our simulations, limiting spending to +/- 10 percent of the previous period can help firms reach the equlibirum. This also prevents small errors in spending from moving too far away from this level. Those familiar with the importance of "last year's budget" in determining "next year's budget" may find some comfort that such managerial practices may have some utility, even in fiercely competitive environments. Of course, "common sense" and natural caution may cause managers to discount opportunities that appear to be too good to be true

Our objective was not to "guide" firms to a NE, but to learn the patterns of competitive activity (spending) that are associated with the naïve profit-maximizing decision processes and the base case of us/(us + them) model of competition. Some of these patterns, we think, may be inherent in the very notion of "us versus them" competition that appears to be so stable and forgiving in the range of two through four competitors and chaotic and turbulent with five or more competitors. Our simulations include firms that use relatively simple optimization rules that result in collective behavior that surprised us with its complexity and unpredictability. Young (1998) writes "one of the central messages of the pure theory of exchange, for example, is the ability of prices and markets to coordinate economic activity without assuming that the agents are anything more than naïve optimizers acting on limited information." Young's diagram of the feedback loop in this kind of system is similar to the way the simulated firms behaved. Historical precedent (competitor spending) is the basis for formulating expectations about the economics of certain actions (their own spending). Collectively, these actions change expectations.

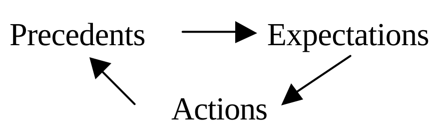

Our model assumes that firms use customer knowledge to select their optimal spending budget—but do so without regard for future competitive movements. Is this a realistic assumption? Certainly real firms often use test markets to estimate customer response as a precursor to selecting spending for the next budgetary cycle. In this process, what assumptions do firms make about the competition? Urbany et al. (2000) suggest that most firms fail to make any explicit assessment of competitive reactions. "Consideration of potential competitive reactions ... still has an incidence below 20 percent."

An empirical test of the propositions in this article will not be easy. There are always issues in defining markets that may include segments. Each segment can be argued to constitute a different market, and no simple guidelines are available to ensure that researchers see markets the same way managers do, or that managers perceive market boundaries in the same way. It is possible that some natural experiments might be found in periods after deregulation opened up geographic and product markets to more competitors. Banking might be one such example. Perhaps a leading indicator of the potential for turbulence might be the synchronization of competitive moves. In our models, the competition was drawn into "lock step" by the logic of the MSA model and the enforcement of a decision time-frame. In the real world, there are aspects that may encourage this synchronization: profitability reporting by standardized calendar periods, market research companies and customers that supply competitors with the same information at the same time, holidays and seasonal demand peaks and valleys that force competitors to all act at the same time.

References

Bell, David E., Ralph L. Keeney, and John D.C. Little (1975), "A Market Share Theorem," *Journal of Marketing Research*, 12 (May), 136-41.

Basuroy, Suman and Dung Nguyen (1998), "Multinomial Logit Market Share Models: Equilibrium Characteristics and Strategic Implications," *Management Science*, 44-10 (October), 1396-1408.

Bischi, Gian Italo, Laura Gardini, and Michael Kopel (2000), "Analysis of Global Bifurcations in a Market Share Attraction Model," *Journal of Economic Dynamics and Control*, 24, 855-879.

Cooper, Lee G., and Masao Nakanishi, *Market-Share Analysis*, Kluwer Academic Publishers, 1988.

Devancey, R.L. (1989), *An Introduction to Chaotic Dynamical Systems*, The Benjamin/Cummings Publishing Co., Menlo Park.

Farris, Paul, Willem Verbeke, Peter Dickson, and Erjen van Nierop (1998), "Path Dependencies and the Long-term Effects of Routinized Marketing Decisions," *Marketing Letters*, 9(3), 247-268.

Gruca, Thomas and Sudharshan, D. (1991), "Equilibrium Characteristics of Multinomial Logit Market Share Models," *Journal of Marketing Research*, Vol. XXVIII (November 1991), 480-2.

Henderson, Bruce D. (1976), "The Rule of Three and Four," in Carl. W. Stern and George Stalk, Jr., (editors), *Perspectives on Strategy from The Boston Consulting Group*, John Wiley & Sons, 1998, p. 31.

Karnani, Aneel (1984), "The Value of Market Share and the Product Life Cycle – A Game Theoretic Model," *Management Science*, 30 (June), 696-712.

Karnani, Aneel (1985), "Strategic Implications of Market Share Attraction Models," *Management Science*, 31(5), 536-547.

Lambkin, Mary (1992) "Toward a Theory of Market Shake-outs" University College Dublin Working Paper, 37 pages, July 1992.

May, R.M. (1976), "Simple Mathematical Models with Very Complicated Dynamics," *Nature*, 26, 459-467.

May, Robert (1986), "When Two and Two Does Not Make Four: Nonlinear Phenomena in Ecology," *Proceedings of the Royal Society*, Vol B228, 1986. p. 241, as quoted in

"The Chaotic Rhythms of Life," Robert May, in *The New Scientists Guide to Chaos*, Nina Hall ed, Penguin Books, London, 1992

Mills, H. D. (1961), "A Study in Promotional Competition," in F. M. Bass et al. (Eds.), *Mathematical Models and Methods in Marketing*, Richard D. Irwin, Homewood, Ill.

Mira, C. (1987), *Chaotic Dynamics*, World Scientific, Singapore.

Monahan, George E. (1987), "The Structure of Equilibria in Market Share Attraction Models," *Management Science*, 33 (February), 228-43.

Steckel, Joel (1984), "Using a 'Robust' Response Function to Allocation Resources in a Competitive Environment," Working Paper, Graduate School of Business, Columbia University.

Urbany, Joel E., David B. Montgomery, and Marian C. Moore (2000), "Competitive Reactions and Modes of Competitive Reasoning," Notre Dame, IN: University of Notre Dame, Mendoza College of Business, Working Paper, as cited by Urbany, Joel E., "Justifying Profitable Pricing," Marketing Science Institute Working Paper, Report No. 00-117, 2000.

Young, H. Peyton (1998), *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions*, Princeton University Press, Princeton, 1998.