Local Hillclimbing on an Economic Landscape

David Kane

SFI WORKING PAPER: 1996-08-065

Local Hillclimbing on an Economic Landscape

David Kane
John F. Kennedy School of Government
Harvard University, Cambridge, MA 02138
David_Kane@harvard.edu

Abstract

Profit maximization is difficult. Sophisticated and experienced managers often disagree about which action is most likely to maximize profits for a given firm. Economic models of profit maximization, on the other hand, are—in general—easy to solve. Well-trained economists can readily discern the action which maximizes the firm's objective function. The global maximum is unique and achievable because the objective function is designed to have this property. This paper weakens the assumption of analytically tractable objective functions. I propose a model of profit maximization in which it is, essentially, impossible for the firm to discover the global maximum. Firms have no choice but to, in the words of Charles Lindblom, "muddle through" in their attempt to find the optimal budgetary allocation in an extremely complex economic landscape [9]. Computer simulations provide details of that landscape as well as evidence that certain strategies may be more effective in difficult environments. Specifically, "patience" may be a virtue that applies to firms as well as to people.[1]

1 Introduction

Since the original formalization of the utility function approach to economic theory by von Neumann and Morgenstern in 1944, it has been standard to assume that the objective function is *tractable* [13]. That is, the preferences of an agent may be represented by a function which the agent is able to maximize, subject to certain constraints. An agent knows which specific values of the inputs maximize her objective function. For a large class of functions, this is a perfectly reasonable assumption. But, for an equally large class of alternative functions, it is not. The question is: Which class of functions—those which are easy to maximize or those which are difficult—is more useful for modeling the behavior of economic agents?

I propose that functions which are difficult to maximize will prove to be better suited for modeling purposes than functions which are not. The primary reason for this lies in the irreducible complexity confronting economic agents in the real world. In order to understand how and why agents make decisions and to suggest methods by which those decisions might be improved, it is necessary to model that complexity. This is the flip side to Simon's description of bounded rationality [12]. Failure to maximize can be explained in two complementary ways: either agents are not intelligent enough to find a findable maximum or the maximum is difficult for agents, no matter how intelligent, to find. See [11] for a survey of recent work focusing on the first explanation. This paper details a model of the second.

The next section provides the theoretical background to the classic example of profit maximization. A firm has a fixed budget which it must allocate among a specified set of production inputs. Different budgetary allocations yield different profit levels, and the firm would like to discover which combination of inputs produces the most profit. The third section describes a model of economic "landscapes." A landscape is a metaphor which highlights the importance of local knowledge in the search for increased profits. Firms presumably know more about the profit characteristics of input combinations which are "close" to their current position. There is a "neighborhood" of alternate input combinations, alternate points on a profit landscape, which are similar to the current combination. Firms search this neighborhood, move to points with higher profits, and thereby climb the peaks of the profit landscape. The landscape framework provides for an explicit model of the evolution of firm behavior [1]. The fourth section applies the landscape model to questions of firm strategy. Certain strategies for traversing the profit landscape are more successful, on average, than others. In particular, a "patient" firm which looks for gradual improvements in profits will achieve higher long-run profits than a "greedy" firm which insists on always moving to the highest available point on the landscape. The fifth section concludes.

[1]Published in *Evolutionary Programming V: Proceedings of the Fifth Annual Conference on Evolutionary Programming.* Edited by Lawrence J. Fogel, Peter J. Angeline and Thomas Bäck. MIT Press: Cambridge, Massachusetts, 1996, pp. 9–15.

2 Theoretical Background

To be concrete, consider the example of a firm which seeks to maximize its profits by selecting the optimal allocation of input spending subject to a budget constraint.

Definition 2.1 *The budget,* $B \in \{0, 1, 2, \ldots\}$.

The budget constraint and all input variables are restricted to the set of non-negative integers.

Definition 2.2 *The set of inputs,* $X = \{x : x = x_1 x_2 \ldots x_n \text{ with } x_i \in \{0, 1, 2, \ldots\}\}$.

The profit function, Π, is drawn from the class of functions which map a vector of N non-negative integers to $\mathcal{R}$.

Definition 2.3 *The budget constraint,* $\sum_{i=1}^{n} x_i = B$.

2.1 Complexity

For an arbitrary function Π, the firm's profit maximization problem is well posed but difficult. The space of possible solutions is huge and, without assuming global quasiconcavity for Π, it is difficult to know how neighboring points are related to one another. Profit maximization becomes a problem of combinatorial optimization. Indeed, von Neumann and Morgenstern themselves insisted that "The emphasis on mathematical methods seems to be shifted more towards combinatorics and set theory—and away from the algorithm of differential equations which dominate mathematical physics." [13, p. 45] Consider:

Claim 2.1 *The number of different combinations of input allocations which satisfy the budget constraint is* $\binom{N+B-1}{B}$. *For a proof, see [3].*

Even for relatively modest values of N and B, this is a very large number. For example, if $N = 20$ and $B = 50$ then there over 4.6×10^{16} possible combinations of integer budget allocations which exactly satisfy the budget constraint. This number increases approximately exponentially in N.

Moreover, the optimal allocation is extremely difficult to find if Π is non-trivial. The problem is NP complete. In contrast, assuming global quasiconcavity of Π allows a firm to solve for the profit maximizing budget allocation easily. Global quasiconcavity is important, and almost universally assumed, because it makes finding the optimal input combination a reasonable goal. Consider a profit maximization problem involving a budget of 50 to be allocated among 20 different inputs. Without global quasiconcavity or some other restricting assumption, it is impossible to know whether or not any particular local maximum is also the global maximum.

2.2 The Landscape Metaphor

In answering this question, it is helpful to picture profit functions as landscapes. Visualize the case for $N = 3$. All possible integer combinations of the three inputs which satisfy the budget constraint constitute a two dimensional grid on a simplex. Each combination of inputs generates a specific amount of profits via the Π function. Define the height of the landscape at any given point to be the value of profits generated by the combination of inputs specified by that point. A discontinuous set of points floats above the grid of input combinations. Higher points correspond to higher profits. Connect neighboring profit points to each other if their corresponding budget points are next to each other on the simplex. The resulting profit landscape describes the set of all possible profit amounts along with a method for moving from one level of profit to another by moving from one combination of inputs to another around the simplex. The metaphor, if not the easy visualization, can be extended to an arbitrary N.

The landscape metaphor originates in the study of population genetics [14]. Wright described a fitness landscape in which each specific genotype had an associated fitness, a measure of the reproduction likelihood of the organism defined by that genotype. Evolution could then be viewed as a search for higher points on the fitness landscape.[6] In political science, the quintessential landscape would be electoral; a particular position on the issues yields a specific number of votes [4]. Kollman *et al.* perform computer simulations of a complex political landscape to model party behavior [7]. Jones provides an excellent overview and analysis of the landscape metaphor [5].

3 An Economic Landscape

Whatever the merits of the landscape metaphor in biology and political science, however, the question remains as to whether or not landscapes are useful in economics. We can already imagine the traditional profit function as a surface in the three dimensional (or more) space above a grid of acceptable input combinations. How does calling this profit function a "landscape" provide additional insight into the underlying economic phenomenon of interest?

The answer is that the landscape metaphor highlights the idea of a "neighborhood" to which a firm could move and, therefore, the importance of uncertainty surrounding the status quo. Presumably, firms have better information about the profitability of input combinations near to their current position. Because they do not know the global optimum, they must search for it—or at least for points which improve on the status quo. Conceptually, that search occurs on a profit landscape.

3.1 Model

Consider a firm faced with the single period profit maximization problem defined in Section 2. The key unspecified part of the model is Π:

Definition 3.1 *The profit function,* $\Pi(X) = c_1x_1 + c_2x_2 + \ldots + c_nx_n + c_{n+1}x_1^2 + c_{n+2}x_2^2 + \ldots + c_{2n}x_n^2 + c_{2n+1}x_1x_2 + c_{2n+2}x_1x_3 + \ldots + c_{n*(n+3)/2}x_{n-1}x_n$ *with* $X = \{x : x = x_1x_2 \ldots x_n$ *with* $x_i \in \{0, 1, 2, \ldots\}\}$ *such that* $\sum_{i=1}^{N} x_i = B \in \{0, 1, 2, \ldots\}$ *and with* $c_i \in [-1, 1]$.

In words, the profit function is linear in the input variables, their squares and cross products. Because the profit function is the heart of the model, it requires a few words of explanation and justification. Assume that there is a "true" profit function. What are the characteristics of this (unknown) function and how does it compare with our proposal?

Approximation: Whatever the form of the true function, it may be approximated with a polynomial. Consider Π to be such an approximation. Moreover, the specification allows us to place the problem of price effects within the overall profit function. Depending on the how much of a given input a firm purchases, the price which it pays for that input may change. Definition 3.1 allows us to focus just on the share of the budget which is devoted to a particular input; the effect of any gains or losses associated with price movements caused by changes in budget allocations are subsumed within the profit function.

Returns: Standard economic theory describes firms as facing decreasing or constant returns. An additional unit of a given input therefore does not increase profits. However, recent work has helped to point out the importance of increasing returns [2, 8]. Clearly, there are situations in which increasing the budgetary share of a given input yields increasing profits, just as there situations in which increasing the same item yields constant or even decreasing profits. The proposed profit function mimics this characteristic.

Multiple local optima: It often appears to a firm that its particular combination of inputs is optimal because small departures from that combination lead, or can confidently be forecast to lead, to lower profits. One of the primary goals of the day-to-day functioning of the firm is, in fact, to maximize profits taking a particular budget allocation as given. Firms of the same size and in the same industry often allocate their budgets in very different ways. If all are optimizing, then perhaps all are at different local optima. For non-trivial choices of the c_i terms, the proposed profit function will feature numerous local optima.

Difficulty: The most important characteristic of our profit function is that it is difficult to solve. Even a firm which knows exactly what it is trying to maximize will have trouble finding the global optimum, or even knowing whether or not the specific local optimum which it has discovered is the global one. The primary contributors to this difficulty are the intermixing of decreasing, constant and increasing returns along with the existence of multiple local optimum. This economic landscape is rugged. Any landscape which is not rugged, on which the global optimum is easy to find, fails to capture the complexity of actual economic life.

There are obviously an infinite number, not just of functions but of classes of functions, which map an N-tuple of integers into the reals. A particular realization of our function, defined by a specific set of values for the c_i terms is an example drawn from one such class. A profit function, to be realistic, to adequately capture and model the difficulties faced by firms, must have certain attributes. First, it must be difficult, where difficulty implies the existence of numerous local optima and a hard-to-find global optimum. Second, that difficulty should derive from the underlying economics of the model. There are many rugged landscapes which have little connection to economics. The profit landscape outlined above is rugged because it allows for increasing and decreasing returns along with a variety of complimentarities—both beneficial and harmful—among various inputs. Third, the profit function, or the model in which it is embedded, should capture the importance of the status quo and of the information gleaned from the neighborhood surrounding it. In their search for higher profits, firms have to start from where they are. We now formalize the importance of vicinity to the status quo.

3.2 Neighborhoods

The landscape metaphor is unintelligible without the notion of a neighborhood. For a certain point to be a "peak," all the surrounding points must yield lower profits. But precisely which points are "surrounding," are in the "neighborhood," of a putative peak?

In the context of our model, a fairly straightforward definition of neighborhood is:

Definition 3.2 *The neighborhood of an input vector,* $X = \{x : x = x_1x_2 \ldots x_n$ *with* $x_i \in \{0, 1, 2, \ldots\}\}$ *consists of all those points,* $X' = \{x' : x' = x'_1x'_2 \ldots x'_n$ *with* $x'_i \in \{0, 1, 2, \ldots\}\}$ *with exactly one* $x'_i = x_i - 1$ *and* $x'_{i+1} = x_{i+1} + 1$ *and, for all* x'_j *such that* $j \neq i, i+1$ *then* $x'_j = x_j$.

In words, the neighborhood of a given budget allocation X is the set of points such that all the individual allocations are identical except for two; and, for those two, one input receives one more unit—call them dollars—and the other input receives one less unit. It is clear that if X satisfies the budget constraint, then all points in the neighborhood of X will also satisfy the constraint. Note that the redistribution of one dollar cannot go from a given input to any arbitrary input. Instead, each input is "partnered" with one other input to which it may transfer dollars. Each input may also, therefore, receive dollars from one other input. The subscripts denote this transfer scheme. Input i may receive dollars from input $i-1$ and transfer dollars to input $i+1$. (For $i = N$, transfers go to $i = 0$, thereby completing the cycle.) Any circular transfer scheme may be modeled by renumbering the inputs.

Other definitions of neighborhood are possible, and may even be more appropriate in certain circumstances. For example, the network of possible transfers could be more dense. Each input could transfer/receive dollars from two other inputs. This would double, in general, the number of points in the neighborhood surrounding any given point. Another change to the definition of a neighborhood would involve allowing larger changes in the budget allocations than just a single dollar. Consider a firm with the ability the re-allocate up to two dollars, from a single budget item, to other inputs. A neighborhood would then include the same set of points as before plus points which involve larger redistributions. Not only could the firm take one dollar from x_i and redistribute it to x_{i+1}, it could do the same with two dollars.

It is important to note that changing the definition of a neighborhood does more than simply change the framework for considering a given landscape; it fundamentally alters the landscape itself. A peak with one dollar redistribution may or may not be a peak under two dollar redistribution. A point which is "near" the global optimum under one definition of neighborhood may be "far" from the global optimum under another. It is necessary to include in the definition of landscape a definition of neighborhood. A landscape without neighborhoods is a landscape without structure. Neighborhoods are the conceptual device which transform standard profit functions into profit landscapes.

3.3 Descriptive Statistics

Consider a profit landscape with one dollar redistribution and one connection per input, under the conditions specified in Section 2 and using the profit function from Definition 3.1. All points on the landscape are accessible from any given starting point in a finite number of moves. For any particular realization of the landscape, the set of coefficients, the c_i's, are randomly drawn from the uniform interval $[-1, 1]$. The firm has a budget of 50 dollars which it seeks to allocate among 20 inputs in order to maximize its next period profits. There is no uncertainty. For any specified allocation, the firm knows what level of profits will result. All of the following results are relatively robust to changes in the size of the budget and the number of inputs.

3.3.1 Local Optima

We know that the space of possible budgetary combinations is finite but immense. Specifically, there are over 4.6×10^{16} possible budgets, all of which are accessible from any starting point on the profit landscape. The first question is: How many local optima are there? If there are very few, then the firm might consider a simple hillclimbing strategy. If there are many local maxima, then this strategy is likely to lead to a profit level which is below the global maximum.

In fact, out of $2,000,000$ individual budget combinations tested ($1,000$ unique points on $2,000$ different landscapes), only 8 were found to be local maxima—points such that each of the surrounding points yields lower profits. However, given the size of the entire landscape, this result still suggests that the expected number of local maxima in the entire space is at least 1.8×10^{11} unique input allocations. Even though local maxima form an extremely small subset of all points, the large number of points means that, in absolute terms, the number of local maxima is large as well. In general, firms that hillclimb from an arbitrary point will reach a local maximum which is both lower than the global maximum and different from the local maximum which they would have reached if they had started from a different initial point.

3.3.2 Comparative Statics

Of interest are the "comparative statics" of the model as three key parameters are varied: the budget, the number of inputs, and the number of connections to each input. It is obvious that, as the budget increases, profits should increase as well. To adjust for this effect, it is convenient to normalize the profit level by dividing by the maximum possible profit (the budget plus the square of the budget, which would occur if the coefficients for both the linear and quadratic terms of one input were both equal to 1 and if the entire budget were spent on this input).

Consider the expected profit of a firm which starts from an arbitrary budget allocation, surveys all the neighboring points, moves to the ones with the highest expected profit, surveys again, and so on, until it reaches a local maximum. How does the expected profit change as the size of the budget, the number of inputs, and the density of connections is varied?

Budget: In the absence of normalization, profits increase approximately as the square of the budget. With the normalization, however, the level of profit at a local maximum averaged over 100 different landscapes is statistically indistinguishable as the

		Number of Connections Per Input				
		1	2	3	4	5
Strategy	Steepest Ascent	15.6 (0.2)	24.7 (0.3)	32.1 (0.3)	37.9 (0.4)	43.0 (0.4)
	Median Ascent	16.5 (0.2)	25.3 (0.3)	32.2 (0.3)	39.0 (0.4)	43.7 (0.5)
	Least Ascent	19.1 (0.2)	30.1 (0.3)	38.4 (0.4)	43.1 (0.5)	48.8 (0.6)

Table 1: Mean normalized profits over 1,000 landscapes for different firm strategies and connections per input. Coefficients for the profit function are drawn from $[-1, 1]$. Standard errors in parentheses.

budget increases from 20 through 250, with a constant 20 inputs and 1 connection per input.

Inputs: As the number of inputs increase beyond a "small" number, normalized profits begin to fall. This is a generic result for any model in which each input is connected to less than half the other inputs. The intuition is that an increase in inputs, beyond some minimum, makes it more difficult to transfer dollars from one input to another which is "far away," given the constraint of the specified connections.

Connections: More connections mean that each point on the landscape has a larger number of neighboring points. Hillclimbing becomes easier because there are more directions in which to go. However, once the number of connections surpasses half the number of inputs, the landscape becomes so densely connected that no further improvements are possible.

4 Application to Firm Strategy

No one believes that an actual firm functions as a mechanical hillclimber, estimating all the points, and only those points, in the neighborhood of its current position and then moving to the best available alternative. The real economy is, obviously, infinitely more complex than this model. But it is possible that firms behave "as if" they were simple hillclimbers on a rugged profit landscape. In that event, it is interesting to investigate which strategies are most likely to yield large profits.

4.1 Computational Aspects

Consider three simple strategies for hillclimbing firms. Steepest Ascent (SA) means that a firm always moves to the highest point in its immediate neighborhood until it arrives at a local maximum. All the statistics in section 3 were generated under SA. Under Median Ascent (MA), a firm selects, from the subset of points higher than its current location, the median point and then moves. Least Ascent (LA), as its name implies, causes the firm to take the smallest upward step at each iteration. All three strategies iterate until a local maximum is reached. In these experiments, we are only interested in the average value of the local maximum reached—not in the number of steps required to reach it. A more complete analysis would consider the cost of sampling alternate points and the time spent at points below the eventually reached maximum.

Table 1 reports the mean (of $1,000$ landscapes) profit of the local maximum reached by each of the three strategies in a model with 20 inputs and a budget of 50 with increasing numbers of connections per input. (To make the following tables easier to read, all profits were multiplied by 100. This simple scaling has no effect on the conclusions.)

The most interesting result is the clear superiority of Least Ascent. A firm will perform much better if it is "patient," if, instead of always making the "best" move at each step across the landscape, it moves upward in the most gradual fashion possible. The underlying structure of the profit landscape is such that extremely gradual slopes tend to lead to higher local maxima.

The problem with computationally derived results, like those above, lies in their possibly knife-edge quality. It is possible that a small change in the model—in the coefficients of the profit function, say—will lead to completely different results. Fortunately, this danger can be minimized by exploring alternative formulations of the same model. If the results are the same for reasonable changes in the original assumptions, than we may be fairly sure that are conclusions are robust, that they are a generic feature of models like this one and not just of this one in particular.

Consider a different version of the profit function, one in which the coefficients for the linear and quadratic terms are drawn from $[0, 1]$ and the coefficients for the cross products from $[-1, 0]$. Even though the resulting profit function has more structure—i.e. makes more assumptions—than before, the complexity of interactions creates a landscape on which the global optima is every bit as hidden. Table 2 reports the results from a comparison of the three strategies on this new landscape.

Again, we see the superiority of Least Ascent. The same normalization correction is used in this table and the next, even though it is less well suited to these profit functions. However, doing so maintains comparability across tables.

Finally, consider a profit function in which there is a

		Number of Connections Per Input				
		1	2	3	4	5
Strategy	Steepest Ascent	-6.5 (0.2)	1.0 (0.2)	7.9 (0.3)	13.4 (0.3)	19.9 (0.4)
	Median Ascent	-5.7 (0.2)	1.9 (0.2)	8.3 (0.3)	14.4 (0.4)	20.9 (0.5)
	Least Ascent	-2.3 (0.1)	8.8 (0.3)	22.5 (0.5)	34.8 (0.7)	41.6 (0.7)

Table 2: Mean normalized profits over $1,000$ landscapes for different firm strategies and connections per input. Coefficients for the profit function are between $[0,1]$ for linear and squared terms and $[-1,0]$ for cross products. Standard errors in parentheses.

		Number of Connections Per Input				
		1	2	3	4	5
Strategy	Steepest Ascent	-6.52 (0.2)	-1.70 (0.1)	0.002 (0.1)	0.65 (0.1)	1.06 (0.1)
	Median Ascent	-6.06 (0.2)	-0.86 (0.1)	0.59 (0.1)	1.08 (0.1)	1.34 (0.1)
	Least Ascent	-5.43 (0.1)	0.10 (0.1)	0.74 (0.1)	1.08 (0.1)	1.09 (0.1)

Table 3: Mean normalized profits over $1,000$ landscapes for different firm strategies and connections per input. Coefficients for the profit function are 0 for linear terms, $[-19,0]$ for squared terms, and $[0,2]$ for cross products. Standard errors in parentheses.

penalty for input values which differ from one another. This Leontief-type specification derives its formulation from a penalty function in which the square of the distance from each input to every other input matters. It follows that there will be $N-1$ squared terms for each input (each with a negative coefficient) and two cross product terms for each pair of inputs. To include this assumption within the more general framework, we may use 0 as the coefficient for each linear term, add together $N-1$ separate draws from $[-1,0]$ for the squared terms, and combine two separate draws from $[0,1]$ for the cross products. The results are presented in Table 3.

Again, Least Ascent does at least as well as MA and SA, except for the five connections case. In general, the performance of all three strategies converge as the number of connections increase. Once the landscape becomes densely connected enough, all hillclimbers reach relatively high points.

4.2 Practical Aspects

How relevant are these results to the functioning of actual firms?

First, this model is most appropriate for firms which operate in extremely difficult environments. The above analysis suggests that firms should not rush in pursuit of the currently most profitable opportunities. A project involving the transfer of funds from one input to another which increases profits gradually is better than a project which does so quickly—not because the quick project is more risky, but because the gradual project is likely to lead to more opportunities for further improvement.

Second, this model works for the non-profit sector as well. As long as a firm has a well-defined objective function which depends on budgetary inputs and may be approximated by a second-order polynomial, then the above analysis is applicable. Most notably, it appears that the desirability of Least Ascent is independent of the details of the objective function. For every tested objective, Least Ascent is at least as good as any other strategy. In general and on average, it is better.

Third, a CEO need not explicitly insist that less profitable projects be funded. Instead, she can organize the firm to make this outcome more likely. For example, if decision making is decentralized, then smaller divisions will be able to devote parts of their budget to projects which are less profitable than projects found in other divisions. Decreased central control and evaluation leads to a more gradual walk through the profit landscape without the need for any individual to explicitly decide against a more profitable option.

The real allure of computer simulation is that it is most useful in precisely those contexts in which traditional economic modeling breaks down. Whether or not this model, in particular, captures any of the important underlying structure of the environment in which firms must operate is an open question. There can be little doubt, however, that models like this one provide a new method for understanding the complex problems faced by economic agents.

5 Conclusion

The fundamental assumption of this paper is that firms face problems which are difficult rather than easy. Firms do not know which combination of inputs will maximize profits. We have sought to model this search within the classical framework of a single period input allocation subject to a budget constraint. Our model assumed a profit function which is significantly more complex than those traditionally employed and is analytically intractable. Analytical intractability is a virtue, not a vice. Using the techniques of computational modeling allows us to derive conclusions which follow from the assumptions in exactly the same manner that analytically derived conclusions follow from simpler assumptions. In a discussion of bounded rationality in the context of evolutionary game theory, Mailath writes:

> Modelers do not make assumptions of bounded rationality because they believe players are stupid, but rather that players are not as sophisticated as our models generally assume. In an ideal world, modelers would study very complicated games and understand how agents who are boundedly rational in some way behave and interact. But the world is not ideal and these models are intractable. In order to better understand these issues, we need to study simple models which we can solve [10, p. 265].

Mailath mistakenly assumes that because something is analytically intractable, it is completely intractable. This is not true. Our profit landscape is designed to be as "complicated" as real world economic problems and our search strategies are designed to be as "sophisticated" as actual economic agents. It is possible, even likely, that we are off on both counts. That is, economic optimization problems may be more or less complex than our profit landscape and economic agents may be more or less efficient at searching for improvements than our computational strategies. The key point is that both the landscape and the strategies are tunable. We can make the landscape more or less complex while holding the strategies constant. We can make the strategies more or less sophisticated while holding the landscape constant. We can change both simultaneously. In all of these cases, the resulting model is computationally, albeit not analytically, tractable. Well-defined assumptions about landscapes and strategies yield precise conclusions about outcomes and comparisons. Mailath's "simple models" may be sufficient for doing economic analysis, but they are not necessary. Computational techniques extend economic analysis to traditionally intractable models.

Interesting economic problems are difficult. Employing "simple models"—models in which the analytic solution can be derived with pencil and paper—has proved to be a fruitful method for understanding these problems. Employing complex models—models in which there is no analytic solution—has become widely practicable only recently. Our hope is that models which capture, indeed embrace, the complexity of real world economic scenarios will prove at least as fruitful for research and understanding as models which assume away that complexity. If parsimony has its costs, then perhaps complexity will have its benefits.

Acknowledgments

Financial support from the Santa Fe Institute and the Olin Foundation is gratefully acknowledged. David Fogel, Gary King, Kieron Meagher, John Miller, Ken Shepsle and Richard Zeckhauser provided helpful comments on an earlier draft.

References

[1] Armen Alchian. Uncertainty, evolution and economic theory. *Journal of Political Economy*, 58:211–221, 1950.

[2] W. Brian Arthur. Competing technologies, increasing returns, and lock-in by historical small events. *Economic Journal*, 99:116–131, 1989.

[3] Peter J. Cameron. *Combinatorics: Topics, Techniques and Algorithms.* Cambridge University Press, New York, 1994.

[4] Anthony Downs. *An Economic Theory of Democracy.* Harper Collins, New York, 1957.

[5] Terry Jones. *Evolutionary Algorithms, Fitness Landscapes and Search.* PhD dissertation, The University of New Mexico, Alberquerque, New Mexico, May 1995.

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

[7] Ken Kollman, John Miller, and Scott Page. Adaptive parties in spatial elections. *American Political Science Review*, 86:929–937, 1992.

[8] Paul Krugman. Increasing returns and economic geography. *Journal of Political Economy*, 99:483–499, 1991.

[9] Charles Lindblom. The science of 'muddling through'. *Public Administration Review*, 19:79–88, 1959.

[10] George Maliath. Introduction: Symposium on evolutionary game theory. *Journal of Economic Theory*, 57:259–277, 1992.

[11] Thomas Sargent. *Bounded Rationality in Macroeconomics.* Clarendon, Oxford, 1993.

[12] Herbert A. Simon. *Administrative Behavior.* The Free Press, New York, third edition, 1976 [1947].

[13] John von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior.* Princeton University Press, Princeton, New Jersey, third edition, 1953 [1944].

[14] Sewall Wright. The roles of mutation, inbreeding, crossbreeding and selection in evolution. In *Proceedings of the Sixth International Conference of Genetics*, number 1, pages 356–366, 1932.