The Emergence of Firms in a Population of Agents: Local Increasing Returns, Unstable Nash Equilibria, And Power Law Size Distributions

Robert Axtell[†]

Center on Social and Economic Dynamics Working Paper No. 3 June 1999

© Brookings Institution 1999

⁺ Robert Axtell is a fellow at the Brookings Institution and a member of the Santa Fe Institute. This paper is dedicated to the memory of Ben Harrison, teacher and friend, for whom firms were all about adaptation and innovation, not equilibrium and stasis. For helpful conversations and comments Robert Axtell thanks Todd Allen, David Audretsch, Bob Axelrod, Per Bak, Margaret Blair, Kathleen Carley, Michael Cohen, Art DeVany, Bill Dickens, Joshua Epstein, Richard Florida, Duncan Foley, Joe Harrington, Jack Hirshleifer, Steve Klepper, Paul Kleindorfer, Axel Leijonhufvud, Roger Lewin, Blake LeBaron, Bob Litan, Francesco Luna, Igor Matutinovic, Chris Meyer, Benôit Morel, Scott Moss, Nienke Oomes, Scott Page, Teresa Parker, Kai Shih, Leigh Tesfatsion, Sid Winter, Peyton Young, and seminar participants at Brookings, Carnegie Mellon, George Mason, Harvard, Michigan, UCLA, and Wharton. Preliminary versions of this paper were presented at conferences, including INFORMS (Fall 1996), the New England Complex Systems Institute (1997), and the Society for Computational Economics (1998). Financial support from the MacArthur Foundation is gratefully acknowledged. This research was partially supported by the National Science Foundation under grant IRI-9725302.

Abstract

A model in which heterogeneous agents form firms is described and empirically tested. Each agent has preferences for both income and leisure and provides a variable input ('effort') to production. There are increasing returns to cooperation, and agents self-organize into productive teams. Within each group the output is divided into equal shares. Each agent periodically adjusts its effort level to maximize utility non-cooperatively. Agents are permitted to join other firms or start up new firms when it is welfare-improving to do so. As a firm becomes large, agents have little incentive to supply effort, since each agent's share is relatively insensitive to its effort level. This gives rise to free riders. As free riding becomes commonplace in a large firm, agents migrate to other firms and the large firm declines. It is demonstrated analytically that there exist Nash equilibrium effort levels within any group, but these are (1) Pareto-dominated by effort configurations that fail to be individually rational, and (2) dynamically unstable for sufficiently large group size. The out-of-equilibrium dynamics are studied by an agent-based Individual firms grow and perish, there is computational model. perpetual adaptation and change at the micro-level, and the composition of each firm at any instant is path-dependent. However, at the aggregate-level stationary firm size, growth rate and lifetime distributions emerge. These are compared to data on U.S. firms. In particular, the power law character of empirical firm size distributions is reproduced by the model. Log growth rates are distributed as a Laplace (double exponential) distribution, while the standard deviation in growth rates scales (decreases) with firm size, both in agreement with recent empirical analyses. Constant returns obtain at the aggregate level, in contrast to the increasing returns of the micro-level. A portrait of this agents-withinfirms economy is developed by analyzing typical firm life cycles, typical agent careers, and through cross-sectional analysis. The model Right-skewed size parameterization is systematically investigated. distributions are robust to a variety of alternative specifications of preferences, compensation, interaction structure, and bounded rationality. The role of residual claimants within firms is briefly explored. Finally, it is argued that any theory of the firm based on microeconomic equilibrium is unlikely to explain the empirical data on firm sizes, growth rates, and related aggregate regularities. In the multi-agent perspective developed here, successful firms are ones that can attract and retain productive workers, a picture very different than the conventional profit maximizing view of firms.

<u>Keywords</u>: endogenous firm formation, increasing returns, bounded rationality, unstable Nash equilibria, power law size distribution, double exponential growth rate distribution, agent-based computational model, non-equilibrium game theory, path-dependence, economic complexity, group selection

<u>JEL</u> classification codes: C63, C73, D23, L11, L22

CONTENTS

1	Introdu	action	1
	Var	ious Theories of the Firm	3
2	A Vari	able Effort Model of Firm Formation	7
	2.1	Set-Up	8
	2.2	Equilibrium	9
		<i>Example:</i> Graphical depiction of the solution space, 2 agents with θ =	0.5
		11 Singlatan Firma	10
		Singleton Firms.	1Z
		Homogeneous Groups	13
	23	Stability of Equilibrium	11 17
	2.0	Analysis for Fixed Group Size	
		<i>Example:</i> Onset of instability in a homogeneous group having $\theta = 0.7$	20
		Analysis for Variable Group Size	23
		Unstable Equilibria and Pattern Formation Far From	
		Equilibrium	23
3	Compu	Itational Implementation with Agents	25
	3.1	Set-Up of the Computational Model	26
	3.2	Object-Oriented Implementation	
	3.3	Aggregate Dynamics	31
		Number of Firms	33
		Firm Size	34
		Effort Levels	35
		Output	36
		Income	
		Utility	
	3.4	Distributions of Firm Sizes, Growth Rates and Lifetimes	39
		Firm Size Distribution	
		Froductivity Firm Crowth Bata Distribution	42
		Firm Lifetime Distribution	43
	35	Dynamics of Individual Firms: The Firm Life Cycle	40 47
	3.6	Dynamics of Individual Agents	<u>-</u> 7
	3.7	Overall Behavior of Agents I (Population Cross-Section)	51
	3.8	Overall Behavior of Agents, II (Cross-Section by Firms)	61
	0.0	Group Selection	67
	3.9	A Speculative Welfare Analysis	67
4	Variati	ons and Elaborations	69
	4.1	The Importance of Locally Purposive Behavior	71
	4.2	Effect of the Increasing Returns Parameters	71
	4.3	Alternative Specifications of Preferences	73

4.4	4.4 Effect of the Extent and Composition of Agent Social		
	Networks	74	
4.5	Effect of Agent	76	
4.6	Bounded Rationality: Groping for Better Effort Levels	76	
4.7	Alternative Compensation Schemes	78	
4.8	Firm Founder Acts as Residual Claimant	79	
4.9	Firm Founder Sets Hiring Standards	80	
4.10	Effort-Level Monitoring within Firms	81	
4.11	Putting it all Together: Compound Variations	82	
5 Conclu	sions	83	
5.1	The Emergence of Firms and "Universality"	85	
5.2	What Constitutes a 'Theory of the Firm'?	87	
5.3	The Comparative Advantages of Computational Agents	90	
Appendix.		92	
A	More General Preferences	92	
В	A More General Model of Effort Level Adjustment	93	
References	3	94	

ANIMATION

Animation 1: Snapshots from a typical firm formation process over 8 consecutive time	
periods; the agent who has been with a firm the longest (typically the firm	
founder) is colored red, others blue	\$2

FIGURES

Figure 1 : Dependence of e_i^* on $E_{\sim i}$; $a = b = 1$, $\theta_i \in \{0.50, 0.80, 0.90, 0.95\}$ 10
Figure 2 : Effort level space for two agents each having $\theta = 0.5$
Figure 3 : Optimum utility as a function of θ , parameterized by $b \in \{0, 1, 2\}, a = 1$ 13
Figure 4 : Nash equilibrium effort levels as a function of <i>θ</i> , for homogeneous groups of size <i>N</i> , with <i>N</i> varied parameterically from 1 to 1015
Figure 5 : Individual utility levels at Nash equilibrium effort levels in homogeneous groups over a range of group sizes, for various θ 15
Figure 6 : Optimal equilibrium size in homogeneous groups as a function of θ 16
Figure 7 : Optimum utility in homogeneous groups of optimal size, as a function of θ 16
Figure 8: Phase space of effort level adjustment 17
Figure 9 : Onset of instability in homogeneous groups; for each θ the Nash equilibrium effort level is unstable for sufficiently large group size
Figure 10 : Typical time series of the total number of firms (upper line) and the number of new firms (lower line)
Figure 11: Power spectrum of the time series for the total number of firms
Figure 12: Typical time series for average firm size (inset) and largest firm size
Figure 13 : Typical time series for average effort level in the population (thick line) and in the largest firm (thin line)
Figure 14: Typical time series for total output of all firms (thick line) and output of the largest firm (thin line)
Figure 15: Typical time series for average income in the population (thick line) and in the largest firm (thin line)37
Figure 16: Typical time series for average utility in the population (thick line) and in the largest firm (thin line)
Figure 17: Stationary firm size distribution (probability mass function)40
Figure 18: Stationary distribution of firm output (probability mass function)42
Figure 19 : Productivity exhibits near constant returns at the aggregate level, despite increasing returns at the micro-level
Figure 20: Stationary firm growth (probability mass function)45
Figure 21 : Dependence of the standard deviation in <i>r</i> on firm size46
Figure 22: Firm lifetime-rank distribution 47

Figure 23 : Typical evolution of firm size over the life cycle of the firm, by number of agents (thick line) and output (thin line)
Figure 24 : Typical evolution of firm composition by agent type over the firm life cycle48
Figure 25: Typical evolution of effort levels in a firm over the life cycle49
Figure 26 : Typical evolution of the number of free riders in a firm over the life cycle49
Figure 27 : Typical evolution of agent income by agent type over the firm life cycle50
Figure 28 : Typical evolution of agent utility by agent type over the firm life cycle51
Figure 29 : Career path over 30 periods' of an agent for whom $\theta = 0.7$
Figure 30 : Effort level history over 30 periods of an agent for whom $\theta = 0.7$
Figure 31 : Income history over 30 period of an agent for whom $\theta = 0.7$
Figure 32 : Utility history over 30 periods of an agent for whom $\theta = 0.7$
Figure 33: Effort level histogram
Figure 34: Effort levels by agent type; averages ± 1 standard deviation
Figure 35: Distribution of income in the agent population 57
Figure 36 : Income by agent type; averages ± 1 standard deviation57
Figure 37: Utility histogram
Figure 38: Utility by agent type; averages ± 1 standard deviation
Figure 39 : Decision utility by agent type; averages ± 1 standard deviation60
Figure 40 : Tenure by agent type; averages ± 1 standard deviation60
Figure 41 : Size of an agent's firm, by agent type; averages ± 1 standard deviation61
Figure 42 : Histogram of average θ in firms
Figure 43 : Histogram of standard deviation of θ in firms
Figure 44: Distribution of average effort levels in firms
Figure 45: Histogram of standard deviation of effort levels in firms
Figure 46: Distribution of income across firms 65
Figure 47: Distribution of average utility in firms
Figure 48: Distribution of standard deviation in utility in firms 66
Figure 49 : Utility derived from working in single agent firms, in optimal size homogeneous firms, and in non-equilibrium multi-agent firms, by agent preference

PSEUDO-CODE

Pseudo-code block 1: Agent object	29
Pseudo-code block 2: Firm object	30
Pseudo-code block 3: Pseudo-code for the model overall	

TABLES

Table 1 : Onset of instability in a group of agents having $\theta = 0.7$; Nash equilibrium effort levels in groups beyond size 6 are unstable
Table 2: 'Base case' configuration of the computational model
Table 3 : Dependence of the scaling law exponent, μ , on the increasing returns exponent, β 72
Table 4 : Dependence of the scaling law exponent, μ , on the coefficient <i>b</i> , of the increasing returns term
Table 5 : Dependence of the scaling law exponent, μ , on the distribution of agentpreferences
Table 6 : Dependence of scaling law exponent on the size of social networks, where social networks are composed of agents (friends), specified exogenously
Table 7 : Dependence of scaling law exponent on the size of social networks, where social networks are ccomposed of randomly selected firms
Table 8: Dependence of scaling law exponent on agent loyalty
Table 9 : Dependence of scaling law exponent on stickiness of effort adjustment, various β 77
Table 10 : Dependence of scaling law exponent on stickiness of effort adjustment, various β 77
Table 11: Dependence of scaling law exponent on hyperbolic compensation functions
Table 12 : Dependence of scaling law exponent on type of base compensation
Table 13: Dependence of scaling law exponent on firm target output level

1 Introduction

This essay describes a model in which firms self-organize within a heterogeneous population of boundedly rational agents who interact locally. In the terminology of Epstein and Axtell [1996], firms "grow" from the bottom-up. In particular, we study a team production environment characterized by increasing returns and various rules for dividing group output. Placing agents who are local utility-maximizers in such an environment proves sufficient for the emergence of multi-agent groups. Such groups, it will be seen, have characteristics that are suggestive of firms.

There are two primary motivations for building a model in which firms emerge. First, to the extent that the received theory of the firm is concerned with why and how firms happen,¹ a model in which firms emerge can be used to test the extant theory, and to assess its generality by relaxing assumptions of the theory. For example, do theoretical claims concerning what constitutes the core elements of firms, such as economizing on transaction costs, represent statements about necessity or sufficiency? What happens as rationality assumptions are progressively relaxed? As complete rationality gives way to bounded rationality and finally to mere purposiveness, at what point are firm-like multi-agent organizations no longer observed? Of crucial importance, are the various 'theories of the firm' even sufficiently well specified that one can build more or less complete microeconomic models of them?² A living model, in which some firms grow and prosper while others do not, can function as a laboratory in which systematic experiments, designed to test some aspect of the theory, are conducted. At the very least such experiments will shed light on the strengths and weaknesses of the extant theory. However, a working model could do much more, perhaps serving as a catalyst for new conceptualizations of the firm.

The second motivation for building a model of firm formation is to contribute to the development of the methodology of agent-based computational modeling. Creating agent models requires explicit specification of how agents interact at the micro level. The resulting model is then spun forward in time and one looks for pattern and structure to emerge from the interactions of the agents. Today, much is known about how to get markets to emerge in artificial agent models (cf. Arthur *et al.* [1994], Epstein and Axtell [1996], Chen and Yeh [1997], Kirman and Vriend [1998]). Too, computational organization theory has progressed as a modeling discipline by

¹ Arguably, the theory of the firm, such as it exists today, is predominantly concerned with intra-firm organization, and less about the origin of firms, although these matters are not unrelated.

² Models without explicit dynamics constitute, at best, an incomplete explanation.

taking organizational forms as given (cf. Prietula, Carley and Gasser [1998]). However, little is known about how to get multi-agent organizations to form and evolve endogenously.³ Operational multi-agent firms are a necessary step along the road to the creation of a full-blown agent-based economy in software.⁴

These two motivations—to test the economic theory of the firm and to add to the methodology of agent-based computational modeling—are intrinsically related. On the one hand, any computational model of firm formation must have economic foundations. On the other hand, we shall see that the out-of-equilibrium character of the economic model is efficaciously studied through the agent-based computational approach

But how will we know when we have succeeded? That is, by what criterion might we evaluate the performance of one model of firm formation against another? Axtell and Epstein [1994] describe a variety of ways to assess the performance of agent-based models. Our approach here will be to compare model output to *aggregate* statistical data. In particular, the empirical firm size distribution—well known to be right skewed—is shown to be similar to the size distribution that emerges in the population of firms of the computational model. Further, the model yields firm growth rate statistics that are closely related to recent empirical analyses (Stanley *et al.* [1996]).

Ostensibly, the best current explanations for the observed firm size distribution are, in essence, phenomenological in nature since they are written in terms of aggregate variables. Beginning with Gibrat⁵ and continuing in the efforts of Simon and co-workers (Ijiri and Simon [1977]), as well as others, there exists a body of stochastic process models in which random draws from a symmetric distribution of growth rates yield distributions of firm sizes that are right skewed, following a lognormal, Pareto, or Yule distribution, depending on the exact structure of the process.

³ Padgett [1998] has modeled the formation of networks of complimentary skills within a population and describes the local clustering of particular skills as a simple kind of economic organization. Luna [1998] investigates problem solving by heterogeneous teams of neural networks and interprets the results in terms of firms. Axelrod [1995] and his student, Cederman [1997] have each built models of state formation and evolution that are similar in spirit to the model of firms described here, involving heterogeneous agents having limited information who engage in purposive behavior through local interaction.

⁴ Building an entire artificial economy from agents is an active research program, first described by Lane [1993]; see also Lewin [1997]. An early effort in this direction is Basu and Pryor [1997], although firms are unitary actors in their model, not multi-agent organizations.

⁵ For reviews of Gibrat's contributions see Steindl [1965] or Sutton [1997].

Apparently, there does not exist today a purely microeconomic explanation for the overall firm size distribution.⁶ The inability of the neoclassical theory of the firm—with its U-shaped cost functions and perfectly informed and rational managers—to render a plausible explanation of the empirical size distribution has been caustically critiqued by Simon [Ijiri and Simon, 1977: 7-11, 138-40].⁷ Nor do the transaction cost, principal-agent, and other more recent theories of industrial organization seem to place significant restrictions on firm size and growth distributions. These varied approaches to understanding firms are briefly reviewed in the next section.

Various Theories of the Firm

The *neoclassical theory* of the firm distills the multi-agent character of real firms down to a single rational actor, who faces completely specified technological options, and who acts to maximize profit by choosing inputs that minimize cost. There are many problems with this firm-as-production-function picture, from its static technology to its hyper-rationality. But, as Winter [1993: 181-2] has carefully argued, it is also something of a methodological curiosity. Throughout economics the common principle of explanation involves *methodological individualism*, from the theory of consumer behavior to general equilibrium theory and beyond. But where are the individual firm members in the neoclassical cost minimization model of the firm? More generally, the neoclassical model seems to have little to say about intra-firm organization. It is perfectly consistent with an economic configuration in which the economy is one giant firm and has as its divisions and plants what today we call firms. That is, the neoclassical theory of the firm fails to explain the boundary between market and firm.

Transaction cost theories, together with incomplete contracting approaches, constitute perhaps the dominant paradigm in the modern theory of the firm. Originating with Coase [1937] and continuing with Williamson [1975, 1985], this approach explicitly recognizes the connection between organization and cost, and argues that different ways of organizing transactions leads to systematically different cost structures. The incomplete

⁶ Sutton [1998] is an ambitious exception to this statement. He has developed a simple game theoretic model of firm competition that determines bounds on the extent of intra-industry concentration, thus constraining the shape of the size distribution. A variety of other recent research, also written at the firm level, addresses the origins of the size distribution. Kwasnicki [1998] obtains right-skewed firm sizes in an evolutionary model of technological change. Mazzucato [1998] explores the evolution of market concentration and market share instability using a model in which firms experience falling costs. Aslund [1998] describes a game theoretic model of competition between firms and finds that firm sizes within industries should be similar. He then goes on to advance the seemingly novel hypothesis that right-skewed firm size distributions result from skewed distributions of market size.

⁷ For a more recent statement of these views see his Mattioli lectures (Simon [1997b]).

contracts model further argues that it is prohibitively expensive to write comprehensive contracts, so transactions are intrinsically costly. This leads to the view of the firm as a nexus of contracts. While ostensibly quite broad in scope, the operational content of the transactions cost approach has not always been clear, the extent to which it views agents as rational and the firm as profit maximizing is ambiguous, and it seems to have little to say about concentration, i.e., relative firm sizes.

Partly in recognition of the defects in the textbook orthodoxy, a body of results on *principal-agent theory* has been developed in the context of the theory of the firm. This work relates, principally, to optimal incentive systems within firms and optimal arm's-length contracts between firms. Results have been derived for cases involving multiple principals, multiple agents, reputation effects and so on. However, this literature too, fails to shed much light on either intra-firm organization or firm-market boundaries [Hart 1995: 20].

A rather different theory of the firm is the coalitional or general equilibrium view. Here, agents are treated as heterogeneous, each with unique preferences and abilities. A firm then becomes a (stable) coalition of such agents; see, for example, Kihlstrom and Lafont [1979], Kleindorfer and Sertel [1979, 1982], Lucas [1978], and Laussel and LeBreton [1995]. The formation of such firms can be considered endogenous (Hart and Kurz [1983]), although the process by which such coalitions might assemble themselves is largely unspecified since this general equilibrium conception of a firm is a completely static notion. When it comes to theorizing the dynamics of coalition formation a variety of problems are encountered. While it is possible to describe certain aspects of a dynamical theory formally (Roth [1984]), for even small agent populations the number of coalition structures is so vast that it is simply not feasible that any significant subset of them could ever be sampled. Thus, it is not credible that a particular set of firms represents anything like an optimal coalition structure (De Vany [1993a, 1993b, 1993c, 1996]).8

Related to the coalitional approach to firms is the recent flowering of the economics of information processing within organizations. Here the firm is viewed as an information processing organization with some topology of interaction among the agents (Radner [1993], Van Zandt [1996, forthcoming], Miller [1996], DeCanio and Watkins [1998]). The comparative efficiency of organizations having different topologies and alternative incentive structures is the primary object of study in this sub-field. The related view of

⁸ Recent work on the computational complexity of bounding optimal coalition structures can be viewed as impossibility results for the efficient determination of tight bounds (see Sandholm *et al.* [1998], Shehory and Kraus [1993], and Klusch and Shehory [1996a, 1996b]).

Finally, the firm plays a central role within the broad field of *evolutionary economics*. Here the firm is viewed neither as a production function, nor as a nexus of contracts, but as an amalgam of operating rules and heuristics (Nelson and Winter [1982]). Instead of interpreting firm structure and function to be the result of some optimization program, the evolutionary approach suggests that founder effects and path dependence are determinate. Firm behavior is modeled as incrementalist and profit seeking instead of profit maximizing. But the evolutionary theory of the firm as it stands today, as with the transactions cost approach, seems to come up short operationally. That is, its empirical relevance remains largely to be developed.

Overall, and taken together, these varied theoretical perspectives have much to offer the student of the firm. They present frameworks for building models, as well as a vocabulary in which alternative hypotheses can be expressed. These theories seem to explain certain stylized facts, and it is possible to interpret some of them empirically, both through case studies as well as econometrically (cf., Joskow [1993]). However, to greater or lesser degrees, all these approaches view the firm as a reasonably small group of more or less homogeneous agents who behave rationally.⁹ Furthermore, an implicit assumption in most of these approaches is that equilibrium behavior constitutes the observed behavior of firms, and thus comparative statics is an appropriate tool for studying alternative form, function and fitness of firms.

This paper describes a model that draws together various threads from these competing theoretical literatures. From the neoclassical tradition the notion of a production function is preserved, albeit in a somewhat modified form. The model is written at the level of individual agents, and incentive problems and related ideas from the principal-agent literature manifest themselves in important ways. The agents in the model work in perpetually novel environments, in which contracts are of necessity incomplete, and so transaction costs are intrinsic. Each firm at each instant will be composed of a coalition of agents, so the general equilibrium approach is relevant. Finally, the way in which agents make decisions, and the way firms grow and decline, is very much in the spirit of evolutionary economics.

Although the model described below is situated conceptually within existing theories of the firm, the main results are developed using a methodology that is largely unfamiliar to many economists—so-called agentbased computational modeling. In agent-based computational models a

⁹ The coalition formation approach relaxes the small group assumption, while rationality is commonly considered bounded in the evolutionary approach and, perhaps to a lesser degree, in transactions cost theorizing.

population of data structures representing individual agents is instantiated and permitted to interact. One then looks for systematic regularities, often at the macro-level, to emerge from the interactions of the agents. The shorthand for this is that macroscopic regularities "grow" from the bottomup. No equations governing the macro social structure are solved in multiagent computational modeling. Typically, the only equations present are those used by individual agents for decision-making. Nor are agents assumed to have complete information that they can costlessly process. Instead, agents glean data concerning their economic environment from members of their social networks, that is, through local interactions. This relatively new methodology facilitates modeling agent heterogeneity, non-equilibrium dynamics, local interactions (network/spatial processes), and boundedly rational behavior (cf. Epstein and Axtell [1996]: Chapter I).

The model of firm formation elaborated below consists of a population of heterogeneous agents who have preferences for income (derived from work) and leisure (all time not spent working). There are increasing returns to cooperation, so agents who work together can produce more output per unit of effort than if they work alone. However, agents act non-They continually select effort levels that maximize cooperatively.¹⁰ individual welfare, and may migrate between firms or start-up new firms. Thus, entry decisions are endogenous. Firm output is divided equally among the workers. Firms form in the model due to the increasing returns, but, since agents are constantly adjusting their effort levels, large firms are not stable. This is because once a firm becomes large each agent's share is only weakly related to its effort level, and so free-riding sets in. Agents eventually move out of firms 'infected' with free riders. Exit decisions are, therefore, also endogenous. It is demonstrated analytically that there do not exist stable equilibria in this environment. Furthermore, it is argued that the nonequilibrium regime provides greater welfare for the agents than would equilibrium even if it were stable. An agent-based computational model is used to study the non-equilibrium dynamics, in which firms are perpetually born, growing and perishing. After an initial transient period there results stationary distributions of firm size (by both number of employees and output) and growth rate. Firm size follows a scaling (power law) distribution, in accord with empirical data.¹¹ In fact, for certain parameterizations the power law exponent estimated from the model 'data' is similar to that for U.S. firms. In contrast to increasing returns at the firm level constant returns

¹⁰ For a cooperative game theoretic view of firms see Ichiishi [1993].

¹¹ As this draft was being completed the paper of Laherrère and Sornette [1998], on so-called 'stretched exponential distributions' (aka Weibull distributions) has come to my attention. Preliminary results of fitting this functional form to firm data indicate that it is somewhat superior to a pure power law. It is anticipated that the implications of this will be developed in a subsequent draft.

obtain at the macro-level. The computational model generates empirically testable patterns and regularities, about which there seem to be very little data, such as the distribution of firm lifetimes. Finally, there is a sense in which the model supports the idea that intra-firm cooperation *between* agents is a by-product of inter-firm competition *for* agents.

In the next section (§ 2) some analytical results are obtained for a representative agent/representative firm version of the model. Then, in § 3, the computational model is described and its output analyzed. Section 4 describes some extensions of the basic model. Finally, § 5 summarizes the main findings and draws conclusions.

2 A Variable Effort Model of Firm Formation

In his famous paper on firms, Coase [1937] wondered why there were firms at all, and not merely market relations between individuals. Thinking of firms as "islands of cooperation" (Robertson [1930]), and markets as institutions for competition (between individuals, between firms, and between individuals and firms), a different way to phrase Coase's query is 'What are the boundaries between cooperation and competition?'

An important answer to this question is Alchian and Demsetz [1972]. These authors argued that firms are vehicles for managing incentive problems in team production, and that the production environment determines the most efficient organizational form. Holmström [1982] subsequently formalized certain of these ideas.

Recently, a variety of models of individual behavior within groups have appeared in which the tension between cooperation and competition is explicitly considered, and the resulting group dynamics are studied analytically. The general character of the model described below—variable agent effort leading to more or less cooperation and fluctuating sizes of groups—is a variation on the formulation of Canning [1995], which is quite similar to Huberman and Glance [1998].¹² It extends these analyses by focusing not merely on the unstable equilibria of group formation (Canning) or on fluctuations about equilibrium (Huberman and Glance). Rather, recourse is made to agent-based computational modeling, which yields a complete dynamical picture of the solution space. But first, the structure of the model is described analytically.

¹² See also Glance *et al.* [1997].

2.1 Set-Up

There is a finite, fixed set of agents, A, each of whom works with some effort level $e_{i \in A} \in [0, 1]$. Consider a representative group composed of N agents. The total effort level in the group is simply

$$E = \sum_{i=1}^{N} e_i . \tag{1}$$

The group produces output, *O*, as a function of *E*, according to

$$O(E) = aE + bE^2.$$
⁽²⁾

This represents the group's production function.¹³ The case of b = 0 corresponds to constant returns to cooperation, while b > 0 amounts to increasing returns.¹⁴ The increasing returns to production means, essentially, that agents working together can produce more than they can as individuals.¹⁵ To see this, consider two agents having effort levels e_1 and e_2 . As individuals they produce total output $O_1 + O_2 = a(e_1 + e_2) + b(e_1^2 + e_2^2)$, while working together they make $a(e_1 + e_2) + b(e_1 + e_2)^2$. Clearly this latter quantity is at least as large as the former since $(e_1 + e_2)^2 \ge e_1^2 + e_2^2$.¹⁶

The agents in a group share total output equally. That is, at the end of each period each agent receives an O/N share of the total output.¹⁷

¹³ This terminology is at least somewhat problematical. While (2) relates inputs to outputs, in concert with the neoclassical view of a production function, the inputs are not explicit choice variables of a decision-maker within the firm, since they are determined by the individual agents who make up the firm. Therefore, there is no sense in which the production function given by (2) can be made the object of a mathematical programming problem, to minimize costs, for example, as in conventional production theory. However, since (2) describes the effect on firm output of agent behavioral adjustments, as well as the effect of changing firm composition, there is a definite sense in which it describes production possibilities, albeit dynamic ones, and so we call it a production function.

¹⁴ The notion of increasing returns at the firm level goes at least back to Marshall [1920], and proved to be fodder for various theoretical controversies in the 1920s (Sraffa [1926], Young [1928]). More recent work on increasing returns is reprinted in Arthur [1994] and Buchanan and Yoon [1994]. Colander and Landreth [1999] give an interesting account of the history of the idea.

¹⁵ There are many ways of justifying increasing returns between individuals, such as when agents are solving so-called 'four hands problems.' These will not be pursued here.

¹⁶ Note that the increasing returns character of (2) is preserved if a = 0. Although setting a to 0 would simplify somewhat the analytical results that follow, it will subsequently prove insightful to compare the results for increasing returns with those that obtain in the case of constant returns; this can only be accomplished with $a \neq 0$.

¹⁷ It will eventually be demonstrated that the model yields a more or less constant amount of total output, as well as a stationary distribution of income. Thus, in a competitive market the price of the output will be more or less constant. Since there are no fixed costs, the output shares sum to total cost which equals total revenue, and therefore profit equals zero. The income shares can be thought of as either uniform wages in pure competition or equal

Each agent has Cobb-Douglas preferences for income and leisure.¹⁸ All time not spent working is spent in leisure, thus agent *i*'s utility can be written as a function of its effort level, e_i , as

$$U^{i}\left(e_{i};\theta_{i},E_{\sim i},N\right) = \left(\frac{O\left(e_{i};E_{\sim i}\right)}{N}\right)^{\theta_{i}}\left(1-e_{i}\right)^{1-\theta_{i}},$$
(3)

Note that in this expression the group output, *O*, has been written as a function of e_i , with the remainder of the total group effort, $E_{\sim i}$, considered a parameter; $E = e_i + E_{\sim i}$.

2.2 Equilibrium

Let us say that each agent knows its preferences, θ_i , the size of the group to which it belongs, N, and the output of the group, O, from which it can determine E and thus $E_{\sim i}$. Individual effort is not observable. Furthermore, each agent, i, selects the effort level, e_i^* , that maximizes its utility; formally,

$$e_i^* \equiv \arg \max_{e_i} \left[U^i \left(e_i; \theta_i, E_{\sim i}, N \right) \right].$$
(4)

It is straightforward, albeit somewhat tedious, to show that the solution to equation 4 is given by

$$e_{i}^{*}(\theta_{i}, E_{\sim i}) = \max\left[0, \frac{-a - 2b(E_{\sim i} - \theta_{i}) + \sqrt{a^{2} + 4ab\theta_{i}^{2}(1 + E_{\sim i}) + 4b^{2}\theta_{i}^{2}(1 + E_{\sim i})^{2}}}{2b(1 + \theta_{i})}\right]$$
(5)

for b > 0, and

$$e_i^*(\theta_i, E_{\sim i}) = \max\left[0, \theta_i - E_{\sim i}(1 - \theta_i)\right]$$
(6)

in the case of constant returns (i.e., b = 0).¹⁹ Note that these results do not depend explicitly on *N*, the size of the group. However, they do depend on

profit shares in a partnership. Alternative agent compensation policies are studied in § 4.7 below.

¹⁸ In the appendix a more general model of preferences is specified. All the main results of this paper obtain in the general case.

¹⁹ In the case of constant returns it is never strictly individually-rational for agents to work together when the option of working alone exists. To see this, note that if in a group of *identical* agents each works at the same level it would if working alone then the output share each receives is identical to what each agent would produce on its own; indeed, this is the definition of constant returns. Thus there are no welfare advantages to cooperation. In a group of *heterogeneous* agents those with relatively smaller preference for income work less than those who prefer income to leisure. Therefore, under the equal share rule for dividing output, participation in such a group is not incentive compatible for agents with relatively larger preference for income. Results for constant returns will be shown for

 $E_{\sim i}$ —the amount of effort put in by the other agents. In order to develop some intuition for the general dependence of e_i^* on its parameters, we plot it for a = b = 1 in figure 1 below, as a function of $E_{\sim i}$, for several values of θ_i .



Figure 1: Dependence of e_i^* on $E_{\sim i}$, a = b = 1, $\theta_i \in \{0.50, 0.80, 0.90, 0.95\}$

Note that the optimal effort level decreases monotonically as 'other agent effort,' $E_{\sim i}$, increases, and that for each $\theta_i < 1$, there exists some maximum amount of $E_{\sim i}$ beyond which it is rational for agent *i* to put in no effort. Stated differently, for a fixed $E_{\sim i}$, there exists a cut-off value of θ , call it θ_c , such that for $\theta_i \leq \theta_c$, $e_i^* = 0$, while for $\theta_i > \theta_c$, $e_i^* > 0$. That is, for agents who face the same $E_{\sim i}$, the only agents who put in positive effort are those having preference for income above θ_c , It is possible to develop an expression for θ_c by setting the second argument on the RHS of (5) equal to 0. This yields

$$\theta_{c} = \frac{a + bE_{\sim i}}{a(1 + 1/E_{\sim i}) + b(2 + E_{\sim i})}$$

As $E_{\sim i}$ gets large θ_c approaches 1, meaning that the only agents who put in non-zero effort under such circumstances are those who do not care for leisure. Note also from figure 1 that each of the curves is nearly linear over the entire range of feasible solutions. In the case of constant returns, (6) says that each is exactly linear, with slope $1 - \theta_i$.

Equilibrium in a group corresponds to all agents computing e_i^* from equation 5, using E_{-i}^* in place of E_{-i} such that

purposes of completeness, and to compare and contrast with results from the increasing returns case.

$$E^*_{\scriptstyle \sim i} = \sum_{j \neq i} e^*_j \; .$$

From the continuity of the RHSs of (5) and (6) and the compactness of the space of effort levels, it is clear that solutions to this set of equations exist. Further, given that the RHSs are monotone decreasing (see figure 1) the solution is unique. Such an equilibrium configuration is a Nash equilibrium, since once it is established no agent can make itself better off by working at some other effort level.

However, this Nash equilibrium in effort levels is not efficient. In general, there exists a continuous set of agent effort levels that Pareto dominate the Nash equilibrium, as well as a subset—also having cardinality of the continuum—that are Pareto optimal. These solutions all (a) involve larger amounts of effort than the Nash equilibrium, and (b) are not individually rational. To see (a) note that

$$dU^{i}\left(e_{i}^{*};\theta_{i},E_{\sim i}^{*},N\right) = \frac{\partial U^{i}}{\partial e_{i}}de_{i} + \frac{\partial U^{i}}{\partial E_{\sim i}}dE_{\sim i} > 0$$

since the first term on the RHS vanishes at the Nash equilibrium and

$$\frac{\partial U^{i}}{\partial E_{\sim i}} = \frac{\theta_{i} \left[a + 2b(e_{i} + E_{\sim i}) \right] \left(1 - e_{i} \right)^{1 - \theta_{i}}}{N^{\theta_{i}} \left[a(e_{i} + E_{\sim i}) + b(e_{i} + E_{\sim i})^{2} \right]^{1 - \theta_{i}}} > 0 \,.$$

Part (b) is true as a result of the fact that each agent's utility is monotone increasing on the interval $[0, e_i^*)$, and monotone decreasing on $(e_i^*, 1]$. Therefore,

$$\frac{\partial U^i}{\partial e_i} < 0 \forall e_i > e_i^*, E_{\sim i} > E_{\sim i}^*.$$

This region of effort levels that Pareto dominate the Nash equilibrium is the space in which firms live.

Example: Graphical depiction of the solution space, 2 agents with $\theta = 0.5$

Consider two identical agents having preference for income of 0.5. Solution of equation 5, with a = b = 1 yields $e^* = 0.4215$, corresponding to utility levels of 0.6704. Effort level deviations by either agent alone are Pareto dominated by the Nash equilibrium. For example, decreasing the first agent's effort to $e_1 = 0.4000$, with e_2 at the Nash level of 0.4215, yields utility levels of 0.6700 and 0.6579, respectively, generating welfare losses for both agents. An effort level perturbation in the opposite direction, to $e_1 = 0.4400$ with e_2 remaining at 0.4215, produces utility levels of 0.6701 and 0.6811, respectively, a loss for the first agent while the second agent gains from the first's additional effort. If both agents decrease their effort levels the utility of each falls, while joint increases in effort are welfare-improving for both. There exists a symmetric Pareto optimal solution in the case of identical agents. For the problem at hand this amounts to each putting in effort level 0.608 and receiving utility of 0.7267. However, no solutions involving more effort than the Nash equilibrium are individually rational: from any of these solutions each agent gains utility by reducing its effort.

All of this is depicted in Figure 2, below. It is a plot of iso-utility contours as a function of effort levels for two identical agents having $\theta = 0.5$. The lines that are 'U' shaped with respect to the page refer to the first agent, with utility increasing to the right. The 'C' shaped curves correspond to the second agent, with utility increasing up the page. The point labeled 'N' is the Nash equilibrium. The 'core' shaped region extending above and to the right of 'N' is the set of effort levels that Pareto dominate the Nash equilibrium. The set of effort levels on the curve from 'P' to 'P' are Pareto optimal.



Figure 2: Effort level space for two agents each having θ = 0.5; thin lines are iso-utility contours, 'N' corresponds to the Nash equilibrium, and the heavy line from P-P corresponds to Pareto optimal solutions

For two agents with distinct preferences for income the qualitative structure of the solution space shown in figure 2 is preserved, but the symmetry is lost. Essentially, increasing returns insures the existence of solutions that Pareto dominate the Nash equilibrium.

For more than two agents the Nash equilibrium and Pareto optimal solutions continue to be distinct. For N = 3, figure 2 can be thought of as the $e_3 = 0$ solution space. Then, because e^* is decreasing in $E_{\sim i}$, for $e_3 > 0$ the effort levels of agents 1 and 2 that correspond to the Nash and Pareto optimal solutions are lower than in the $e_3 = 0$ case.

Singleton Firms

The $E_{\sim i} = 0$ solution of (5) or (6) corresponds to agents working alone in single agent firms. For this case the expression for the optimal effort level can be written as

$$e_{i}^{*}(\theta_{i},0) = \frac{-a + 2b\theta_{i} + \sqrt{a^{2} + 4b\theta_{i}^{2}(a+b)}}{2b(1+\theta_{i})}.$$
(7)

for b > 0 and, in the case of constant returns, $e_i^*(\theta_i, 0) = \theta_i$. Note that in the limit of $\theta_i = 0$, equation 7 gives $e_i^* = 0$, while for $\theta_i = 1$ we have $e_i^* = 1$. That is,

these two limiting cases of (7) correspond to constant returns. For $\theta_i \in (0, 1)$ it can be shown that the optimal effort level with increasing returns is always greater than that corresponding to constant returns, for the same value of *a*.

Expressions for the utility corresponding to the optimal effort levels can be written down, but these are unwieldy. Attempting to learn how the utility function depends on its parameters *a* and *b* yields, upon differentiation, even less useful results due to the large number of terms involved. However, the general shape of the optimal utility as a function of θ can readily be discerned by simply plotting it numerically. Furthermore, its dependence on *b* is established by making multiple plots. These are shown in figure 3.



Figure 3: Optimum utility as a function of θ , parameterized by $b \in \{0, 1, 2\}, a = 1$

Note that for constant returns (b = 0), agents with extreme preferences have the highest utility levels when working alone. Then, as increasing returns become progressively more prominent, higher and higher utility levels accrue to agents who prefer income to leisure. Also note that this figure makes clear that utility is monotone increasing in b. It turns out that it is also an increasing function of a, in accord with intuition.

Example: Nash equilibrium with free entry and exit

Four agents having preferences for income {0.6, 0.7, 0.8, 0.9} work together in a group in which a = b = 1. Equilibrium in such a group, from equation 5, corresponds to agents working with effort levels {0.15, 0.45, 0.68, 0.86}, respectively, producing 6.74 units of output. The corresponding utilities of the 4 agents are {1.28, 1.20, 1.21, 1.32}, respectively. If each of these agents were to leave the group to found their own firm they would, according to equation 8, put in effort levels {0.68, 0.77, 0.85, 0.93}, respectively, generating outputs of {1.14, 1.36, 1.58, 1.80} and total output equal to 6.07. Their respective utility levels would be {0.69, 0.80, 0.98, 1.30}.

That is, working together they make more output, with each agent putting in less effort and receiving greater reward. This is the essence of team production.

Now say that an agent with income preference of 0.75 joins the group. The 4 original group members now adjust their effort levels to {0.05, 0.39, 0.64, 0.84}—i.e., all work less—while total output rises to 8.41. The utility levels of the original agents become {1.34, 1.24, 1.23, 1.33}, respectively, meaning that all members benefit from the new arrival. This new agent has an equilibrium effort level of 0.52 and utility level 1.23. It is individually rational for the newest agent to join since the utility it gets working alone is just 0.88.

Next, imagine that another agent having preference for income of 0.75 joins the group. The new Nash equilibrium effort levels among the original 4 group members is then {0.00, 0.33, 0.61, 0.83}, while the two new agents each put in effort of 0.48. The total output rises to 10.09. The corresponding utility levels are {1.37, 1.28, 1.26, 1.34} for the original agents and 1.26 for each of the two agents having $\theta = 0.75$. Overall, even though the addition of this sixth agent causes one of the first agents to free ride—that is, put in no effort—the net effect of this agent on the group is welfare-improving for all.

Finally, imagine that an agent having income preference of 0.55 arrives in the group. Such an agent will engage in free-riding and so will not effect the total effort or output levels, thus the individual effort levels of the extant group members will not change. However, since the output must be shared by one additional agent all utility levels fall. For the 4 original agents the new utility levels become {1.25, 1.15, 1.11, 1.17}. For the two agents having $\theta = 0.75$, their utility falls to 1.12. Overall, the addition of this last agent reduces the welfare of all. Furthermore, it lowers the utility of the $\theta = 0.9$ agent below what it can obtain working alone (1.17 versus 1.30). If agents may exit the group freely, this agent would find it rational to do so, causing all agents to readjust their effort levels. In the new equilibrium the three remaining original agents would now work with efforts {0.10, 0.42, 0.66}, respectively, while the agents having $\theta = 0.75$ would put in effort of 0.55. The newest agent would free-ride. The new output level would be 7.52, yielding utility of {1.10, 0.99, 0.96} for the original three, 0.97 for the $\theta =$ 0.75 agents, and 1.13 for the free-riding agent. Unfortunately for the group, the $\theta = 0.8$ agent now finds that it too can do better—utilities of 0.96 versus 0.98—by leaving the group to work alone (or joining with the θ = 0.9 agent). This induces a further re-equilibration of the remaining group, so that the original two members work with effort levels of 0.21 and 0.49, respectively, the two $\theta = 0.75$ agents put in effort equal to 0.61, and the $\theta = 0.55$ agent rises out of free-ridership to work at the 0.04 level. The total output drops to 5.80. The utilities of the original two are now 0.99 and 0.90, respectively, 0.88 for the $\theta = 0.75$ agents, and 1.07 for the newest agent. In this equilibrium the $\theta = 0.75$ agents are essentially indifferent between staying in the group and moving off to work alone.

Homogeneous Groups

It is informative to consider a group composed of agents all of a single type, that is, all having the same preference for income. In a homogeneous group each agent works with the same effort in equilibrium. This can be determined from (5) above, by substituting $(N-1)e_i^*$ for $E_{\sim i}$, and solving for e_i^* . Doing this yields the following expression for e_i^* :

$$\frac{-a\theta_{i}+N(2b\theta_{i}-a(1-\theta_{i}))+\sqrt{4abN\theta_{i}(2\theta_{i}+N(1-\theta_{i}))+(a\theta_{i}-N(2b\theta_{i}-a(1-\theta_{i})))^{2}}}{2bN(2\theta_{i}+N(1-\theta_{i}))}$$
(8)

for b > 0. For a homogeneous group of type θ , the Nash equilibrium effort levels can be computed directly from (8). These are shown in figure 4 as a function of θ , with a = b = 1 and group sizes, N, varied from 1 to 10.



Figure 4: Nash equilibrium effort levels as a function of θ , for homogeneous groups of size *N*, with *N* varied parameterically from 1 to 10

The corresponding utility levels are then obtained from (3). Utility levels are plotted in figure 5 below for $\theta \in \{0.5, 0.7, 0.8, 0.9\}$, with *N* varying from 1 to 25.



Figure 5: Individual utility levels at Nash equilibrium effort levels in homogeneous groups over a range of group sizes, for various θ

Note that each curve is single-peaked so there is an optimal group size for every θ . Optimal sizes are displayed in figure 6 below as a function of θ , in a semi-log plot.



rigare o. Optimal equilibrium size in nontogeneous groups as a function of o

Optimal group sizes are relatively small—less than 10—for agents having θ < 0.85, then rise quickly for agents having larger θ . The utility levels corresponding to these optimal sizes are shown in figure 7, also in semi-log coordinates. The singleton utility levels from figure 3 (b = 1) have been superimposed for purposes of comparison.



Figure 7: Optimum utility in homogeneous groups of optimal size, as a function of θ

Comparing figure 7 with figure 3 one sees that the gains to be had from participation in a homogeneous group are much more pronounced for agents

having high θ . Figures such as these will eventually serve as the basis for a welfare analysis of this model (see § 3.8).²⁰

2.3 Stability of Equilibrium

While a unique Nash equilibrium always exists in this model, it is not true that this equilibrium is always stable dynamically. That is, while the definition of Nash equilibrium guarantees that no effort level deviations are individually welfare improving at the agent level, such an equilibrium may or may not represent a stable outcome at the population level. Indeed, it will presently be shown that for a sufficiently large group the Nash equilibrium described above is unstable.

There is a simple intuition for this result. In a group of agents out of Nash equilibrium, each agent will adjust its effort level each period. As long as the adjustment functions are decreasing in other agent effort then one might expect the Nash equilibrium efforts to establish themselves over time. Because the aggregate effort level is a linear combination (simple sum) of the individual efforts, the adjustment dynamics can be conceived of in aggregate terms. In particular, the total effort level at time t + 1, E(t+1), is a decreasing function of the current period's total effort, E(t), as depicted notionally in figure 8 for a 5 agent firm. Indeed, from figure 1 it is clear that individual effort levels are nearly linear in $E_{\sim i}$, and so the way in which E(t+1) depends on E(t) is approximately piecewise linear.



Figure 8: Phase space of effort level adjustment

²⁰ As a computational aside, it is tempting to compute the optimal sizes and utilities of figures 6 and 7 directly by differentiation with respect to N. However, this yields a vast expression that cannot be explicitly solved for N, and thus recourse to numerical methods is necessary. It is much more efficacious to obtain these optimal sizes by simple enumeration.

The intersection of this piecewise linear function with the 45° line is the equilibrium total effort. However, if the slope of this function at the intersection point is less than -1 then the equilibrium will be unstable. It will presently be demonstrated that in any group there exists a maximum stable size, dependent on the characteristics of the agents composing the group, beyond which the Nash equilibrium in effort is dynamically unstable.

Analysis for Fixed Group Size

Consider the *N* agent group to be in some state other than equilibrium at time t, described by the vector of effort levels, $e(t) = (e_1(t), e_2(t), ..., e_N(t))$. Now suppose that at t+1 each agent adjusts its effort level according to (5) above using the previous period's value of $E_{\sim i}$, that is,²¹

$$e_{i}(t+1) = \max\left[0, \frac{-a - 2b(E_{i}(t) - \theta_{i}) + \sqrt{a^{2} + 4ab\theta_{i}^{2}(1 + E_{i}(t)) + 4b^{2}\theta_{i}^{2}(1 + E_{i}(t))^{2}}}{2b(1 + \theta_{i})}\right].$$

Since each agent is engaged in effort level adjustment in this way there results an N-dimensional dynamical system, the stability of which is assessed from the eigenvalues of its Jacobian matrix.²² The Jacobian is formed by differentiating each of the RHSs with respect to each agent's effort. Doing this yields, for the case of $i \neq j$:

$$J_{ij} \equiv \frac{\partial e_i}{\partial e_i} = \frac{-1 + \theta_i^2 \frac{a + 2b(1 + E_{\sim i}^*)}{\sqrt{a^2 + 4b\theta_i^2(1 + E_{\sim i}^*)[a + b(1 + E_{\sim i}^*)]}}{1 + \theta_i},$$
(9)

1

while $J_{ii} = 0$. Since each $\theta_i \in [0, 1]$ it can be shown that $J_{ij} \in [-1, 0]$, and J_{ij} is monotone increasing with θ_i as appeal to figure 1 makes clear. Furthermore, for either *b* or $E_{\sim i} \gg a_i$,

²¹ Other effort level adjustment functions produce results qualitatively similar to those described below as long as they are decreasing in $E_{\sim i}$ and increasing in θ_i , both reasonable conditions in this strategic situation. Note that while this is a dynamic strategic environment, and ideas from dynamic games might apply, we eschew such notions here. Although agents have full information about their own payoffs they do not know very much about the other agents, and make no attempt to deduce an optimal multiple period strategy. Rather, at each period they 'best respond' to their environment and are thus myopic. This simple behavior is sufficient to produce very complicated dynamics. The intrinsic complexity of the environment, therefore, suggests it is highly unreasonable that anything resembling sub-game perfect strategies, or even mixed strategies, could ever emerge in the agent population.

²² Technically, agents who put in no effort do not contribute to the dynamics, so the effective dimension of the system will be strictly less than N when such agents are present.

$$J_{ij} \approx \frac{\theta_i - 1}{\theta_i + 1}.$$
 (10)

That is, when either increasing returns are manifest or group effort levels are high, the value of J_{ij} is independent of all parameters except agent *i*'s type. Note that in general the RHS of (9) is independent of *j*, so each row of the Jacobian contains the same value off the diagonal, i.e., $J_{ij} \equiv k_i$ for all $j \neq i$. The overall structure of the Jabobian is thus:

$$J = \begin{bmatrix} 0 & k_1 & \cdots & k_1 \\ k_2 & 0 & \cdots & k_2 \\ \vdots & & \ddots & \vdots \\ k_N & \cdots & k_N & 0 \end{bmatrix}.$$

Stability of equilibrium requires that the dominant eigenvalue, λ_0 , of this matrix have modulus strictly inside the unit circle. It will now be shown that this condition holds only for sufficiently small group sizes. First, since each entry of *J* is non-positive, it will be convenient to work with *-J*, and thus each k_i is positive; *J* and *-J* have eigenvalues of the same magnitude. Now we will establish bounds on $|\lambda_0|$ and see how these change with the group size, *N*. Call r_i the row sum of the *i*th row of *J*. It is well-known (Luenberger [1979: 194-195]) that

$$\min r_i \le |\lambda_0| \le \max r_i \,. \tag{11}$$

Given that the rows of J are comprised of identical entries, (11) amounts to

$$(N-1)\min k_i \le |\lambda_0| \le (N-1)\max k_i$$
. (12)

ı.

Considering the lower bound it is clear that whenever the least $k_i > 0$ there will always be some value of N beyond which $|\lambda_0| > 1$ and the solution is unstable. Furthermore, since small k_i correspond to agents with high θ_i , it is the most productive members of a group who determine its stability. From (12) it is possible to develop an expression for the maximum stable group size, N^{max} , by setting the lower bound equal to 1 and rearranging:

$$N^{\max} \le \left\lfloor \frac{1}{\min_{i} k_{i}} + 1 \right\rfloor,\tag{13}$$

where $\lfloor z \rfloor$ refers to the largest integer less than or equal to *z*. Groups larger than N^{max} will never be stable, that is, (13) is an upper bound on group size.

In the special case of either large *b* or high 'other agent effort,' $E_{\sim i,}$ (10) can be used in (13) to obtain an expression for N^{\max} in terms of agent preferences, as²³

$$N^{\max} \le \left\lfloor \frac{2}{1 - \max_{i} \theta_{i}} \right\rfloor.$$
(14)

So if the agent with the highest preference for income in the group is known the maximum stable group size can be readily established.

Returning to the general case, other bounds on $|\lambda_0|$ can be obtained through reference to the column sums of *J*. Noting the *i*th column sum by *c_i*, we have

$$\min_{i} c_{i} \leq \left| \lambda_{0} \right| \leq \max_{i} c_{i} ,$$

which, given the structure of *J*, means that

$$\sum_{i=1}^{N} k_i - \max_i k_i \le |\lambda_0| \le \sum_{i=1}^{N} k_i - \min_i k_i.$$
(15)

These bounds on $|\lambda_0|$ can be written in terms of the group size by substituting $N \bar{k}$ for the sums. Then an expression for N^{max} can be obtained by substituting $|\lambda_0| = 1$ in the lower bound of (15) and solving for the maximum group size, yielding

$$N^{\max} \le \left\lfloor \frac{1 + \max_{i} k_{i}}{\overline{k}} \right\rfloor.$$
(16)

Firms beyond this size are not stable.²⁴

Example: Onset of instability in a homogeneous group having $\theta = 0.7$

Consider a group of agents having income preference $\theta = 0.7$, with a = b = 1. For homogeneous groups the size bounds given by (13) and (16) are identical. The approximation given by (14) reveals that the maximum stable group size, $N^{\text{max}} = 6$. Let us consider how instability sets in as the group grows in size. For a single agent working alone the optimal effort level can be computed directly from (7) above, and is equal to 0.770, while its resulting utility is available from (3) and amounts to 0.799. Now imagine two agents of this type working together. From (8) the Nash equilibrium effort levels are 0.646 and (3) now yields greater utility for each of 0.964. Each element of the Jacobian (9) is identical when agents are

20

²³ Remember that *-J* is being used instead of *J* here, so expression (10) must be multiplied by -1.

²⁴ Since the largest k_i corresponds to the smallest θ_i , it is the agents who most prefer leisure to income who determine the value of the numerator in (16). But note that it is not merely the agent with the minimum θ who matters, since such an agent will not generally be putting in any effort whatsoever. The agent who is putting in strictly positive effort and has the smallest θ is the one who determines the RHS of (16).

homogeneous; call this *k*. For N = 2, k = -0.188. The bounds on the dominant eigenvalue given by (12) are tight when agents are homogeneous. This yields $\lambda_0 = k = -0.188$. For a group composed of three such agents the optimal effort of each is further reduced, utility is higher, and the dominant eigenvalue is -0.552. The same qualitative results—progressively lower effort levels and higher utility levels—hold for groups of size 4 and 5, with the modulus of the dominant eigenvalue increasing monotonically toward 1. At group size of 6 effort levels have declined further, yet now each agent's utility is somewhat lower than in the size 5 case. Finally, for N = 7, both effort and utility levels have further declined and the dominant eigenvalue now has modulus greater than unity. Thus, this group size is unstable in the sense that any perturbation of the Nash equilibrium effort levels would cause effort level fluctuations that grow without bound. These results are summarized in the following table.

Ν	e *	U(e*)	k	$ \lambda_0 = (N\text{-}1)k$
1	0.770	0.799	not applicable	not applicable
2	0.646	0.964	-0.188	-0.188
3	0.558	1.036	-0.184	-0.368
4	0.492	1.065	-0.182	-0.547
5	0.441	1.069	-0.181	-0.726
6	0.399	1.061	-0.181	-0.904
7	0.364	1.045	-0.180	-1.082

Table 1: Onset of instability in a group of agents having θ = 0.7; Nash equilibrium effort levels in groups beyond size 6 are unstable

All groups of greater size are, of course, also unstable in this sense. For homogeneous groups having lesser preference for income the onset of instability occurs for smaller sizes, while groups composed of agents having relatively greater preference for income can support larger numbers.²⁵

Calculations of the type illustrated in the example can be performed for homogeneous groups having any θ . In figure 9, below, the maximum stable group size is shown as a function of θ , along with the smallest group size at which instability occurs.

²⁵ This can be seen from (10), remembering that θ_i corresponds to higher |k| and thus higher $|\lambda_0|$, other things being equal.



Figure 9: Onset of instability in homogeneous groups; for each θ the Nash equilibrium effort level is unstable for sufficiently large group size

According to this analysis, groups larger in size than 200 would be rarely observed in a population of homogeneous agents. Note the qualitative similarity of figures 9 and 6. It seems that the optimal size of a homogeneous group is very nearly at the stability boundary of the group. For example, in the $\theta = 0.7$ case the optimal size is 5 while the maximum stable size is 6. This suggests that firms operating at or near optimal size and who did not know they were near a stability boundary would be vulnerable to destabilization by addition of one or a few agents. Alternatively, were group size a choice variable in a stochastic environment it might be dangerous to select the optimal group size, for a small perturbation to the next larger size could bring on unstable oscillations and, presumably, demise of the group.

The bounds given by (13) and (16) are the same for homogeneous groups, but are generally different when groups are heterogeneous. For any particular group—particular set of θ_i s—one of these expressions will be larger than the other and thus not binding. It is possible to establish which expression is binding as a function of certain statistics of the agents who constitute the group. In particular, (16) is in effect when

$$\frac{1}{\min_{i} k_{i}} + 1 \ge \frac{1 + \max_{i} k_{i}}{\overline{k}}$$

This expression can be rearranged to yield

$$\bar{k} \ge \min_{i} k_{i} \frac{1 + \max_{i} k_{i}}{1 + \min_{i} k_{i}}.$$
(17)

Since $k_i \in [0,1]$, the fraction on the RHS $\in [1,2]$, and so if $\bar{k} > 2 \min_i(k_i)$ then (16) is the appropriate bound on group size. Note that for $\min_i(k_i) \approx 0$, (17) reduces to $\bar{k} \ge \min_i(k_i) [1 + \max_i(k_i)]$.

Analysis for Variable Group Size

The above analysis has been conducted assuming that the group in question is of some fixed size, N. Now the effect of variable group size will be investigated. It will be shown that if the composition of agents in the group is permitted to change over time then the combined effort level-group size equilibrium will generally be unstable. In particular, let us permit an agent leave the group to start-up a new firm, consisting of only itself, whenever it is utility-maximizing to do so. The optimal effort level in a single agent firm is given by (7) above, and the utility associated with this can be calculated directly from (3), with N = 1 and $E_{\sim i} = 0$.

First, consider the case in which agents having $\theta_i < \theta_c$ are permitted to join the group. Since it is optimal for such agents to put in no effort, they do not contribute to the firm's *E*. The only effect of such agents is to dilute each agent's share of the output and thus depress the utilities of all agents in the group. Clearly, as more and more agents with low preferences for income join the group then the productive (high θ_i) agents would eventually leave. Thus the group size-effort level equilibrium is not stable in this situation.

Next, consider what happens when an agent with a high θ_i joins a group that is in equilibrium. Such an agent will put in positive effort, thus increasing *E* initially. However, this will lead all other agents to adjust their efforts downward. The net effect on *E* is ambiguous; it depends on the exact composition of the group. But assume that the bound given by (16) is binding. Then, the addition of a high θ_i (low k_i) agent has the effect of decreasing \bar{k} and thus increasing the maximum stable group size, N^{max} . As agents having θ_i at or above the average θ are progressively added, the firm remains safely removed from the stability boundary. At some point, though, agents with relatively low θ_i begin to free ride, and the group first passes through its optimal size, beyond which some agents may leave to join other firms or start up new firms. If the group size subsequently passes through the maximum stable size then unstable effort level oscillations set in. This may lead to the group shedding particular dissatisfied agents and re-establishing itself at a stable size or, if the oscillations are severe, to break-up of the group altogether.

Unstable Equilibria and Pattern Formation Far From Equilibrium

Now, one reaction to unstable equilibria is that the model is problematical. Implicit in this is the assumption that social phenomena are

the equilibrium outcome of some game.²⁶ But this reaction is certainly too narrow to be defensible. Games in which the optimal choice is a cycle of strategies have been known since early on in the development of game theory (e.g., Shapley [1964]; see Shubik [1997] for an overview). It is possible to define more general solution concepts to include such possibilities (see, for instance, Gilboa and Matsui [1991]).²⁷ Furthermore, the basic structure of game theory assumes that the strategic environment is fixed and known by all players. However, in multi-agents systems it is common for agents to be placed in combinatorially rich environments in which they are continuously confronted with novel circumstances. Such is the case here.

In particular, when it comes to firms it would seem that any model that suggests there is a single optimal size and effort level would stand in blatant opposition to the facts (Ijiri and Simon [1977: 9-10]). For there do not exist specimens of firms having constant composition and output. As firms grow, agent-agent dynamics shift, some agents leave because they do not like their new co-workers, more or less shirking arises, and more agents may need to be hired.²⁸ Indeed, there is vast turnover in actual firms. My colleague, Margaret Blair, has recently concluded a study of firm dynamics [Blair *et al.*, 1997], in which it is reported that of the largest 5000 U.S. firms in 1982, in excess of 65% of them no longer existed as independent entities by 1996! There is enormous flux in real firms, volatile micro-dynamics for which term 'turbulent' may be appropriate (Beesley and Hamilton [1984], Ericson and Pakes [1995], Sutton [1997]).²⁹

²⁶ For example, Osborne and Rubinstein [1994: 5] seem to suggest that the 'steady-state' interpretation of game theory implies that any empirical regularity is necessarily an equilibrium. They cite Binmore [1987, 1988], who first describes Simon's distinction between *substantive* and *procedural* rationality and acknowledges that the former notion is essentially a static one. He then goes on to distinguish *eductive* and *evolutive* (aka steady-state) ways that players might arrive at equilibrium in a game, claims that each of these describes "a dynamic process by means of which equilibrium is achieved" (Binmore [1987: 184]), but never attempts to justify the focus of game theory on equilibrium.

²⁷ Two points of terminology. First, 'dynamically unstable Nash equilibria' have been referred to in various places. To some this may seem like a bad choice of terms, since if agents have incentive to deviate from these unstable equilibria then they cannot be Nash equilibria in the first place. To the present writer the terminology adopted here seems no more problematical than calling any dynamically unstable equilibrium an equilibrium-such points solve first order conditions but are not rest points. Second, generalizations of Nash equilibrium that permit, for example, cycles, are commonly termed equilibria. But to call such outcomes 'equilibria' is something of a curiosity, one that has been described elsewhere (Epstein and Axtell [1996: 137]) and so will not be further elaborated here.

²⁸ Good arguments against equilibrium theorizing in the context of the firm are given by Kaldor [1972] and Lazonick [1991].

²⁹ Mandelbrot [1998] has recently claimed that the term 'turbulence' does not derive from physics, as one might reasonably suppose, but rather has economic origins. The British hydrodynamicist Reynolds, whose name is now closely associated with certain classes of

In other economic contexts, non-equilibrium models have proven useful, at least for theoretical work. For example, Papageorgiou and Smith [1983] model agglomeration as a local instability of a uniform equilibrium. More recently, Krugman [1996] makes use of the so-called 'Turing instability' in a spatial model of economic activity; see also Heikkinen [1997].

While non-equilibrium models may be a novelty in economic theory, they are well-established in other branches of science. To cite but one example, in mathematical biology the instabilities of various partial differential equation systems that couple chemical reaction and diffusion processes—including the Turing stability—are the basis for a variety of pattern formation processes (cf. Murray [1989]). These models explain certain features of embryonic development, neuronal activity, cardiac rhythms, and even the patterns on animal coats (e.g., 'how the leopard got its spots').

3 Computational Implementation with Agents

The motivation for a computational version of the team production model described above is simple. Since there do not exist stable equilibria of the model we must resort to studying the model's non-equilibrium dynamics. Perhaps this dynamical picture will contain patterns of firm formation and growth that are recognizable vis-a-vis actual firms. Such non-equilibrium patterns are usually very difficult to discern analytically, leaving computational models as the only practical technique for systematically studying such phenomena.³⁰ In what follows we discover computationally that stationary non-equilibrium patterns *do* exist in this model and that they have a close connection with empirical data.

turbulent flows, apparently borrowed the term from contemporaneous descriptions of the London stock market (c. 1880s) in order to communicate, by analogy, the highly irregular nature of the fluid flows he was studying. For more on turbulence see the next footnote.

³⁰ There is a close analogy to be made between the approach we are taking herein and computational physics. In particular, in fluid mechanics laminar flow of an ordinary fluid is the regime that solves the so-called Navier-Stokes equations with all transient terms set to 0. This solution can be shown to be unstable beyond a critical value of a dimensionless parameter of the flow, the so-called Reynolds' number. For Reynolds' numbers beyond critical it is not the case that the flow blows up—a conclusion one might reach if the linearized solution were interpreted literally (!)—but merely that the fluid reorganizes itself into a non-laminar configuration in which transients persist. For sufficiently large values of the Reynolds' number the flow becomes fully turbulent and transient phenomena exist on all length and time scales. These non-laminar regimes, despite the fact that the governing equations are well known, have proven nearly impervious to analytical investigation. Today, computational techniques are the method of choice for the systematic study of non-laminar flows.

3.1 Set-Up of the Computational Model

In the analytical model above the focus was on a representative group. In the computational model a population of agents will form many groups and interactions between agents and groups will be explicitly modeled.

The set-up is essentially identical to the above. Total output of a firm is similar to before, consisting of both a linear term and a term responsible for increasing returns. However, it is desirable to work with a slightly more general form of (2), in which the quadratic nature of the second term is relaxed as

$$O(E) = aE + bE^{\beta}, \qquad (2')$$

 $\beta \ge 1$. Initial realizations of the computational model will be made with $\beta = 2$ and the effect of changing β will be studied subsequently.

Preferences are heterogeneous in the agent population; for all the results described below $\theta \sim U[0, 1]$. Each agent belongs to a social network insofar as it maintains reference to some number, v_i , of other agents. These are assigned randomly at time 0 and do not vary over time; $v_i = 2$ for all *i*. Later, the effect of varying this parameter will be studied.

Agents are activated at random, that is, each agent has a Poisson clock that wakes it up periodically. When an agent is activated it looks up the size and output of its firm as well as its own previous period effort level, and then uses this information to select the effort level that maximizes its utility.³¹ Associated with this optimal effort is an optimal utility level. The agent then repeats this calculation for (1) starting up a new firm, in which it is the only agent, and (2) joining each of its v_i friends' firms. The agent then acts based on which of these options yields the greatest utility. That is to say, it either stays in its current firm and adjusts its effort level accordingly, or it leaves its current firm for another firm, where its best effort level is prescribed similarly. Since agents are not evaluating all possible firms this is a form of limited information processing.³²

A time period consists of some number of agents being activated. The case of all agents being activated exactly once during a single period is known as *uniform* activation. Alternatively, when agents are selected randomly for activation then the number of activations across the agent population has some distribution with non-zero variance. For example, when agents are

³¹ In order for the computational model to be applicable to agents having arbitrary preferences, each agent performs a line search over the feasible range of efforts in lieu of using the expression given by (5) above.

³² In § 4.6 below agents will be made boundedly rational. Instead of calculating its optimal effort level, each agent must grope for effort levels which yield utility improvements. It will be shown there that the qualitative character of the overall model is robust to this respecification of behavior.

selected with uniform probability then the overall distribution of agent activations follows a binomial distribution. It is common to call this *random* activation. It there is some reason to believe that agent activity is homogeneous in a population then uniform activation is appropriate. Otherwise, it usual to employ random activation in agent-based models, in which some agents are more active than others during any particular period. Asynchronous activation is used herein, with a time period defined as 1000 agents being active.

The initial configuration of the agent population has each agent working alone in its own firm. Thus there are *A* firms initially, of mean size 1 with no variance. Over time multi-agent firms form, grow, and perish, and more or less stationary distributions of firm size, firm growth rate, and firm lifetime emerge.

Model Attribute	Value	
agents, A	1,000	
constant returns coefficient, a	1	
increasing returns coefficient, b	1	
increasing returns exponent, β	2	
distribution of preferences, θ	U[0, 1]	
sharing rule	equal shares	
number of neighbors, v	2	
agent activation	random	
initial condition	all agents in singleton firms	

The parameterization of the computational model described above is summarized in table 2, which we shall term the 'base case' parameterization.

Table 2: 'Base case' configuration of the computational model

Initial realizations of the computational model will be made using this base case. Then, in § 4 many variations on the base case will be studied systematically.

The essential feature of this model is that it is described at the level of individual agents. Thus it is common to call such models 'agent-based' or 'individual-based'.³³ The only equations present in the model are those

³³ While agent-based models are methodological individualist, they need not suffer from the various problems that often beset analytical models of individual behavior, where agents are frequently treated as interacting not with other agents but with aggregate statistical measures. For more on the relation of methodological individualism and agent-based models see Epstein and Axtell [1996: 16-17].

governing individual agent decision-making. No attempt has been made to mathematically aggregate the agents' behaviors, and therefore there are no equations governing agent-agent interactions. Rather, "solving" an agent-based model amounts merely to iterating it forward in time and observing the evolution of the agent population, both at the individual and aggregate levels. In this, agent-based computational modeling is similar in spirit to traditional OR simulation. However, in most ways that agent-based computational models have been used to date this methodology is quite unlike conventional simulation.³⁴

Typical dynamics of the model are described in § 3.3 below. Then, in § 3.4 we study the distributions of firm sizes, growth rates and lifetimes that the model yields and compare these with empirical data. § 3.5 gives a picture of the typical firm life cycle, and § 3.6 investigates agent welfare in the model. But first, § 3.2 describes the actual computational implementation of the model.

3.2 Object-Oriented Implementation

There are many ways to computationally implement the model just described. This can be done more or less easily in any modern programming language, as well as with any number of mathematical or simulation software packages. However, since the model is stated in terms of individual agents, it turns out that there is one idea from modern computer science that renders the implementation both transparent and efficacious. This is the notion of object-oriented programming.

Objects are contiguous blocks of memory that contain both data—socalled instance variables—as well as functions for modifying this data—the so-called methods. This ability of objects to hold both data and functions is called *encapsulation*. Agent-based models are very naturally implemented using objects by interpreting an object's data as an agent's state information, while the object's functions become the agent's rules of behavior.³⁵ A population of agents that have the same behavioral repertoire but local state information is then conveniently implemented as multiple instantiations of a single agent object type or class.³⁶

The model described above has been implemented using objectoriented programming. Not only are individual agents objects, but individual firms are objects too, albeit of a different class than agent objects.

³⁴ For more on the distinction between simulation and agent-based computational models see Axtell [1997].

³⁵ Other features of the object model, including *inheritance* and *polymorphism*, seem to be less relevant to agent-based computational models than encapsulation.

³⁶ For a discussion on the distinction between object and agent, see Jennings *et al.* [1998].

In fact, it turns out to be convenient for the population of agents as a whole to be an object as well, as is the population of firms.

The agent object has a variety of state variables and behavioral methods. The state information each agent has includes its preferences and its current effort level. It is also useful to keep track of each agent's income and utility associated with production from the previous period, as well as a running total of its wealth. All of this information is stored locally in the agent object as real numbers. Each agent also keeps track of some number of other agents that are identified as its social network. This data is maintained in an array of pointers to other agent objects. Finally, each agent must also keep track of the firm to which it belongs, and possibly-depending on details of implementation—other agents in the master agent population list as well as in its firm. The behavioral abilities that each agent possesses in the present model include the facility to compute its utility (given some income and effort level), the capability of determining its optimal effort level in its present firm as well as in a friend's firm, as well as the capacity to leave its present firm and to join another firm. These are the agent object's methods. This agent object specification is summarized in pseudo-code block 1.

```
OBJECT agent;
  preferences;
  neighbors;
   effort_level;
   last_income;
   last_utility;
  wealth;
   firm_to_which_agent_belongs;
  next_agent_in_agent_list;
  next_agent_in_firm;
  FUNCTION initialize;
  FUNCTION compute_utility;
  FUNCTION compute optimal effort level in present firm;
  FUNCTION compute_optimal_effort_level_in_other_firm (the_firm);
  FUNCTION leave_present_firm;
  FUNCTION move_to_new_firm (the_new_firm);
  FUNCTION draw.
```

Pseudo-code block 1: Agent object

In practice it makes sense to implement as private some of these data and methods, while others are public, although this is not essential.³⁷

The agent population is also conveniently implemented as an object. The details are less important here and so pseudo-code will not be given. The main data of the population object is the data structure—commonly either an array or linked list—that holds either the individual agents themselves, or

³⁷ Private data and methods are accessible only by the agent instance to whom they belong, unless other objects are given special access privileges.
reference to them. The agent population object's methods include agent access routines as well as a host of routines for computing various statistical measures of the agent population, such as the average effort level, average income and so on.

Firms are also conveniently implemented as objects. Firm objects are created whenever an agent decides to start up a new firm. They are destroyed whenever a firm consists of a single agent and that agent decides to join some other firm. Firm-specific data that each firm object holds includes the current size of the firm (an integer) and the growth rate and output from the last period (real numbers). The agent who founded the firm (pointer to agent object) is also stored, as well as a data structure that keeps track of the agents currently in the firm. The methods of the firm object include routines for adding and removing agents from the firm, accessing individual agents, computing output and dividing it up among the agents, and various statistical calculations, such as computation of the mean and variance in effort levels. The firm object is summarized as pseudo-code block 2.

```
OBJECT firm;
founder;
agent_list;
size;
growth;
last_output;
FUNCTION initialize;
FUNCTION compute_total_effort;
FUNCTION compute_average_effort;
FUNCTION compute_output;
FUNCTION compute_output;
FUNCTION allocate_income_to_agents;
FUNCTION add_agent_to_firm (the_agent);
FUNCTION remove_agent_from_firm (the_agent);
FUNCTION draw;
FUNCTION dispose.
```

Pseudo-code block 2: Firm object

As with the agent object, in practice it is usually useful to make some of these fields private. The population of firms is also an object. Similarly to the agent population object, the firm population object holds the data structure that maintains reference to each individual firm, routines for accessing individual firms, as well as a variety of functions for computing statistical properties of the firm population.

Putting all of this together the computational model simply amounts to (1) initializing all agents and firms; (2) activating M agents sequentially, and letting each one choose its optimal effort level, migrating between firms if necessary; (3) computing firm output; (4) periodically gathering statistics on the populations of agents and firms. This is summarized in pseudo-code block 3.

```
PROGRAM firms;
   initialize agents;
   initialize firms;
  repeat:
      select M agents at random;
      for each agent selected:
         compute effort level to maximize welfare at current firm;
         compute effort level to maximize welfare at neighboring firms;
         compute effort level to maximize welfare in start-up firm;
         move to firm where welfare is greatest;
         draw agent;
      for each firm:
         compute output;
         for each agent in firm:
            allocate income;
            compute welfare;
      compute statistics;
      check for user input;
  until user terminates.
```

Pseudo-code block 3: Pseudo-code for the model overall

The object model is largely responsible for the relatively short description of this code.³⁸ In the next section we describe typical realizations of this model.

3.3 Aggregate Dynamics

Initially, all agents are entrepreneurs, working for themselves in oneperson firms. However, as agents are activated they discover that they can do better by cooperating with one of their friends to jointly produce output. We depict this by coloring firm founders red, and initially lining up all agents from top to bottom in a window. As agents join existing firms, they leave their location, are colored blue, and placed in the first available position to the right of and on the same line as the founder of the firm they are joining. Thus, a long horizontal blue line represents a large firm. A typical evolution is shown in animation 1.³⁹

³⁸ The actual source code runs to some 5000 lines and compiles in the CodeWarrior environment for the Macintosh. A Java implementation is available; see the next footnote.

³⁹ Readers interested in running the model for themselves will soon find a Java implementation at http://www.brook.edu/es/dynamics/models/firms. At present this site is under construction.



Animation 1: Snapshots from a typical firm formation process over 8 consecutive time periods; the agent who has been with a firm the longest (typically the firm founder) is colored red, others blue

Initially, many firms grow and patterns emerge. Such a situation is shown in the first frame of animation 1 corresponding to t = 1000. Firms grow from left to right, with the largest firms truncated by the right side of the frame. There are several large firms displayed there, where size is synonymous (for now) with the number of agents. Over time these firms expand, as agents find it in their best interest to join, and then contract as free riding becomes rampant. New firms are born as discontented agents form start-ups. Notice the large firm about 1/4 of the way down in snapshots 1-3. By frame 4, t = 1003, it has begun to decline. Alternatively, the large firm that arises at the top of the second frame survives through the last frame. There is so much flux in firm composition that it does not seem unreasonable to call the micro-dynamics of

Number of Firms

From animation 1 it can be seen that the total number of firms varies over time. This is due both to the entry of new firms—agents leaving failing firms to form start-ups—as well as the demise of extant firms, predominately due to the onset of free-riding. As the model spins forward in time it is an easy matter to cull, without sample error, each period's data concerning the total number of firms, the number of new firms, and the number of failed firms. It is instructional to plot these as time series; this is done in figure 10 below for the realization depicted in animation 1.



Figure 10: Typical time series of the total number of firms (upper line) and the number of new firms (lower line)

Note from this figure that there are significant fluctuations about more or less constant levels of total firms and new firms. The intermittent bursts of firm start-up activity are closely related to the onset of decline in large firms—many of the agents in the failing firm find it optimal to form new firms.

In order to better understand the dynamical structure of figure 10, the power spectrum of the total firm number time series has been computed. This is shown in figure 11 below, in log-log coordinates.

⁴⁰ Although, it will eventually be demonstrated that stationarity obtains at the macro-level.



Figure 11: Power spectrum of the time series for the total number of firms

Note that there appears to be no single time scale for these fluctuations. Rather, components of all frequencies, f, over more than a decade are present in this series. There exists a power law tail in this distribution that scales like $1/f^2$, meaning that the total firms time series is Brownian in character. A line having slope -2 has been superimposed on the data in figure 11.

Firm Size

The average firm size fluctuates over time, as does the size of the largest firm present. Typical time series for average firm size and maximum firm size are shown in figure 12; this corresponds to the realization depicted in animation 1.



Figure 12: Typical time series for average firm size (inset) and largest firm size

Note that a large firm with size in excess of 200 grew up in the first 200 periods, then abruptly perished. While the maximum firm size fluctuates dramatically, the average firm size overall is not highly variable.

Effort Levels

Given the utility maximization calculus that all agents execute, the reason to join with other agents in production is that the effort level one utilizes when working alone yields greater production when working cooperatively, and thus cooperating agents achieve higher individual utility. Stated differently, agents put in relatively less effort for the same income when working together, and this leaves more time for leisure and is therefore welfare-enhancing.

Agents in small firms tend to have relatively high effort levels, albeit not as high as if they were working alone. As a firm grows, and production rises superlinearly with total effort, production shares also rise. Each agent experiences increasing insensitivity of its share to its effort level, and so its optimal effort level progressively falls.

Time series plots for average effort levels in the population as whole as well as within the largest firm show this process graphically. For the realization described by animation 1, figure 13 gives the corresponding evolution of effort levels, with the heavy line corresponding to the average effort overall and the thin line referring to the average effort in the largest firm.



Figure 13: Typical time series for average effort level in the population (thick line) and in the largest firm (thin line)

Comparing this figure with figure 1 we see that for the large firm that existed through t = 200, its average effort level plummeted to near zero before it finally collapsed. In fact, the minima in the figure 13 effort plot for large

firms correspond almost exactly to peaks in the largest firm size of figure 12. Note that while effort in large firms fluctuate severely, the overall average effort level is quite stable.

Output

Just as effort levels change continuously, so too do output levels. In particular, it is revealing to compare a time series plot of the total output of all firms with the output coming from the largest firm. Large firms can have very high output levels, due to the increasing returns of production. So, before free riding sets in, it can be the case that a significant fraction of total output is due to a single large firm.

Figure 14 gives the time series of total output (thick line) and output of the largest firm (thin line) for the realization described in animation 1. Here we see that over most of the 1000 periods shown, the output of the largest firm did make up a significant fraction of the total. This is especially true early on, when the single large firm that lived through t = 200, ended up producing nearly 50% of the total output, this despite the fact that even at its peak it never contained more than about 20% (200/1000) of the agent population.



Figure 14: Typical time series for total output of all firms (thick line) and output of the largest firm (thin line)

Interestingly, note that the output of this first largest firm is not monotonic. Prior to its ultimate demise at around t = 250, it suffered an output decrease right at t = 200, but weathered this storm and reached a new output high, before succumbing. Another feature of this plot is the output plateaus of large firms. Several of the output maxima are less peaks than flat tops. It is as if firms facing declining effort levels maintain total production, perhaps through the incorporation of new agents.

Іпсоте

If we divide the total output of figure 14 by the number of agents we get the average income (share) and its evolution over time. Once again, it is instructive to compare this with incomes in the largest firms. This is done in figure 14, where the thick line corresponds to average income in the population overall and the thin line is income in the largest firm. While the income (heavy) line of figure 15 looks like a compressed version of the output (heavy) line in figure 14, the same is not true of the income and output lines for the largest firm (thin lines in the figures). There is substantial variation in income from period to period in the largest firms. This is because membership in these firms is constantly changing. Note that large drops in income correspond to the onset of free-riding (falling effort levels) and ultimately to firm dissolution.



Figure 15: Typical time series for average income in the population (thick line) and in the largest firm (thin line)

At first blush it seems something of a curiosity that large firm incomes fall below the average quite regularly in figure 15. Indeed, they nearly reach zero on several occasions. Why would agents stay in large firms if they could get more income at the 'average' firm? This is explicable by remembering that agents are *utility maximizers*, and that low income levels correspond to near zero effort levels. But agents who are, in effect, not working, are deriving great utility from leisure, and so are willing to forego income for leisure. Such agents typically have a relatively high preference for leisure over income, and do not leave these debilitated firms until the very end.

Utility

Now, it would be paradoxical if agents in the largest firms received below average *utility* as firms declined. This can be tested explicitly with this computational model, by making time series plots of average utility in the population overall and in the largest firms. It is usual to think of utility as being ordinal, not cardinal, of course. But since our agents all have Cobb-Douglas preferences, it seems not too unreasonable to compute the average utility by merely summing up individual agents' utilities and dividing by the number in the population, *A*. Similarly for average utility in the largest firms.

Figure 16 shows the result of these computations, with the thick line representing the average utility overall, and the thin line the average utility in the largest firm. In many ways figure 16 is very similar to figure 15. However, note that the large firm mean utility time series does not spend much time below the average utility in the population.



Figure 16: Typical time series for average utility in the population (thick line) and in the largest firm (thin line)

That is, as soon as utility levels in declining, large firms reach the level of utility that is generally available in the population, agents leave *en masse* for better opportunities. Occasional dips below the average level are to be understood as unlikely outcomes, in which agents are trapped in failing firms and, by virtue of circumstance—unsuccessful friends—cannot immediately secure alternative employment.

Animation 1 and figures 10-16 have been presented to build-up the reader's intuition about typical dynamics of firm formation, growth and dissolution. They are a 'longitudinal' picture of typical micro-dynamics of the agent population, and the corresponding behavior of firms. We now turn to aggregate statistical properties of typical realizations.

3.4 Distributions of Firm Sizes, Growth Rates and Lifetimes

The distribution of firm sizes across industries has a very characteristic shape, that of a power law (also commonly known as a scaling law or Pareto distribution). Data on firm sizes, measured variously, have proven to have broadly robust power law behavior over many decades and across national borders, this despite waves of mergers and acquisitions. Power law distributions of firm sizes are also characteristic of the non-equilibrium model described above. This is studied in the next section. Following that, the distribution of firm growth rates yielded by the model is compared with data. Then, the distribution of firm lifetimes the model produces is studied. These also follow a power law.

Firm Size Distribution

At any instant of time there exists a distribution of firm sizes in the model, as measured by the number of agents in each firm. Noting the probability of realizing a firm of size s as p(s), we employ the following parameterization of the discrete power law distribution (probability mass function):

$$p(s;\mu,s_0) = \left(\frac{s}{s_0}\right)^{-(1+\mu)},$$
(18)

where s_0 is a certain minimum size.

From table 2 above, all firms are of unit size at t = 0 in the model. Therefore, the size distribution starts out as a Dirac measure. Over time it acquires a more or less stationary configuration, with a few large firms and progressively greater numbers of smaller and smaller ones. Data from the model are shown in doubly logarithmic (log-log) coordinates in figure 17, for the realization described above.⁴¹

⁴¹ It is perhaps useful to point out that this is noiseless data, in the sense that it contains no sample error. To some readers this point will seem utterly obvious and underserving of mention. However, it is important to point out that the deviations of the data from econometric fits to simple functional forms are due completely to random variation and misspecification errors. As sample sizes become large—and in computational models samples can often be made arbitrarily large—the residual errors are solely due to misspecification.



Figure 17: Stationary firm size distribution (probability mass function)

The average firm size is about 4. In the small populations employed herein this distribution will, at each time, have lots of gaps. That is, since large firms are rare, it is unlikely that two large firms exist at any instant. Therefore, the data shown in figure 17 are the result of repeated sampling.⁴² Notice that the data are quite irregular for low probability events.⁴³ Dropping the data that occur with frequency less than 10⁻⁵, as well as the N = 1 datum, it is possible to fit a scaling (power law) distribution to these data with high confidence.⁴⁴ In particular, OLS yields the probability of a firm of size *s*, *p*(*s*), as proportional to $s^{-2.28}$ (adjusted $R^2 = 0.99$), thus $\mu = 1.28.^{45}$ Our estimate depends on the

⁴² In order to assure that the samples are independent, the sampling period is greater than the maximum firm lifetime. This procedure does not bias the data since these scaling distributions are stable.

⁴³ Data for the largest firm sizes shown in figure 17 represent single observations, and these occur with frequency $O(10^{-6})$. It is possible to realize larger firms merely by running the model longer. Since the frequency scales approximately as s^{-2} , firms of size 1000 are four times less common than firms of size 500, while firms of size 10,000 are 100 times less common than firms of size 1000. Stated differently, it would be necessary to run this model 400 times longer than was done in preparing figure 17 in order for the expected maximum firm size to reach 10,000. While large firms are certainly important, as a practical matter approximately 80% of the U.S. workforce is employed in firms size 500 or smaller (Acs and Audretsch [1990: 31]).

⁴⁴ When fitting data to scaling laws it is common that data at the extremes do not fit well due to so-called finite size cut-offs. The 'stretched exponential distributions' mentioned in footnote 11 attempt to account for the concavity in data produced by finite size cut-offs.

⁴⁵ It is probably important to emphasize that this is a purely computational result. No analytical derivations are being offered to link the dynamically unstable Nash equilibria analyzed in § 2 and the power law size distribution just described, although such a connection surely exists mathematically and it is an open question just how to establish this

parameterization of the model and we shall systematically study this dependence in § 4 below. Here we simply note that empirical data on firms also display scaling structure. In particular Simon and Ijiri [1977] report that data on U.S. firms c. 1955 are well fit by μ = 1.23, while for British firms μ = 1.11.⁴⁶

Power law distributions of this type have been systematically studied in recent years by Bak and co-workers at the Santa Fe Institute [Bak 1996]. While the ubiquity of such distributions has been well known from at least the time of Pareto, basic explanations of their origins are largely due to physicists working with highly idealized systems, e.g., sandpiles.⁴⁷ In this body of theory resides the general idea that power law size distributions result from 'selforganized critical' processes, in which a system arranges itself such that large fluctuations are always possible, and these occur not infrequently. Intrinsic to this theory is that the system is far from equilibrium, a necessary condition in order to observe non-Gaussian fluctuations. Given the power law character of actual firm size distributions, it would seem that equilibrium theories of the firm, such as principal-agent models, will never be able to grasp this Of course, the computational model essential empirical regularity.⁴⁸ developed herein, which yields power law distributions of the type shown in figure 17, is a non-equilibrium one.

Firm size has been interpreted so far as referring to the number of agents in a group. For modern economies it turns out to not matter whether one uses number of employees or firm sales volume in constructing firm size distribution functions, as the power law exponent of the distributions are approximately the same. This invariance across size measure can be checked here as well. The distribution of firm output too follows a power law in the

formally. Here we see the power of the agent-based computational approach to connect the micro-world of individual agents to empirically-relevant, macro-statistical relationships. Stated differently, the computational approach serves as an engine of (exact) aggregation, surmounting the myriad mathematical difficulties in the formal theory of aggregation; on mathematical intractability of certain related, albeit simpler, problems see footnote 47.

⁴⁶ I have been unable to find published results using more recent data, and so am presently engaged in empirical work on this topic. It seems that since the 1960s essentially all econometric work on firm sizes has focused on intra-industry distributions and measures of size inequality; examples of this include Quandt [1966], Kwoka [1982]; see also the review articles of Curry and George [1983] and Schmalansee [1989: 994].

⁴⁷ Highly idealized models of sandpiles, ecological systems, earthquakes, and related physical phenomena that yield power law size distributions have been the objects of intense study recently. Interestingly, while these models can often be stated in just one or two sentences, and implemented computationally in a few dozen lines of code, they have resisted complete analysis; for more on this, see Bak [1996: 62-64].

⁴⁸ For more on the relation of 'self-organized criticality' to the firm size distribution see Morel [1998].

present model, and its exponent can be estimated. This is done in figure 18 below for the same realization that yielded figure 17 above.



Figure 18: Stationary distribution of firm output (probability mass function)

The mean output is reasonably small, at 8.88. Fitting a power law to the data reveals that the OLS exponent is 0.88; adjusted $R^2 = 0.99$. This value is somewhat different from that obtained from firm size measured by number of agents (employees).

Productivity

One way to interpret the two power law exponents that obtain for firm size measured by number of agents and total output is as a measure of returns at the aggregate level. That is, constant returns to scale implies that doubling the size of a firm yields twice as much output, increasing returns means that more than twice as much is produced, and decreasing returns is the opposite (sub-additive production). Interesting here is the fact that the output distribution exponent is somewhat less than for the distribution according to the number of agents, which seems to imply that there exist *decreasing returns* at the aggregate level.

The relation between firm output and size by number of employees can be studied by looking at firm productivity. Figure 19 is a log-log plot of average output versus size (by number of agents) for the realizations that produced figures 17 and 18, above. Clearly these two measures are highly correlated. There is significant variability at large sizes since these data bins hold few observations.



Figure 19: Productivity exhibits near constant returns at the aggregate level, despite increasing returns at the micro-level

Also shown in figure 19 is the OLS line for a power law fit of the data. This line is given by 0.58 $s^{1.15}$; adjusted $R^2 = 0.94$. Thus, slightly increasing returns apparently describe this data, although the hypothesis of constant returns cannot be rejected. Constant returns is also a feature of the empirical data, since the size distribution is the same whether employees or sales are used as measures of size.

Near *constant returns* at the aggregate level occurs despite, of course, *increasing returns* at the micro-level, suggesting the difficulties of making any inferences concerning micro-level behavior from aggregate data.⁴⁹ A qualitative explanation of why this occurs is apparent. As the increasing returns-induced advantages that accrue to a firm with size are consumed by free-riding behavior, agents migrate to more productive firms. Each agent who changes jobs acts to 'arbitrage' the returns to marginal size changes between firms. Thus, even though the agents are heterogeneous, returns to adding an additional agent are more or less equal across the firm population.

Firm Growth Rate Distribution

Just as there exists a distribution of firm sizes at each period, so is there a distribution of firm growth rates. Call *r* the logarithm of a firm's growth rate, i.e., the log of the ratio of its current size to its previous size. Clearly, firms that are expanding have r > 0, while contracting firms have r < 0.

⁴⁹ The general problem of making micro-level inferences from macro-level data is known as the 'ecological inference' problem; see Achen and Shively [1995].

The so-called 'growth of firms' literature explains firm size distribution data by appeal to Gibrat's law, also known as the law of proportional growth, which states, roughly, that if growth rates are independent of firm size then the size distribution of firms will be right skewed. Operationally, this has usually been taken to mean that if growth rates are normally distributed then firm sizes will be lognormally-distributed, with 'fat' right tails. Recently, however, Stanley *et al.* [1996] report that the Laplace (double exponential) distribution, that is,

$$\frac{1}{\sqrt{2}\sigma_r} \exp\left(-\frac{\sqrt{2}|r-\bar{r}|}{\sigma_r}\right),\tag{19}$$

better fits the data on U.S. firms, where \bar{r} is the average log growth rate and σ_r is the standard deviation of r.⁵⁰

Figure 20 shows the growth rate data in log-log coordinates, from the realization described above, together with fits to Gaussian and Laplace distributions,⁵¹ the parabolic and tent-shaped curves, respectively.⁵² The average growth is 1.000.⁵³

⁵⁰ Related work by these and other researchers has revealed that growth rates of countries are also well described by this distribution; see Amaral *et al.* [1997], Buldyrev *et al.* [1997], Lee *et al.* [1998] and Canning *et al.* [1998]. It is an hypothesis of this research that "the evolution of organizations with complex structure is governed by similar growth mechanisms" (Lee *et al.* [1998]: 3275). Such a mechanism is proposed in Amaral *et al.* [1998]. This mechanism is not written in terms of the individuals who make up organizations, but rather in terms of organizational sub-units. As such, it is rather different in character from the model described here.

⁵¹ These two distributions are closely related and were both recommended by Laplace for purposes of describing empirical data (Johnson *et al.* [1995]).

⁵² The growth rates shown in figure 20 are approximately twenty times larger than in the empirical data (cf. Stanley *et al.* [1996]) for comparable frequencies. This large discrepancy arises simply as a result of the definition of time scale in the model. In particular, since 1000 agents being active was equated with a single time period, if time is rescaled such that activation of 1000/20 = 50 agents represents a single period then the growth rates shown in figure 20 would be comparable to the actual data.

⁵³ There is no net growth on average since the population of agents is fixed in size. Real data must first be detrended. Presumably, net positive average growth rates could be worked into the model simply by permitting the population of agents to grow over time.



Figure 20: Stationary firm growth (probability mass function)

Note that the normal substantially misses the datum in the center bin, where nearly all the probability mass lies. The Kolomogorov-Smirnov test reveals that (16) better fits the data than does the Gaussian distribution, but the extent to which it is the better fit depends sensitively on how the data are binned.⁵⁴

Stanley *et al.* [1996] further find that the standard deviation in the distribution of log growth rates, σ_r , decreases with size according to a power law, i.e.,

$$\sigma_r \propto s^{-\gamma} \tag{20}$$

In particular, they estimate that $\gamma = 0.16 \pm 0.03$ for size based on number of employees, and $\gamma = 0.15 \pm 0.03$ for size based on sales.⁵⁵

It is a relatively easy matter to glean data from the computational model concerning the variation in growth rates with firm size. This has been done for the realization described above. Estimates of the standard deviation of r conditional on firm size are shown in figure 21, in log-log coordinates in order to facilitate estimation of γ .

⁵⁴ One difficulty here arises from the fact that the firm sizes being generated by the model are relatively small—at least in comparison to real firms—and therefore there is a certain amount of discreteness to the growth data. For example, a firm of size 10 might grow to size 11 or 12, representing a 10% or 20% growth rate, respectively. But it cannot grow by 15%.

⁵⁵ Interestingly, variance in growth rate data for *countries* is also well-described by (20) with $\gamma = 0.15$ (*cf.* Canning *et al.* [1998]). See footnote 50 for more on the empirical relationship between firms and countries.



Figure 21: Dependence of the standard deviation in *r* on firm size

Note that while there are significant exceptions, overall σ_r falls with firm size, as in the empirical data. For the smallest firms—sizes 1, 2 and 3— σ_r rises with size, while for relatively large firms—greater than about 300—the data are very noisy, owing to small sample sizes. The functional form given by (20) has been estimated by OLS and is shown in figure 21. This estimation yields $\gamma = 0.174 \pm 0.004$, and was accomplished by dropping the first two data points as well as the outliers occurring at large firm size. These results are in close agreement with the empirical data.

Firm Lifetime Distribution

It is easy to study the distribution of firm lifetimes that comes out of this model.⁵⁶ To do this, the firm object holds as data the year it was born, and then later, as it dies, it computes its age and sends this information to the object that maintains the distribution of firm lifetimes. In this way, data is kept on each firm that forms, and sampling issues do not arise.

For the realization described above, the distribution of firm lifetimes is shown in figure 22 as a size-rank plot, with the rank given in log coordinates. Data on nearly 10⁶ firms are reported in the figure.

⁵⁶ It is notoriously difficult to come up with a good measure of firm lifetime in practice, given that mergers, acquisitions, takeovers, bankruptcies, voluntary liquidations, and private buy-outs are common among real firms. Thus empirical work has focused on the rather restricted categories of exit (cf., Schary [1991]), failure (cf. Lane and Schary [1991]), new firm survival (cf. Audretsch and Mahmood [1995]), turnover and mobility (cf. Caves [1998]), among others.



Figure 22: Firm lifetime-rank distribution

The average lifetime is approximately 23.4 periods while the standard deviation is estimated to be 27.1. The three firms having the longest lifetimes, of approximately 485, 370, and 340 periods, are outliers from a linear dependence of lifetime on log(rank), having OLS slope of -70. This means that as one moves out an order of magnitude in rank, firms live for 70 fewer periods; this behavior is robust over 4 decades of rank.

We have interpreted firm growth and demise as a process in which agents are attracted to high-income firms, these firms grow, and once they become large get over-run with free-riders. In the next section this basic picture is elaborated quantitatively by studying the life of a typical firm.

3.5 Dynamics of Individual Firms: The Firm Life Cycle

Some indication of the life cycle of a firm was apparent above in the evolution of the large firm that survived through time 200, shown in figures 10-16. But this firm was hardly representative due to its extreme size. In this section the life cycle of a typical firm is studied.⁵⁷

There is something like a characteristic firm life cycle in this model, involving the growth of the firm, its change in composition, the onset of free riding, and its ultimate demise. Figure 23 below gives a typical evolution of firm size over the life of a particular firm, with size represented by the number of agents as well as the total output. For firms that grow to greater or lesser sizes the corresponding figures would be scaled up or down, but the qualitative shape remains the same.

⁵⁷ In reality the life cycles of firms are thought to be intimately intertwined with product and industry life cycles; on the former see, for instance, Vernon [1966] or Klepper [1996] while Abernathy and Utterback [1978] is a standard reference on the latter.



Figure 23: Typical evolution of firm size over the life cycle of the firm, by number of agents (thick line) and output (thin line)

This firm lived from time 9 to time 28. Its growth was rapid and nearly exponential in the early stages, followed by a brief period of stagnation around time 20. Subsequently, as output began to fall agents left the firm, creating a vicious cycle responsible for rapid decline. However, the firm remained in existence at a very small size for some time.

Over the life cycle the composition of the firm by agent type is changing. Figure 24 depicts this evolution of firm composition by plotting as a function of time the minimum, average, and maximum preference for income, θ , among the agents in the firm.



Figure 24: Typical evolution of firm composition by agent type over the firm life cycle

Initially, the firm is dominated by agents with high preference for income. Over time there is a clear tendency for the firm to include agents with lower and lower θ s—agents who value progressively value leisure relatively more than income. Sown in this picture are the seeds of the firm's ultimate destruction, for agents with low θ put in relatively low amounts of effort and

are among the first to engage in free-riding. This can be seen in figure 25 below, which is a time series plot of the average effort level in the firm, with bars above and below the average representing the maximum and minimum effort levels, respectively.



Figure 25: Typical evolution of effort levels in a firm over the life cycle

Overall, effort levels are decreasing approximately linearly and essentially monotonically, on average. Note that at all times there are some agents who contribute no effort to production. But this plot does not make clear how many agents are behaving in this way, i.e., how many agents are free-riding. Data from the model on the number of free-riders in this firm is depicted in figure 26.



Figure 26: Typical evolution of the number of free riders in a firm over the life cycle Note that the number of free riders peaks at approximately time 21. That is, by the time the firm reaches its maximum size free-riding has become rampant. The overall shape of this figure resembles closely the firm size trajectory, measure by the number of agents (figure 23, thick line).

Given that average effort levels are falling almost monotonically, yet output is single-peaked, what is happening to the output shares in this firm? Figure 27 shows what happens in the base case.



Figure 27: Typical evolution of agent income by agent type over the firm life cycle

Although total output is rising from time 10 to 20, the number of agents in the firm is also rising, causing each agent's share to be progressively diluted. This situation is particularly problematical for the agents who joined the firm early on, who have high preference for income. Such agents garner little utility from the opportunities for leisure that are rational to pursue in a firm dominated by free-riding, and thus can be expected to be the first to flee such a firm once decline sets in. However, the leisure-loving, low θ agents who flock to this firm over time find falling income to be only a minor problem as they bask in high leisure levels. Overall, utility levels drop among the agents in the firm over time, in a way that mirrors the fall in output share, as shown in figure 28.



Figure 28: Typical evolution of agent utility by agent type over the firm life cycle

Here, minimum and maximum utilities are shown as well. Note that the lowest utility level did not significantly change up through time 21, that is, up through the time at which the firm began declining.

3.6 Dynamics of Individual Agents

In the previous section the dynamical behavior of a typical firm was investigated. Here a similar analysis is rendered for a typical agent. Over time, each agent adjusts its effort level and consistently considers the opportunities that are available with 'neighboring' firms. If such opportunities dominate those at the firm where it is currently employed it changes jobs. So each agent's effort level and employer are changing over time as a result of its purposive behavior. Too, its income and utility change, as a result of both its own behavior as well as that of its co-workers, who are constantly adjusting their behavior as well. Overall, the exact way in which any individual's behavior changes is a very complicated function of its strategic environment. In this section we attempt to paint, in broad brush strokes, a portrait of these dynamics. It is hoped that this analysis is both illustrative of the kind of analysis that can be performed with computational models of firms in addition to being of interest in its own right.

To accomplish this analysis agents are simply selected at random and various of their state variables are displayed as a function of time. After viewing many of these plots a specific agent is deemed to be, in strictly heuristic and qualitative terms, representative of the population as a whole. In particular, figures 29-32 give various data on an agent having $\theta = 0.70$, sampled over some 30 periods, from time 52 to 82. In figure 29 a typical career path is shown, where each step represents a job change by the agent in question.



Figure 29: Career path over 30 periods of an agent for whom $\theta = 0.7$

This agent changes jobs nine times during this period. It spends approximately 3.5 periods with each firm on average, with a standard deviation of 2.5. Its longest tenure at any firm is 10 years, while it has several 1 and 2 year tenures. At time 76 it started up a new firm, acting as the founder, but just three periods later left for greener pastures in a rapidly growing eight-agent firm.

It is impossible to discern from a plot such as figure 29 the exact reason why an agent decides to change jobs at a particular time, or just why it stays put on another occasion. The details that go into each agent's decision are highly context dependent—who else is in the agent's firm, what is the average effort level there?—and idiosyncratic, depending intrinsically on its preferences. That is, the agent's decisions are explicable only by reference to the group in which it finds itself. Therefore, subsequent figures attempt to describe the context for the agent of figure 29.

Figure 30 illustrates how this agent's effort changed over this same course of time. If working alone this agent would, according to equation (6), put in effort of 0.77. From the figure it can be seen that the agent starts out working well below this level, then increases its input variously, including when it starts up its own firm (t = 76). It turns out that the agent works in relatively large firms during the first periods shown, and in relatively small firms during the later periods.

1

0.8

0.6

0.4

0.2



Time 55 60 65 70 75 80 85

Figure 30: Effort level history over 30 periods of an agent for whom $\theta = 0.7$

Overall, the average effort level of this agent is 0.61 with a standard deviation of 0.18.

Because the sizes of the firms in which this agent works changes significantly over this period, and given the significant effort level changes shown in figure 30, it is reasonable to expect large income swings for this agent over this period. Income data is shown in figure 31.



Figure 31: Income history over 30 period of an agent for whom $\theta = 0.7$

The agent receives fairly high income while working at a large firm initially. During the period 55 to 65, in which the agent does not change jobs, its income first falls off quickly, then recovers, and then dives again. In its next job the agent finds high income initially but then falls off, and there results a saw-tooth pattern subsequently. Finally, the agent receives its highest income at the end of the period in question, nearly 2 units.

Since this agent has a relatively high preference for income, it is reasonable to expect that its utility will more-or-less mirror its income. The actual utility time series is shown in figure 32.



Figure 32: Utility history over 30 periods of an agent for whom $\theta = 0.7$

Qualitatively, the peaks and troughs of figures 31 and 32 are quite similar. Looking at this figure one can learn, to some extent, why the agent may have switched jobs as shown in figure 29. For example, each of the agent's job switches in the t = 70-80 period are correlated with significant decreases in utility. It is harder to infer the agent's reasons for switching jobs, as well as not switching, in the early periods. For example, it would seem that the agent's large loss of utility in the period t = 57-58 would lead it to migrate to another firm. However, this does not happen. There are two possible reasons for this: (1) the agent, when activated, cannot find a better opportunity at another firm, and (2) perhaps the agent is simply not activated—it does not consider its options—during this period. A similar situation arises in the period of the early 60s. The agent could probably find gain utility by working elsewhere but does not change jobs until t = 65.

Time series plots of the type shown in figures 29-32 are very useful in developing some understanding of an agent's behavior. A variety other data would also be helpful in explicating that agent's actions, such as the behavior of other agents in the agent's firms over this period.

3.7 Overall Behavior of Agents, I (Population Cross-Section)

The thirty year excerpt from the agent career described in the previous section is a longitudinal picture of typical agent dynamics in the midst of firm formation, growth and dissolution. As such, it is rich in the contingencies of time and place, details that are always atypical in the same way that any realization of a stochastic process is idiosyncratic. A different picture of agent behavior arises from a cross-sectional perspective. Here, the average behavior of individual agents is determined at a particular instant in time. The resulting 'snapshot' of the agent population enriches our understanding of agent behavior, quantifying typical agent activity to the neglect of dynamics. To develop such a picture the population is first divided into a number of preference 'bins'; between 10 and 20 such bins are used here, depending on the variable in question. The model is run forward in time until the initial transient has worked its way out of the system and a stationary firm size distribution has obtained. Then, statistics on agent effort, income, and utility levels are computed for each bin, and data on agent tenure and firm size are also obtained. This is done several times and the snapshots compared for representativeness. Typical results are reported below.

We begin this analysis with a look at overall effort levels in the population. We know from the time series presented in figure 13 that the average effort fluctuates around 0.3. Figure 33 depicts the distribution of effort in the agent population.





Note that some 2/3 of the agents put in little or no effort, and that the remaining 1/3 of agents are split roughly equally between the remaining bins.

A different way to look at this same data is to disaggregate by preferences. Now, to the agents preferences are unobservable, but they are perfectly observable to us. Figure 34 gives a cross-sectional picture of agent effort levels as a function of θ , with the average value shown together with error bars representing one standard deviation.



Figure 34: Effort levels by agent type; averages ± 1 standard deviation

As expected, agents having the smallest θ put in little or no effort; actually, the mean value in each of these preference bins is strictly greater than 0, but in the first three bins most of the agents put in no effort at all. Effort levels are nearly monotone increasing in θ . The dip around $\theta = 0.70$ is not a robust feature—other snapshots do not commonly display it, although they do typically feature deviations from strict monotonicity. Agents having intermediate preference for income typically put less than 0.20 units of labor into production, and many put in no effort at all. It is only agents with relatively large income preferences who put in large effort levels; only for $\theta > 0.80$ is zero effort level rare. Note also that this data is rather heteroskedastic, that the largest variances occur at medium to large θ s, and as θ approaches 1 the variance shrinks.

Turning now to agent income, its distribution in the population is shown in figure 35 in semi-log coordinates. Interestingly, income is rightskewed.



Figure 35: Distribution of income in the agent population

Since the tail is approximately linear in these coordinates, income is exponentially-distributed beyond about 1.25. Incomes below this level occur with very nearly uniform frequency.

Who are the high income individuals? Is it individuals having high preferences for income? To find this out it is necessary to disaggregate by preference. Consider the data shown in figure 26. There are two prominent features in the data. First, income increases approximately linearly, and essentially monotonically, with preference for income. Second, the variances are nearly constant across preference bins.



Figure 36: Income by agent type; averages ± 1 standard deviation

Now, since income is only weakly related to effort levels, and is homogeneous within firms—all agents in a firm receive the same

income—the increasing income with preferences must derive, at least in part, from high θ agents working in more productive firms.

Next, consider utility. Effort and income combine to produce utility for each agent. The distribution of utility that obtains in the base case is depicted in figure 37.





This data is right-skewed and would perhaps be well fit by a lognormal distribution. Although not apparent in the figure, some agents have utility as high as 4.

We can anticipate how agent utility will depend on preferences by noting that since utility is monotone increasing in income for all agents, figure 36's monotonically-increasing income with θ suggests that utility should be increasing with agent preferences as well. Figure 38 demonstrates that this is so, although utility is a weaker function of θ than is income.



Figure 38: Utility by agent type; averages ± 1 standard deviation

As with income, variances are approximately constant across preferences. The relatively large variation in utility levels within preference bins suggests that agents having the same preferences can experience quite different welfare levels. Similarly, it is not uncommon for a low θ agent to have higher utility than one with high θ , notwithstanding the overall rise of average utility with preference for income.

So far, in all depictions of agent utility, such as figures 16, 28, 32, 37 and 38, it is the *realized* utility that has been shown. That is, each period every firm engages in production and then divides output among the agent workers, and this income combines with agent leisure to produce utility realizations to the agents. However, a somewhat different measure of agent welfare is the utility level that agents employed in making their most recent effort input decisions. This quantity is the direct result of each agent weighing its options of either staying with its current firm or joining a different firm. Call this an agent's decision utility. In general it will be different from the realized utility insofar as a firm engages in production at some time after each of its agents has made its most recent decision. Sometimes an agent's decision may be soon followed by firm production, in which case the two types of utility are highly correlated. At other times many agents may have the opportunity to make their decisions in between the time at which a particular agent decides and production actually occurs. In such circumstances the relationship between realized and decision utilities will be weaker.

Figure 39 is a cross-section of agent decision utilities for the same run of the model that yielded figure 38's realized utilities. It is quite similar in character to the previous figure. However, note that decision utility levels tend to be greater than realized utilities for the same range of θ . Eventually, decision utility shall serve as a basis for welfare analysis of this model.



Figure 39: Decision utility by agent type; averages ± 1 standard deviation

Note that there is little variance in decision utility for agents having small θ .

The number of periods that an agent has been with a particular firm is called its *tenure*. Agents new to a firm have no tenure, while agents who have spent one period have tenure one, and so on. Given what we know about agent careers from the previous section, the average tenure is relatively short in this model. How might this quantity change with agent preferences. We have seen that large firms can become 'infected' with low θ agents, and so this suggests that leisure-loving agents are wont to jump firms often, leading to low tenure, while income-lovers might do the opposite. The tenure cross-section is shown in figure 40.



Figure 40: Tenure by agent type; averages ± 1 standard deviation

Agents having high θ have slightly longer tenure than low θ agents. There is significant variance within each bin, demonstrating the high levels of agent turnover within firms that exist in this model.

As a final measure of the agent population cross-section, firm size as a function of agent type will be studied. Given that income is increasing with θ , it is tempting to think that average firm size may be increasing in θ as well, since the advantages of increasing returns in large firms should lead to higher incomes. Data from the model are shown in figure 41.



Figure 41: Size of an agent's firm, by agent type; averages ± 1 standard deviation

There does exist a general trend toward rising average firm size with θ , although this is rather irregular. Large variances are characteristic of all of these data. Therefore, even though an agent having $\theta = 0.9$ is likely, on average, to work in a larger firm than an agent having $\theta = 0.1$, there is significant probability that the situation is, in fact, the opposite.⁵⁸

3.8 Overall Behavior of Agents, II (Cross-Section by Firms)

The cross-sectional data just described were with respect to the agent population. A somewhat different portrait of agent behavior results from looking at these same variables—preferences, effort, income, utility—within firms and then binning the data with respect to the firm population. This is

⁵⁸ Note that the average firm size in figure 40 is substantially greater than the true average firm size given, for example, in figure 12. The figure 40 mean is weighted by the population, i.e., each firm is counted multiple times, by all the agents who compose such firms.

too a cross-sectional analysis. It permits us to characterize not only the average behavior within firms but also the variation in behavior.

Let us start with preferences, which are uniformly distributed in the population. If they were uniformly distributed *within firms* then it would be the case that each firm would have mean θ of 0.5 and a histogram of θ by firm would have all the mass in the center bin. However, given the power law character of the firm size distribution there are lots of small firms and so we would expect the average θ in a firm to vary substantially from a single spike at 0.5. This is confirmed in figure 42.



Figure 42: Histogram of average θ in firms

Note that this distribution is not symmetric, has mean somewhat greater than its median, and the mode is near 0.5. The standard deviation of the θ distribution is 0.20.

Within each firm there is some variation in the types of agents present. A measure of this variation is the standard deviation in θ , s_{θ}. For the population of firms that yielded the previous figure, the distribution of s_{θ} is given in figure 43.



Note that there are many firms with little or no variation in agent type. Presumably these are the smallest firms, consisting or but 1 or 2 agents. If agents were distributed more or less uniformly in firms across all agent types then $s_{\theta} = 0.288$. In fact, there is a peak in the distribution centered at $s_{\theta} = 0.275$, suggesting that agents are quite uniformly distributed in a significant number of firms. There are a modest number of firms—perhaps 10%—in which there is substantially more variation in preferences than occurs in the uniform distribution, suggesting that some firms may have bimodal or other unusual distributions. Finally, given that the inter-firm standard deviation in θ is 0.20, many firms have the attribute their intra-firm variance in preferences is less than the inter-firm variance.

Next we turn to the distribution of average effort levels, \bar{e} . In figures 33 and 34 we saw that substantial fractions of the agent population put little or no effort into production. Is it similarly true that a large proportion of firms are characterized by vanishing total effort levels? Intuitively, this would seem to be unrealizable, for agents would not continue to inhabit firms with $E \sim 0$. Rather, it seems more likely that while individual firms could harbor large numbers of free-riders, such agents must be working alongside hardworking agents in order to keep such firms viable. The implications of these considerations for the distribution of effort levels are that, while having a large mass of firms in the lowest effort level bin is not ruled out, it is also the case that a distribution much less extreme than figure 33 is expected. In fact, the same data set that yielded the two previous plots produces the distribution of effort levels shown in semi-log coordinates in figure 44.

0.5

0.2

0.1

0.05

0.02

0.2



0.6

0.8

Figure 44: Distribution of average effort levels in firms

0.4

The average effort in firms is 0.27, while the standard deviation in \bar{e} across all firms is 0.22. The frequency, f, is approximately linear in average effort levels, suggesting that \bar{e} is exponentially distributed.

Within each firm there is significant variation in effort levels, a quantity plotted in figure 45.





The shape of this distribution is very similar to that shown in figure 43 for the standard deviation in θ . A significant fraction of the mass is in the smallest effort bin, and presumably these are small firms formed of agents having small θ . Once again, there is an interior peak located quite near the standard deviation associated with a uniform distribution-here, of effort

ē

1

levels—as well as some small number of firms in which there is substantially more variance than that associated with a uniform distribution. For a substantial fraction of firms the intra-firm variance in effort levels is less than the inter-firm variation.

The distribution of income, *I*, across firms is quite similar to the distribution of average effort levels across firms (figure 45) as well as the distribution of income in the entire population (figure 36). In particular, in semi-log coordinates the income distribution is nearly linear, as shown in figure 46, suggesting that income too is exponentially distributed.



Figure 46: Distribution of income across firms

The OLS estimate of the parameter of this distribution is 2.30. Therefore, the probability of being in a firm having income *I* is proportional to exp(-2.30 *I*), meaning that firms having income I + 1 are exp(2.30) \approx 10 times less common than those with income *I*. Since all agents in a firm receive the same income there is no intra-firm income variation, and thus no histogram of such variations possible.

Finally we study the distribution of average utility received by agents across firms. From figure 38 we know that the utility increases very little with θ . Will things be different when the cross-section is performed with respect to firms instead of the population as a whole? In figure 47 we see that utility is apparently exponentially-distributed, in much the same fashion as are average effort and average income.


Figure 47: Distribution of average utility in firms

Relatively low utility levels are common. Presumably these occur in small firms. The concavity at small average utilities is a robust feature of this data, although the dip at the second datum is idiosyncratic. The irregularities at high average utility result from small sample sizes there.

The amount of variation in income also appears to be exponentially distributed, as shown in figure 48. In most groups there is very little variation in utility levels received by the agents who constitute the group.



Figure 48: Distribution of standard deviation in utility in firms

The break in the data just beyond $s_{u} = 1$ represents a single datum, and is thus explained as simply a random event due to small numbers.

Group Selection

The neoclassical justification for treating firms as profit maximizers hinges on a selection argument. Firms that make the largest profits are able to grow at the expense of the less profitable ones, or so the story goes.⁵⁹ This Darwinian picture is either a group selection argument, if one views firms as multi-agent groups, or not, insofar as a firm is considered to be a single agent. As alluded to in the introduction to this essay, the unitary actor model is rife with conceptual and other problems, while the group selection picture has been intellectually out-of-favor for some time.

However, the recent work of Wilson and Sober [1998], among others, has at least partially rehabilitated the idea of group selection.⁶⁰ Here we merely wish to point out that group selection ideas are very relevant to the model of firms described here. In particular, we have seen that firms having low per capita effort levels are selected against. The mechanism of selection is not biological—indeed, the number of agents in the population at any time is constant. Rather, the continual rearrangement of the composition of firms, through competition for workers, is the mechanism of selection. Firms who cannot attract productive workers eventually die. Firms with low levels of free-ridership succeed in attracting good workers and grow. In essence, in many firms there is greater variance in agent behavior (effort levels) between firms than *within* them—see figures 44 and 45—and so group selection can Thus, from the multi-agent perspective adopted here, manifest itself. successful firms are not profit maximizers, but rather are those who attract and retain productive workers. Profit maximization, if it exists at all, is a byproduct of individual utility maximization.

3.9 A Speculative Welfare Analysis

Since the microeconomic dynamics of this model are inherently outof-equilibrium, it is, ostensibly, very difficult to make any definite welfare statements about the observed stationary state. However, some general comments are possible.

Given that each agent, when activated, considers staying in its present firm or, among other things, starting up a new firm, and accepts the result that maximizes its utility, it is easy to see that agent utility is bounded from below by the singleton utility. Therefore, it must be the case that all agents prefer the non-equilibrium state to one in which each is working alone. That

⁵⁹ Friedman [1953] is a well-known account of this. Blume and Easley [1998] give a formal model and find that the neoclassical result applies only to self-financed firms. The presence of a capital market can act to sever the connection between profit maximization and survival.

⁶⁰ Bowles [1998] discusses the relevance of these developments for economics.

is, the state of the economy in which all firms are size one is Paretodominated by the dynamical configurations studied above. But this is a weak result. More important would be a result that utility levels in some equilibrium configuration were dominated by those in a non-equilibrium configuration, especially if this latter regime was stationary at the aggregate level. For then one could claim that the generic behavior of the model yielded outcomes superior to equilibrium in welfare terms. But do equilibrium configurations even exist in this model? If so, what are they? If there are multiple configurations of this type, which one is appropriate for making welfare comparisons?

Imagine if each group could exclude any agent who wishes to join.⁶¹ This is not the behavioral rule employed above, but is one that is easily implementable.⁶² Now, consider groups in which all agents have exactly the same preferences for income, i.e., homogenous groups. Next, let us say that all agents in the population are in homogeneous groups. Further, no group is larger than its maximum stable size. Associated with such groups are the utility levels shown in figure 7 above. Given the possibility of agents to exclude other agents, it is easy to see that this partition of the agent population into homogeneous, stable groups is an equilibrium configuration. For the only kind of agents that a group having preference of θ wishes to include within itself are those who have preference strictly greater than θ . To include an agent with lesser preference would dilute each agent's income and lower the utilities of all group members. But agents with greater θ are better off in optimally-sized groups of their peers than with agents having lower θ .

Figure 49 starts out as a recapitulation of figure 7: a plot of the optimal utility levels for both singleton firms as well as optimal size homogeneous ones, as a function of θ . Overlaid on these smooth curves is the cross-section of decision utilities from figure 39, now in semi-log coordinates. The main thing to notice about figure 42 is that most agents prefer the non-equilibrium world to the equilibrium outcome with homogeneous groups. Indeed, it is only agents for whom $\theta > 0.8$ who prefer to work exclusively with agents having similar preferences.⁶³

⁶¹ For analytical characterization of an equal share (partnership) model with perfect exclusionary power see Farrell and Scotchmer [1988]; an extension to heterogeneous skills is given by Sherstyuk [1998].

⁶² Indeed, it will be systematically studied in § 4.9 below.

⁶³ The error bars in the figure seem to indicate that some agents get less utility in the heterogeneous groups than they do working alone. This is not true, of course, but rather is an artifact of binning the data.



Figure 49: Utility derived from working in single agent firms, in optimal size homogeneous firms, and in non-equilibrium multi-agent firms, by agent preference

Interestingly, while the non-equilibrium configuration does not dominate the equilibrium one, it is also true that dominance does not run in the reverse direction. Therefore, these two alternative economic worlds cannot be Pareto ranked.

4 Variations and Elaborations

In this section the parameterization of the model that has been employed so far is systematically altered. All of the model attributes of the base case configuration described in table 2, as well as some others, are varied and the overall effect on the resulting population of firms is described. For each of these variations we could present all of the data provided in § 3, information on firm size distributions (size by number of agents and output), productivity, growth rates, and firm lifetimes. We could then go on to develop a picture of the firm life cycle, agent careers, and cross-sectional analysis of the agent population, comparing each variation with both the base case as well as other variations. Clearly, this would be a monumental task.⁶⁴ In lieu of this, each variation of the model will be characterized by a single parameter, the power law exponent, μ , of the firm size distribution, with size defined as the number of agents. This is a crude statistic, but one that reveals much about the general effect of the variation in question nonetheless.

⁶⁴ Although this presents no computational problems. Indeed, each parametric variation of the model essentially develops all of this information. Nor would the cost be prohibitive of storing this information in a database, thus permitting run-to-run comparisons.

Three attributes from the table 2 base case were found to have no effect on the general character of the model. First, in moving from 1000 to 5,000 agents, and then to 10,000 and eventually to 100,000 agents and beyond, the primary effect is to make the resulting statistical character of the model more robust, simply due to larger sample sizes. The quantitative characterization of the model described in § 3 above is invariant to this change.⁶⁵ The second attribute that yielded little effect was the nature of agent activation. In other agent-based models synchronous activation, in which all agents update their state simultaneously, has proven to produce significantly different output than asynchronous activation (cf., Huberman and Glance [1993]). Further, with the asynchronous activation model it has been found that it matters whether agents are activated randomly or uniformly-i.e., agents activated at random vs. each agent active exactly once per period, although perhaps in a random order each period (Axtell et al. [1996]). Here we have worked exclusively within the asynchronous activation model and have found essentially no difference in firm size distributions between uniform and random activation. Finally, the way in which the model was started—each agent working alone in the base case—did not seem to matter much beyond a few dozen periods, i.e., the initial transient quickly decays. Other initial conditions tested included random groups and one large group.

In what follows, § 4.1 investigates the importance of locally purposive behavior by permitting the agents to behave randomly in certain ways. The second section varies the increasing returns parameters, *b* and β . Disparate distributions of agent preferences are studied in § 4.3, while in § 4.4 the extent and composition of agents' social networks are modified. Then, in § 4.5 the idea that an agent might be loyal to its firm is introduced. Section 4.6 looks at the effect of further bounded rationality, through effort level adjustments by groping, instead of explicit calculation. Alternative compensation schemes,

⁶⁵ The primary practical effect of increasing the number of agents is to diminish the productivity of the computations. That is, while increasing the number of agents by an order of magnitude increases the sample size each period by this same amount, and thus requiring an order of magnitude fewer model periods in order to reach the same total sample size, it comes at a cost of a substantially greater amount of real time-it just takes longer! The reason for this is that small models can 'live' on the high-speed caches that work in tandem with the current generation of microprocessors. For example, if each agent object is O(100 bytes) and if each firm object is about the same size, then if the average fir m size is, say, 5, then the total amount of memory required for these components of the model, M, is 100 (A+A/5) bytes; for A = 1000, M = 120,000 bytes. Double these values for various operating system overhead and one finds that the base case model can survive entirely on a 256 kilobyte cache. One megabyte caches are about the largest found today so this puts an upper limit of about A = 4000 for rapid model execution. Experience suggests that by the time one has ventured to the A = 10,000 level a performance penalty in the range of 100-200% obtains, while for A = 100,000 a 10x slowdown is not uncommon. The caches are still of considerable help in the intermediate size model, but of little value for large populations, i.e., M » cache size.

beyond simple sharing, are investigated in § 4.7. In § 4.8 firm founders act as residual claimants: agents are paid a fixed wage and the founder allocates bonus pay based on firm output. Section 4.8 looks at firm hiring standards. Issues of intra-firm monitoring are discussed in § 4.10. Finally, § 4.11 investigates simultaneous model variations.

4.1 The Importance of Locally Purposive Behavior

Against this simple model of firm formation it is possible to mount a critique of the following type. Since the Gibrat law is known to yield power law distributions of firm size, and since growth rates in the Gibrat law are random variables, then perhaps the model described above is simply a complicated way to generate random movement of agents between groups. Stated differently, although the agents are behaving purposively, this may be little more than noise in this strategic environment. Of course, at the *aggregate level* something like this must be true insofar as right skewed size distributions result. But the important question for the model is 'What if agent behavior were truly random at the *micro-level*, would this too yield power law size distributions?'

We have investigated this in two ways. First, imagine that agents randomly select whether to stay in their current firm, leave for another firm, or start-up a new firm. If they choose to leave then the firm they migrate to is chosen randomly. However, we suppose that each agent selects the optimal effort level given its 'decision' on which firm to inhabit. It turns out that this model fails to yield a power law size distribution for firm size. In fact, firms greater than 9 or 10 are rarely observed in this variation. Similarly, if agents select the utility-maximizing firm in which to work—we assume they are able to divine this—but then choose an effort level at random, again nothing like power law size distributions result. These changes in the behavioral specifications of the model suggest that any systematic departure from (locally) purposive behavior is unrealistic.

4.2 Effect of the Increasing Returns Parameters

The two parameters that determine the extent of increasing returns in (2') are *b* and β . In all of the above they have been set to *b* = 1 and β = 2. In this section we systematically vary β and study its effect both on the overall firm size distribution, as measured by power law exponent μ , and on γ , the scaling law exponent for the dependence of standard deviation in log growth rates on size. Results are summarized in table 3; the shaded entry refers to the base case studied above.

β	μ	γ
1.7	2.06	0.238
1.8	1.62	0.237
1.9	1.32	0.204
2.0	1.28	0.174
2.1	0.95	0.173
U[1.7, 2.1]	1.12	0.203

Table 3: Dependence of the scaling law exponent, μ , on the increasing returns exponent, β

The pattern of dependence of μ on β is clear. As the effect of increasing returns is made stronger, larger firms can grow and survive, and therefore μ increases (becomes less negative). For $\beta < 1.7$, the firm size distribution that results is too concave to be fit well by a simple power law.⁶⁶ For $\beta > 2.1$, very large firm sizes are possible, and so computational limitations are encountered.⁶⁷ Average firm size increases with β .

Next, the effect of the coefficient *b* in (3) is studied. The case of b = 0 corresponds to constant returns, in which case there are no advantages to be had from agents forming firms. So we estimate the power law exponent, μ , for various b > 0. Results are summarized in table 4 below.

b	μ
0.50	2.09
0.75	1.33
1.00	1.28
1.25	0.91
1.50	0.53
U[0.50, 1.50]	0.89

Table 4: Dependence of the scaling law exponent, μ , on the coefficient *b*, of the increasing returns term

For b < 0.5, the firm size distribution is too concave to be well-fit by a simple scaling law, as was the case for $\beta < 1.7$ in table 3. For b > 1.5, the maximum firm size is so large as to present computational constraints, also in similar

⁶⁶ Concavity is also a feature of the empirical data, so it may be that β s of this magnitude are not unrealistic.

⁶⁷ In order for keep the computational model from running away to a single large firm, a large population of agents must be instantiated for β s of this magnitude, and it is not practical to simulate such large models on workstations today.

fashion to table 3 (large β). Average firm size is also an increasing function of *b*.

The case of $b \sim U[0.50, 1.50]$ is of particular interest. Here each time an agent founds a firm it draws a *b* from this distribution. Note that the power-law exponent for this case behaves much more like b = 1.25, not b = 1.0, the mean of the distribution. This is to say that the kinds of firm that exist at any time have a preponderance of large *b*'s.

4.3 Alternative Specifications of Preferences

Preferences are distributed uniformly on [0,1] in the base case. This yields a certain number of agents having quite extreme preferences: those with $\theta \approx 0$ are leisure lovers and consider income superfluous, while those with $\theta \approx 1$ are income lovers and spend little time in leisure activities. In this section a variety of alternative populations of preferences shall be studied. In particular, we shall first look at the effect of removing the agents with extreme preferences from the population, by truncating the uniform distribution, yielding a distribution with less variance but the same mean. Next, uniformity gives way to a single-peaked distribution of triangular type. Then, the effect of a skewed distribution is studied. Preferences are made somewhat 'smoother' subsequently through use of a truncated normal distribution. Then, measures having special properties—the beta and Dirac delta—are used. Finally, the Cobb-Douglas specification of preferences is relaxed in favor of the CES functional form, i.e.,

$$U^{i}(e_{i};\rho_{i},\delta_{i},E_{\sim i},N) = \left\{\delta_{i}\left[\frac{O(e_{i};E_{\sim i})}{N}\right]^{\rho_{i}} + (1-\delta_{i})(1-e_{i})^{\rho_{i}}\right\}^{1/\rho_{i}},$$

where each $\delta_i \in [0, 1]$ and $\rho_i \in [-1, \infty)$. For $\rho_i = -1$ the utility function is linear in effort, while in the limit of ρ_i approaching 0 Cobb-Douglas behavior obtains (cf. Varian [1984: 30]). One of the variations reported below has ρ_i distributed uniformly on [-1, 0], so that agents have a variety of utility function specifications. As ρ_i gets large, preferences take on a Leontief character, and a case of this type is also studied. The results of these variations are summarized in table 5.

Distribution	μ
C-D with θ ~ uniform on [0, 1]	1.28
C-D with θ ~ uniform on [0.25, 0.75]	1.21
C-D with θ ~ triangular on [0, 1] with mode 0.5	1.31
C-D with θ ~ triangular on [0, 1] with mode 0.75	1.01
CD with θ ~ truncated normal on [0, 1] with variance 1/2	1.30
C-D with θ ~ beta on [0, 1] with parameters 1 and 2	0.99
C-D with θ ~ Dirac delta at 0.75	0.91
CES with δ ~ uniform on [0, 1], ρ ~ uniform on [-1, 0]	1.56
CES with δ ~ uniform on [0, 1], ρ ~ uniform on [0, 10]	1.26

Table 5: Dependence of the scaling law exponent, μ , on the distribution of agent preferences

Overall, the general power-law character of the size distributions remain, although changes in the distribution of preferences have significant effect. CES preferences yield power law exponents that are comparable to those previously obtained, leading to the conclusion that the general character of the model does not depend sensitively on the functional forms employed.

4.4 Effect of the Extent and Composition of Agent Social Networks

In all of the above, each agent had only two friends. The number of friends is a measure of an agent's search or information space, since the agent queries these other agents each period to assess the feasibility of joining their firms. In this section we want to determine the effect of increasing the extent of the agents' social networks.

To do this we simply assign each agent some number v > 2 friends at time 0. Then, when determining its optimal effort level, an agent will consider each of these friends' firms as a possible place to work.⁶⁸ In particular, the number of friends was varied from 2 to 10, and the power law parameter in the firm size distribution function estimated. The results are summarized in table 6.

⁶⁸ Computationally, this slows down execution of the model in proportion to v.

V	μ
2	1.28
4	1.11
6	1.08
8	1.08
10	0.99
U[2, 10]	1.11

Table 6: Dependence of scaling law exponent on the size of social networks, where social networks are composed of agents (friends), specified exogenously

Thus, the overall effect of increasing the size of agent social networks is in the same direction as raising the increasing returns parameter, β . That is, it tends to stabilize large firms, although this effect is relatively weak. Essentially, as V increases the system becomes more "fluid" and agents can better seek out gains from cooperation. Note that the effect of having heterogeneous V is that the system behaves a little more like a society with a smaller network. Presumably this means that agents with large social networks have very little advantage over agents with smaller networks. What matters is that one's social network is 'large enough'.

Next, instead of the social network consisting of exogenously specified agents, here it consists of randomly selected firms. Results are shown in table 7.

V	μ
2	1.28
4	1.22
6	1.18
8	1.07
10	1.03
U[2, 10]	1.02

Table 7: Dependence of scaling law exponent on the size of social networks, where social networks are composed of randomly selected firms

These results are analogous to those of table 6. As social networks get larger somewhat larger firms can be supported. Note that when agents have social networks of various sizes the overall behavior of the model is similar to when agents have relatively large networks. That is, agents with large networks seem to be primarily responsible for the overall structure of the firm size distribution.

4.5 Effect of Agent "Loyalty"

In the basic model, agents are local optimizers, doing the best they can for themselves without regard for others around them. One way to make them more socially aware is to give them loyalty to their firms, where by loyalty we mean that an agent does not move to a new firm even though it has determined that it would be better off by doing so.⁶⁹ Formally, call λ the maximum number of times an agent an computes that the best thing for it to do is move elsewhere, but it does not do so. Once the number of such assessments exceeds λ then the agent, in fact, moves, and resets its λ counter. That is, it now displays loyalty to its new firm.

The setting $\lambda = 0$ corresponds to the base model. Various other values for λ have been experimented with and the effect on the firm size distribution exponent, μ , measured. These results are summarized in table 8.

λ	μ
0	1.28
2	1.14
5	0.85
10	0.77
U[0, 10]	0.79

Table 8: Dependence of scaling law exponent on agent loyalty

Note that increasing λ produces more large firms. That is, loyalty is a stabilizing factor for large firms, in accord with intuition. This is a relatively strong effect. Furthermore, when the agent population is heterogeneous the agents with little loyalty are incapable of destabilizing large firms.

4.6 Bounded Rationality: Groping for Better Effort Levels

In all of the above agents were able to adjust their effort levels to anywhere within the feasible range [0, 1] instantaneously. That is, when an agent was active, it solved an unconstrained optimization problem. Here we explore the notion that agents used to working with effort *e* will make only small changes in *e* each time they are activated. It is as if there was some

⁶⁹ Loyalty is a prominent feature in other agent-based computational models. Tesfatsion [1998] utilizes a notion of worker-employer loyalty in a model of labor markets. In Kirman and Vriend [1998], loyalty between buyers and sellers emerges in bilateral exchange markets.

prevailing cultural norm or *work ethic* which constrains the agents to keep doing what they have been, more or less, despite what is going on around them.

In particular, we require that each agent searches for its utilitymaximizing effort level over a range of 0.10 around its current effort level. That is, an agent working with effort e_i picks its new effort level from the range $[e_i, e_i]$, where $e_i = \max(0, e_i - 0.05)$ and $e_i = \min(e_i + 0.05, 1)$. The overall effect of such sticky effort level adjustment dynamics is given in table 9, where the firm size distribution exponent, μ , is shown as a function of the increasing returns parameter, β .

β	μ	$\mu_{ m sticky}$
1.7	2.06	1.54
1.8	1.62	1.25
1.9	1.32	1.17
2.0	1.28	0.92
2.1	0.95	0.95

Table 9: Dependence of scaling law exponent on stickiness of effort adjustment, various β

The effect is similar to adding agent loyalty, that is, to produce larger firms. The reason for this is intuitively clear. As large firms tend toward non-cooperation, sticky effort adjustment puts the brakes on the downhill spiral to complete free riding.

Next imagine that agents can only sample a single effort level different from their current level and either accept it as the new level or else maintain their previous effort. Call this process effort level *groping*. Table 10 summarizes the effect of this restriction on search on the power law exponent.

β	μ	$\mu_{ m groping}$
1.7	2.06	1.69
1.8	1.62	1.49
1.9	1.32	1.33
2.0	1.28	1.19
2.1	0.95	1.00

Table 10: Dependence of scaling law exponent on stickiness of effort adjustment, various β

The effect is similar to before, just more pronounced. The firm size distributions that result from effort level groping are characterized by fewer smaller firms and greater numbers of large firms.

4.7 Alternative Compensation Schemes

So far all agents in a group have shared total output equally. Here we study a variety of alternative division rules.⁷⁰ The class of rules we investigate involves compensation formulas in which the agents who have been with the firm the longest receive relatively larger shares. Several variations of this are studied. The next section studies a division rule in which each agent receives some fixed income, with equal bonus shares based on overall firm performance. The common effect of these two variations is to help stabilize large firms.

It is an easy matter computationally to keep track of the seniority of the agents within a firm. Thus, it is easy to tie agent compensation to seniority—one needs only a formula for doing so. Call $s_i(O)$ the share of total output that is agent *i*'s, where *i* is an index of the agent's seniority, with *i* = 1 referring to the firm founder, or at least the agent with the greatest seniority in the event that the actual founder has left the firm. Define g(i) as an output division function such that

$$s_i(O) = \frac{g(i)}{\sum_{i=1}^N g(i)} O.$$

We shall here consider a class of 'hyperbolic' compensation policies, division rules which have the general character of allocating more to agents having greater seniority.⁷¹ As an example, consider $g(i; p) = p^{-i}$. Each share in such a system is progressively smaller by 1/p. Results are summarized in table 11, below.

⁷⁰ Encinosa *et al.* [1997] have studied compensation systems empirically for team production environments quite similar to the one considered here. They find that "group norms" are at least as important as conventional economic considerations in determining pay practices in medical groups. Garen [1998] empirically links pay systems to organizational form, finding that monitoring costs largely determine the type of system used.

⁷¹ It might be argued that compensation functions of this type would have been a better choice for our model's base case, insofar as they are more realistic than the equal division one. However, given that these compensation functions are essentially 'power law' in character, the achievement of power law firm size distributions could have been viewed as a possible modeling artifact—a spurious result induced by the choice of compensation function. For this reason the uniform division compensation rule was used in the base case.

g(i)	μ
5- <i>i</i>	1.11
4- <i>i</i>	1.03
3- <i>i</i>	0.96
2- <i>i</i>	0.89
$1^{-i} = 1$	1.28
2-(<i>i</i> +1)	0.99
2-(<i>i</i> +2)	1.07
2-(<i>i</i> +3)	1.04

Axtell

Table 11: Dependence of scaling law exponent on hyperbolic compensation functions

Note that the overall effect of hyperbolic compensation, in comparison to uniform compensation, is to make large firms somewhat more stable. But notice that this effect is not strong.

4.8 Firm Founder Acts as Residual Claimant

Bonus compensation is an important feature of real compensation systems. It is usually coupled to some form of base compensation. Here the equal output shares model is modified by breaking up total compensation into base and bonus amounts, where the latter is what the founder agent allocates to each agent once base compensation is paid. In essence, the equal shares model represents the limit of having 100% of each agent's compensation variable. Call Φ^i the fixed (base) part of income an agent receives. Therefore, the bonus each agent will share is

$$\max\left[0, \frac{O(e_i; \theta_i, E_{\sim i}, N) - \sum_{i \in A} \Phi^i}{N}\right].$$

Overall, each agent is faced with optimizing a utility function that now has the form

$$U^{i}(e_{i};\theta_{i},E_{\sim i},N) = \left(\Phi^{i} + \frac{O(e_{i};\theta_{i},E_{\sim i},N) - \sum_{i \in A} \Phi^{i}}{N}\right)^{\theta_{i}} \left(1 - e_{i}\right)^{1-\theta_{i}}.$$
(21)

The first order conditions for this problem can be written down but are of little use here since the Nash equilibrium of the model is certain to be unstable. Rather, given that each agent is now motivated by (21) we simply spin the model forward in time and study the firms that self-organize.

Base Compensation	μ
0	1.28
20% of singleton income	1.26
50% of singleton income	1.06
80% of singleton income	0.85
singleton income of median agent	0.99
singleton income of mean agent	1.01

Table 12: Dependence of scaling law exponent on type of base compensation

Note that the overall effect of bonus compensation is indeed to stabilize large firms somewhat.

Although tables 11 and 12 demonstrate that compensation rules have systematic effects on the overall stability of firms and the resulting size distribution, it is also true that these effects are relatively small and that the 'equal sharing' compensation scheme is a reasonable starting point for analysis.

Here the residual claimant, in the form of the firm founder, acts only to divide up the firm's surplus over base wages. A richer model would permit the founder to save some of the surplus for future periods, use it to invest in or buy-out another founder, or just keep it for itself. Furthermore, agents would have some expectations as to the disposition of the surplus and would make their effort level adjustments accordingly. Here there are many paths which one might fruitfully with the model, and perhaps these will be the subject of future work.

4.9 Firm Founder Sets Hiring Standards

Perhaps the most unrealistic aspect of this model is that agents can join any firm that they deem to be a better opportunity for them. Here this specification of the model is varied, by having the firm founder put hiring standards in place. In particular, when an agent finds a firm it wants to join it first queries the firm founder, who determines whether or not the new job candidate has a preference for income that is at least $\phi\%$ of its own—income preference is a surrogate for how hard an agent will work.⁷² Results are shown in table 13 below for ϕ varying from 0 to 100%.

$ heta_{\iota}$	μ
$\geq 0\% \ \theta_{\rm founder}$	1.28
$\geq 20\% \ \theta_{\text{founder}}$	1.27
$\geq 40\% \ \theta_{\text{founder}}$	1.22
$\geq 60\% \ \theta_{\text{founder}}$	1.03
$\geq 80\% \ \theta_{\text{founder}}$	1.17
$\geq 100\% \ \theta_{\text{founder}}$	*
\geq U[0%, 100%] θ_{founder}	1.13

 Table 13: Dependence of scaling law exponent on firm target output level

The (*) indicates that the resultant size distribution is not well described by a power law. Note that the overall effect of hiring policies of this type is to first stabilize large firms to some extent (smaller μ), and then to destabilize them (larger μ), although the value of μ for the $\phi = 80\%$ case is perhaps not dependable as this size distribution also departed significantly from power law behavior. The firms in this parameterization tend to grow more slowly, mostly due to the difficulty of finding agents who satisfy the hiring criteria.

4.10 Effort-Level Monitoring within Firms

An unrealistic aspect of this model is that shirking goes completely undetected and unpunished. Such matters are of crucial importance in actually-existing firms, and large literatures have grown up around this question; well-known work includes Olson [1965]. More recent work includes the models of mutual monitoring of Bowles and Gintis [1998] and Dong and Dow [1993b], the effect of free exit (Dong and Dow [1993a]), and endowment effects (Legros and Newman [1996]). Ostrom [1990] describes mutual monitoring in institutions of self-governance that have arisen for managing resource problems. property Getting realistic intra-firm common organizational structures to emerge is an active area of research within the multi-agent computational approach to firm formation and evolution.73

⁷² An alternative way to view the base model, in which agents can join any firm they wish, is that agent type is unobservable and, therefore, job applicants cannot be filtered by preference.

⁷³ Allen [1997] is working on evolving agent-agent monitoring structures by which agents can detect shirking.

However, perhaps one lesson that can be drawn from the results in table 13 is that, while it may be the objective of the management of firms to have *perfect* monitoring of employees, when monitoring is perfect the empirical size distribution fails to obtain.⁷⁴ Stated differently, perhaps all firms suffer, to a greater or lesser extent, from imperfect monitoring, and therefore the creation of economic models in which perfect monitoring obtains in equilibrium is a kind of quixotic undertaking, for which the only possible outcome can be disagreement with empirical data. Indeed, many real-world compensation systems can be interpreted as devices for managing incentive problems by substituting reward for supervision, from efficiency wages to profit-sharing (Bowles and Gintis [1996]). Furthermore, if incentive problems in team production environments were perfectly handled by monitoring then there would be little need for the corporate law (Blair and Stout [1997]).

4.11 Putting it all Together: Compound Variations

So far each variation of the model has been undertaken as a single departure from the base case. We have systematically explored the neighborhood of the base case but each time returned to the safe haven from whence we departed before striking out anew. Here we simply note that it is possible to proceed somewhat differently. Imagine that multiple attributes from the table 2 base case are varied at once. For example, a model quite distant from the base case would result if the coefficients of the production function are selected to be different from unity, the increasing returns exponent is different from 2, the number of neighbors is set to 10, loyalty is uniformly distributed on [0, 10], a hyperbolic compensation policy is invoked, and we start with 10,000 agents in random groups. Given the modest number of variations described in § 4 alone it is possible to construct nearly 200 million distinct models.⁷⁵

Imagine now that some number of such variations had been executed, and for each the power-law exponent of the firm size distribution, μ , estimated. It would then be possible, in principle, to determine how μ depends on the model parameterization. But how would such a model be specified? In general, the way in which μ depends on *a*, *b*, β , the distribution of θ , loyalty (λ), and so on is surely *very* complicated. A naive specification would write μ as a linear function of these parameters. But certainly this is a misspecification. Indeed, we have employed agent-based computation as an engine of aggregation, precisely to surmount the many and deep difficulties

⁷⁴ The table 13 results on hiring restrictions are not directly about monitoring, but are similar in spirit since filtering of workers by ability is a way to control effort levels.

⁷⁵ This quantity results multiplying through the variations in tables 3 - 13, i.e., $6 \times 6 \times 9 \times (6 + 6) \times 5 \times (5 + 5) \times 8 \times 6 \times 7 = 195,955,200$.

associated with the exact aggregation of such a model. Perhaps every functional form is a misspecification, and that the only 'true' model is the agent-based representation that generated the data. If so, we need to learn how to identify and estimate such models, something about which little is known at present.⁷⁶

5 Conclusions

A microeconomic model of firm formation has been described analytically and explored computationally. Stable equilibrium configurations of firms do not exist in this model. Rather, agents constantly adapt to the social circumstances in which they find themselves and periodically jump from one firm to another, or start-up a new firm. This very simple model, consisting only of locally optimizing agents embedded in a world of increasing returns, proves sufficient to generate macro-statistics on firm size and growth rate distributions that closely resemble real-world data. Despite increasing returns at the firm level, approximately constant returns obtain at the aggregate level, a result due in part to agent migration between firms. It is also worth emphasizing that this model does not yield firms who are explicit profit maximizers. While there is certainly no shortage of optimizing behavior present in the model, this occurs solely at the agent level, not the firm level. Yet selection happens at the firm level. Successful are those firms capable of providing utility to income-loving agents. Welfare analysis reveals that firms serve as vehicles through which agents realize greater utility than they might otherwise.

The model yields a reasonably complete picture of the evolution of simple firms having little internal organization. There exists a well-defined firm life cycle in the model, one which shares certain qualitative similarities with actual firm life cycles. Firms are typically founded by agents who prefer income to leisure then, over time, agents with relatively greater preference for leisure join. The number of free riders grows over time in a typical firm, leading to dilution of agent income shares, exit of the most productive agents in the firm, and collapse of total firm output. Firm lifetimes are exponentially distributed.

The model also produces a noiseless picture of agent careers, both longitudinally and in cross-section. Agents move from firm to firm, staying in one place some 3-4 years on average. Overall, agent effort levels, income and utility are positively correlated with preference for income, although there is significant variation among agents having identical preferences. Income-loving agents typically work in larger firms.

⁷⁶ I thank Carl Christ for emphasizing to me the importance of such matters for the continued development of this modeling approach.

The general character of these results was demonstrated to be robust to a great number of variations in the model specification. Altering the increasing returns parameters merely changes the slope of the power law firm size distribution, as do modifications to the extent and composition of agent social networks. Permitting the agents to display 'loyalty' to their firms tends to produce larger firms, as does limiting the agents' abilities to adjust their A variety of alternative compensation schemes were effort levels. investigated, as was the possibility that firm founders act as residual claimants. While the right-skewed character of the size distribution proved to be robust to a number of alternative specifications of agent preferences, it is also true that if agents were homogeneous and did not significantly value income then few large firms resulted. Similarly, when agents did not act purposively, but either chose their efforts randomly or chose to migrate between firms selected at random, then the model failed to generate size distributions that resembled empirical data. Finally, when firms are able to observe the preferences of job applicants then firms tend to be more stable and large firms can survive. However, if founders perfectly filter the population for agents having preferences for income equal to or greater than their own then the power law character of the size distribution breaks down.

Future work includes modeling the output market. There are two ways to interpret how output is converted into income in the present model. First, given that stationary distributions of income and output arise in the model, any market mechanism would yield an essentially constant price for the output good, treated here as a single homogeneous consumption good. Alternatively, imagine that agents have heterogeneous preferences over many goods, and that firms specialize in producing a single one of these goods. Then, as long as the economy is sufficiently large—as long as there are a large number of firms-all products will be made in more or less constant amounts and, given stable agent preferences, there will obtain an equilibrium vector of prices for these goods. While there are likely many interesting facets to actually building an output market, it is also the case that a powerful result of the present model is that through purely local interactions there develops long-range correlations—the power law size distributions have no characteristic scale, and therefore all sizes are present up through the finite size cut-offs. That is, by coupling the agents through the output market we would expect long-range correlations, but such internal structure arises spontaneously in the model even when the agents are not explicitly coupled globally.

5.1 The Emergence of Firms and "Universality"

The essential result of this paper is to connect an explicit microeconomic model of endogenous firm formation to aggregate firm data. It describes a bridge between the micro and macro worlds, accomplished using agent-based computational methods.⁷⁷ But clearly the model is *very* minimal, so spare as to seem quite unrealistic.⁷⁸ How is it that such a stripped-down model of behavior, yielding such 'flat' organizations, could ever resemble empirical data?

One plausible answer to this apparent riddle is what physicists call 'universality'. In certain physical systems the detailed behavior of the constituents of the system do not make much difference in the macroscopic behavior of the system. Perhaps the canonical example of this occurs in phase transitions, where a variety of models of the actual dynamics have more or less equivalent statistical mechanics. More to the point for social systems, it has been demonstrated by Nagel and Rasmussen [1994] that universality exists in a class of agent-based traffic models. It turns out that empirically-relevant distributions of traffic jamming behavior result from a variety of microspecifications of agent driving behavior. One way to think about universality is to view it as a property of highly constrained systems, with few degrees of freedom for the components of the system—each agent in a traffic jam has so few options that it doesn't really matter very much which driving strategy is adopted, all will yield about the same travel time.

We suspect that a very similar phenomenon is at work in the present model. In part B of the appendix it is demonstrated that fairly weak requirements on effort level adjustment functions are necessary in order to produce instability at the agent level. That is, the strategic environment is so highly structured and the feasible behavioral options so constrained that very wide classes of purposive agent action yield the kinds of results reported in this paper.⁷⁹

This would seem to mitigate against criticisms that the agents described herein are too myopic, act too little strategically, and have too little rationality

⁷⁷ It is folk wisdom—ostensibly attributable to no one in particular, although I first heard it from Chris Langton—that agent-based models can function as a kind of 'macroscope,' permitting one to see the macroscopic regularities that result from particular microscopic specifications.

⁷⁸ In this the model is a kin to Schelling's segregation model [Schelling 1978]; for a more recent discussion of the Schelling model, including a modern implementation as well as citations to related work, see Epstein and Axtell [1996: 165-171].

⁷⁹ Universality in this guise seems to suggest that in environments such as the one described here there is little value to be gained by discovering, as through experiments, just how people behave. In essence, as long as behavior has the 'right sign at the margin' then certain patterns of aggregate behavior will always emerge.

to arrive at anything even approximating optimal strategies. Universality aside, it is possible to render at least two additional defenses of the relatively simple agents employed herein. First, the economic environment in which the agents find themselves is *combinatorially too complex* for even highly capable agents to be able to deduce anything like rational behaviors. There are just too many possible coalition structures for anything like an optimal one to ever get considered, to say nothing of sampled. Each agent finds itself in perpetually novel circumstances.⁸⁰ Second, the strategic environment is *too complex dynamically* for agents to make anything like accurate forecasts or predictions, even in the relatively short run.⁸¹ Agents are constantly migrating between firms. Firms are constantly being formed, growing, and dissolving. While there is local quasi-stationarity in particular firms at specific times, overall there is constant flux and adjustment.

These considerations bring into relief a feature of conventional game theory that, although tacitly acknowledged, represents an axiom that is never justified. Here we speak of the equation of social equilibrium with agentlevel equilibrium. That is, given that the goal of game theoretic models is to explain certain social and economic regularities, it is implicitly assumed that 'equilibria' must be the result of microeconomic aggregate such equilibria-game theory treats the micro- and macro-worlds as being homogeneous in this sense. But macroscopic regularities that have the character of statistical equilibria-stationary distributions, for instance-may have two conceptually quite distinct origins. When equilibrium at the agent level is achieved, perhaps as stochastic fluctuations about one or more deterministic equilibria, then there is a definite sense in which macrostationarity is a direct consequence of micro-equilibrium (for examples of this see Young [1993]). But when there does not exist a stable equilibrium, either deterministic or stochastic, then the homogeneity assumption of equilibrium game theory is invalid, yet it may often be the case that stationary configurations—patterns—will appear at the macro-level nonetheless. Furthermore, when equilibria exist and are stable but require, from generic initial conditions, an amount of time to be realized that is long in comparison to the duration of the economic process under consideration, then the homogeneity assumption is again violated, and one may be better off looking for aggregate regularities in the long-lived transients. This critique is particularly relevant to coalition formation games in large agent populations, where the number of coalition structures is given by the unimaginably vast Bell numbers, and it would thus be extremely unlikely that anything like stable equilibrium coalitions could ever be realized. Therefore constant flux

⁸⁰ Anderlini and Felli [1994] ascribe the inherent impossibility of writing complete contracts to the combinatorial complexity of nature.

⁸¹ Anderlini [1998] describes the kinds of forecasting errors that are intrinsic in such environments.

in the composition of firms would seem to be undeniable, leading to the conclusion that microeconomic equilibrium, while it may exist, is never attained.

5.2 What Constitutes a 'Theory of the Firm'?

Extant theories of the firm are steeped in this kind of micro-to-macro homogeneity. They begin innocuously enough, with purposive agents in strategic environments of one kind or another, notionally similar to some known organization form (e.g., hierarchy). They then go on to derive the performance of the resulting firms in response to strategic rivals, uncertainty, information processing constraints, and so on. But these derivations are almost everywhere characterized by equilibrium theorizing, that is, inter-firm stationarity is seen as the result of intra-firm equilibrium and thus the homogeneity assumption is manifest.⁸²

It is a claim of this paper that preoccupation with equilibrium notions is largely responsible for the relative neglect of the overall size distribution of firms in industrial organization.⁸³ While the existence and stability of this distribution have been well-known for decades, and while microeconomic conceptions now dominate industrial organization theory, there apparently does not exist a microeconomic (equilibrium) explanation of the aggregate size distribution. Indeed, perhaps it is the case that no equilibrium theory could ever reproduce the empirical data.⁸⁴ For as alluded to in § 3.4 above, the recent body of work on 'self-organized criticality' suggests that power law distributions are generically *not* the result of perturbations about an equilibrium configuration.

Some readers will find implausible the suggestion that the model described herein should stand on equal footing with any of the conventional theories of the firm. Indeed, there is no firm-specific management, organization, or 'strategy' in the present model. Cost functions, product pricing, profit maximization and other staples of the neoclassical texts are nowhere present. Also absent are explicit transaction costs, physical assets, specialization—even the very notion of ownership is ambiguous. Nor do product or process innovations, or technological change of any kind, for that matter, enter into the picture—indeed, differentiated products are not even

⁸² Least guilty of this charge is the evolutionary paradigm.

⁸³ For example, in the recent textbook of Shy [1995], the overall size distribution is not mentioned at all, and discussions of concentration are both industry-specific and purely descriptive—measures of concentration are defined but no theory of concentration is proffered.

⁸⁴ For a related critique of equilibrium theorizing in the neoclassical theory of the firm, see Lazonick [1991].

part of the model. Yet there are two senses in which the model developed herein is clearly a 'theory of the firm.'

First, supplementing Stanley *et al*. [1996: 806], there are three important empirical facts that any *accurate* theory of the firm must reproduce:

- (a) firm sizes must be right-skewed, approximating a power law;
- (b) firm growth rates must be Laplace distributed;
- (c) the standard deviation in log growth rates as a function of size must follow a power law with exponent -0.15 ± 0.03 .

Condition (a) may or may not be redundant given that (b) is satisfied, depending on the model in question. One additional requirement suggests itself in order to bring the theory of the firm into the company of the rest of modern economics—it should be methodologically individualist:

(d) the model must be written at the level of individual agents.

Aside from the research described in this essay, theories of the firm that satisfy all these requirements are unknown to us.

While today there exists a variety of theories of the firm, none is sufficiently explicit to be operationalized at a level comparable to the model described herein. That is, although each is stated at the microeconomic level-in terms of individual firms or agents-the focus of each on equilibrium means that agent behavior away from equilibrium is left more or In the language of Simon [1976], these theories are less unspecified. substantively rational, not procedurally so. Micromechanisms by which agents might plausibly arrive at the equilibria described in these theories are not usually made explicit.⁸⁵ The second sense in which our model can be considered a 'theory of the firm' is that agent-based computational models always constitute an *explanation* of the phenomenon they reproduce.⁸⁶ In the philosophy of science an explanation is only defined with respect to a theory.87 A theory has to be general enough to permit many provide explanations of whole instantiations—to classes of phenomena—while not being so vague that it can rationalize all phenomena—i.e., is not falsifiable. Each parameterization of an agent-based model is an instantiation of a more general agent 'theory'. Executing a

⁸⁵ I began my investigations of endogenous firm models with the expectation of drawing heavily on existing economic theory in order to extend agent-based computational modeling to include the self-organization of firms and related types of organizations. While I did not literally expect to be able to turn Coase's beautiful prose into operational software line-by-line, I did expect to find significant guidance in the post-1937 literature as to the micromechanisms by which agents form firms. But these expectations were soon dashed by this focus of conventional theories of the firm on equilibrium.

⁸⁶ According to Simon (Ijiri and Simon [1977: 118]): "To 'explain' an empirical regularity is to discover a set of simple mechanisms that would produce the former in any system governed by the latter."

⁸⁷ This is the so-called deductive-nomological (D-N) view of explanation; see Hempel [1966].

particular instantiation yields patterns and regularities that can be compared to empirical data, thus making both the instantiation as well as the overall model falsifiable.⁸⁸

So, what is the *explanation* for firms and the firm size distribution being proposed here? That is, *why* do purposively-behaving agents form firms having a power law size distribution and a Laplace growth rate distribution? The answer to this question is not profitably given by appeal to reduced form equations estimated econometrically. Rather, the ultimate explanation for the regularities that arise in the model lies with the agents themselves: if agents behave in the way we have specified then the various distributions described above result. Agent-based computational models are in this way a kin to laboratory science, where the procedures of investigation are described in as much detail as the discovery itself. Ultimately, we may have a mathematical theory that links the micro-specifications to the macroregularities, but for now one must be content with the discovery that the latter result from the former.

It is sometimes said that the science of thermodynamics owes more to the steam engine than the steam engine owes thermodynamics. Perhaps the same is true of firms and economic theory. Really-existing firms and organizations have provided the important concept of bounded rationality, a notion still far from fully-integrated into economic theory.⁸⁹ More recently, consideration of the economics of information within organizations has led to promotion of both 'information processing' perspectives and 'local interactions' approaches to economic models (cf. Radner [1993], Van Zandt [1996, forthcoming]). It is hoped that the present paper will provide impetus for general development of non-equilibrium theories in economics, in which case computational methods may prove to be relatively more efficacious than purely analytical ones, at least if progress in other branches of science is any guide. But it is probably the case that a move to non-equilibrium models is a larger leap than either the modest hop to local interactions or the somewhat greater jump to bounded rationality, and requires significant re-tooling of intellectual capital. However, once such a transition is underway it will certainly be justified to say that economic theory owes more to firms than the other way around.

⁸⁸ In models that are intrinsically stochastic multiple executions of individual model instances must be made in order to filter out idiosyncratic correlation from regularities that are robust to stochastic variation.

⁸⁹ Simon's early conclusion concerning the unreasonableness of fully-optimizing models was a direct result of his empirical work on organizations [1947].

5.3 The Comparative Advantages of Computational Agents

This model is a first step toward a more realistic theory of the firm, one with explicit micro-foundations, employing heterogeneous agents with bounded cognitive abilities, who interact locally in perpetually novel environments. It has also been demonstrated that this approach is one which produces empirically accurate results.

The technical apparatus utilized to produce these new results is the computer. Historically, computers have been employed by economists to solve equations, from mathematical programming problems to so-called microsimulation techniques to macroeconomic models written entirely in terms of aggregate variables. More recently, numerical methods have worked their way into all sub-disciplines of the field, especially in areas where models are complex or analytically intractable.⁹⁰ Numerical techniques largely complement conventional theorizing, serving as the crank by which computer hardware churns initial conditions into intermediate results and then final answers.⁹¹

The way in which computer power is being harnessed here is very different from such numerical workouts. Implicit in the agent-based modeling techniques on display above is a rejection of the idea that a set of equations of relatively small dimension, whether deterministic or stochastic, can ever be a satisfactory representation of real economic processes. So many heroic assumptions are made in writing down such equations—concerning aggregation, continuity, rationality, information, price-taking—and then again in their solution—homogeneous agents, equilibrium, common knowledge, continuum of agents—as to all but vitiate the results obtained, except to those used to laboring under such heavy burdens and who have acquired the calluses requisite for quieting what would otherwise surely be significant discomfort.

Agent-based computation provides a way out. Agents can be arbitrarily heterogeneous. They can interact with one another directly (local interactions) or indirectly through aggregate economic variables (global interactions). Agents can possess but a limited amount of information and are of necessity boundedly rational, since to model full rationality is too computationally complex in all but extremely simple environments (Papadimitriou [1993]). Aggregation happens automatically in such models, with noiseless cross-sectional information available at each instant. Aggregate relationships *emerge* in such models and are thus not limited *a priori* by what the 'armchair economist' (Simon [1986]) can imagine. Perhaps

⁹⁰ The recent volume of Judd [1998] is testament of the extent to which numerical analysis has penetrated economic theory and practice.

⁹¹ Judd [1997] has argued that numerical methods can also function, in some circumstances, as a substitute for conventional theorizing.

most importantly, there is no need to postulate the existence or attainment of equilibrium. Rather, one merely interrogates the model output for patterns and regularities, which may or may not include stable equilibria. Indeed, as the present model has demonstrated, agent-based computational modeling is a very natural technique for studying economic processes that are far from equilibrium.

The present work has just scratched the surface of the pregnant interface between autonomous agent modeling and the theory of the firm. Much work remains to be done.

Appendix

In this appendix the particular function forms used in § 2 are relaxed. The main reason for adopting explicit functional forms above was to provide a clear benchmark against which the computational model could be compared. Of course, the computational model requires that particular functions be used. In § A, the Cobb-Douglas preferences are replaced by a general form involving only single-peakedness. Then, the effort level adjustment function employed above is generalized in § B.

A More General Preferences

Here a general model of agent preferences is described and results analogous to those of section 2.2 derived. These concern the existence of equilibrium effort levels in a group of fixed size.

Each agent has preferences for income, *I*, and leisure, *L*, with more of either preferred to less, *ceteris paribus*. Agent *i*'s income is monotone non-decreasing in its effort level e_i as well as that of the other agents in the group, $E_{\sim i}$. Its leisure is a non-decreasing function of $1 - e_i$. Overall, the agent's utility can be written as

$$U^{i}(e_{i};E_{i}) = U^{i}(I(e_{i};E_{i}),L(1-e_{i}))$$

with

$$\frac{\partial U^{i}(e_{i}; E_{\sim i})}{\partial I} > 0, \frac{\partial U^{i}(e_{i}; E_{\sim i})}{\partial L} > 0$$

and

$$\frac{\partial I(e_i; E_{\sim i})}{\partial e_i} > 0, \frac{\partial I(e; E_{\sim i})}{\partial E_{\sim i}} > 0, \frac{\partial L(e_i)}{\partial e_i} < 0,$$

Furthermore, assuming that each of these partial derivatives is continuous in its arguments, and that $U^i(0, \cdot) = U^i(\cdot, 0) = 0$, it is easy to show that each utility function is single-peaked in effort level.

Each agent selects the unique effort level that maximizes its utility, as in equation (4). The first-order conditions amount to

$$\frac{\partial U^{i}(e_{i};E_{\sim i})}{\partial I}\frac{\partial I(e_{i};E_{\sim i})}{\partial e_{i}} + \frac{\partial U^{i}(e_{i};E_{\sim i})}{\partial L}\frac{\partial L(e_{i};E_{\sim i})}{\partial e_{i}} = 0$$

which can be written as

$$\frac{\partial U^{i}(e_{i};E_{\sim i})}{\partial I}\frac{\partial I(e_{i};E_{\sim i})}{\partial e_{i}}=-\frac{\partial U^{i}(e_{i};E_{\sim i})}{\partial L}\frac{\partial L(e_{i};E_{\sim i})}{\partial e_{i}}.$$

Axtell

From the inverse function theorem there exists a solution to this equation having the form $e_i^* = \max[0, \zeta(E_{\sim i})]$, analogous to (5). Furthermore, from the implicit function theorem it is straightforward to show that ζ , and therefore e_i^* , is a continuous, non-increasing function of $E_{\sim i}$ (analog of figure 1).

Group effort level equilibrium corresponds to each agent determining its optimal effort level, e_i^* , assuming that the other agents are doing so as well, i.e., substituting

$$E_{\sim i}^* = \sum_{i \neq j} e_j^*$$

in place of $E_{\sim i}$. Since each e_i^* is a continuous function of $E_{\sim i}$ so is the vector of optimal effort levels, $e^* \in [0, 1]^N$, a compact, convex set. Therefore, by the Leray-Schaduer-Tychonoff theorem a solution exists. Furthermore, such a solution constitutes a Nash equilibrium. In general, a Nash equilibrium is Pareto-dominated by other effort level vectors, ones involving larger amounts of effort on the part of all agents. Here the reasoning is identical to that in § 2.2 above.

B A More General Model of Effort Level Adjustment

In section 2.3 conditions were derived relating to the onset of instability in groups, using a particular agent effort adjustment rule. Here it will be shown that an upper bound on agent group size exists for any effort level adjustment rule that possesses certain properties.

In particular, imagine that each agent adjusts its effort level according to

$$e_i(t+1) = h^i(E_{\sim i}(t))$$

where

$$\frac{dh^{i}(E_{\sim i})}{dE_{\sim i}} = \frac{\partial h^{i}(E_{\sim i})}{\partial e_{i}} \le 0.$$
(B.1)

Under these circumstances the Jacobian matrix retains the structure described in § 2.3, i.e.,

$$J = \begin{bmatrix} 0 & k_1 & \cdots & k_1 \\ k_2 & 0 & \cdots & k_2 \\ \vdots & & \ddots & \vdots \\ k_N & \cdots & k_N & 0 \end{bmatrix},$$

where each row contains N-1 identical entries and a 0 on the diagonal. The bounds on the dominant eigenvalue established in section 2.3 guarantee that there exists an upper bound on the stable group size, as long as the k_i do not vanish. Such a condition holds when (B.1) is a strict inequality, thus establishing a sufficient condition for the onset of instability in groups above some maximum size.

References

- Abernathy, W.J. and J.M. Utterback. 1978. "Patterns of Industrial Innovation." Technology Review, 80 (1): 2-29.
- Achen, C.H. and W.P. Shively. 1995. Cross-Level Inference. University of Chicago Press: Chicago, Illinois.
- Acs, Z. and D. Audretsch. 1990. Innovation and Small Firms. MIT Press: Cambridge, Massachusetts.
- Alchian, A.A. and H. Demsetz. 1972. "Production, Information Costs, and Economic Organization." American Economic Review, 62: 777-795.
- Allen, T. 1997. "Decentralized Decision-Making in Labor-Managed Firms." Dissertation Proposal. Baltimore, Maryland: Johns Hopkins University.
- Amaral, L.A.N., S.V. Buldyrev, S. Havlin, H. Leschhorn, P. Maass, M.A. Salinger, H.E. Stanley, and M.H.R. Stanley. 1997. "Scaling Behavior in Economics: I. Empirical Results for Company Growth." *Journal de Physique I France*, 7: 621-633.
- Amaral, L.A.N., S.V. Buldyrev, S. Havlin, M.A. Salinger, and H.E. Stanley. 1998. Physical Review Letters, 80: 1385.
- Anderlini, L. 1998. "Forecasting Errors and Bounded Rationality: An Example." Mathematical Social Sciences, 36: 71-90.
- Anderlini, L. and L. Felli. 1994. "Incomplete Written Contracts: Indescribable States of Nature." Quarterly Journal of Economics, 109: 1085-1124.
- Angeline, P., R.G. Reynolds, J.R. McDonnell, and R. Eberhart, editors. 1997. Evolutionary Programming VI. Springer-Verlag: New York.
- Arthur, W.B. 1994. Increasing Returns and Economic Theory. University of Michigan Press: Ann Arbor, Michigan.
- Arthur, W.B., S. Durlauf, D. Lane, eds. 1997. The Economy as an Evolving Complex System, Volume II: Addison-Wesley: New York.
- Arthur, W.B., J.H. Holland, B. LeBaron, R. Palmer and P. Tayler. 1996. "Asset Pricing Under Endogenous Expectations in an Artificial Stock Market." Working paper #96-12-093. Santa Fe Institute: Santa Fe, New Mexico.
- Asplund, M. 1998. "On the Size Distributions of Firms and Markets." Working paper. Stockholm School of Economics, Stockholm, Sweden.
- Audretsch, D.B. and T. Mahmood. 1995. "New Firm Survival: New Results Using a Hazard Function." Review of Economics and Statistics, 77 (1): 97-103.
- Axelrod, R. 1984. The Evolution of Cooperation. New York: Basic Books.
- Axelrod, R. 1995. "A Model of the Emergence of New Political Actors." In Hillebrand and Stender, eds.
- Axtell, R.L. 1997. "Three Distinct Uses of Agent-Based Computational Models in Economic Theory or When is Computational Modeling with Agents *Not* Simulation" Working Paper. Brookings Institution: Washington, D.C.

- Axtell, R.L., R. Axelrod, J.M. Epstein and M.D. Cohen. 1996. "Aligning Simulation Models: A Case Study and Results." Computational and Mathematical Organization Theory, 1 (2): 123-41.
- Axtell, R.L. and J.M. Epstein. 1994. "Agent-Based Models: Understanding Our Creations." Bulletin of the Santa Fe Institute.
- Bak, P. 1996. How Nature Works. Springer-Verlag: New York.
- Basu,, N. and R.J. Pryor. 1997. "Growing a Market Economy." Sandia Report SAND97-2093 UC-905. Sandia National Laboratories: Albuquerque, New Mexico.
- Beesley, M.E. and R.T. Hamilton. 1984. "Small Firms' Seedbed Role and the Concept of Turbulence." *Journal of Industrial Economics*, 33 (2): 217-231.
- Binmore, K.G. 1987. "Modeling Rational Players, Part I." *Economics and Philosophy*, 3: 179-214.

——. 1988. "Modeling Rational Players, Part II." Economics and Philosophy, 4: 9-55.

- Blair, M.M., D.L. Kruse, and J.R. Blasi. 1998. Is Employee Ownership and Unstable Form? Or a Stabilizing Force? Working paper. Brookings Institution: Washington, D.C.
- Blair, M.M. and L.A. Stout. 1997. "A Theory of Corporation Law as a Response to Contracting Problems in Team Production." Working paper. Brookings Institution: Washington, D.C.
- Blume, L.E. and D. Easley. 1998. "Optimality and Natural Selection in Markets." Working paper #98-09-082. Santa Fe Institute: Santa Fe, New Mexico.
- Bowles, S. 1998. "Cultural Group Selection and Human Social Structure: The Effects of Segmentation, Egalitarianism, and Conformism." Working paper. University of Massachusetts: Amherst, Massachusetts.
- Bowles, S. and H. Gintis. 1996. "Efficient Redistribution: New Rules for Markets, States and Communities." *Politics and Society*, 24: 307-342.
- and . 1998. "Mutual Monitoring in Teams: The Effects of Residual Claimancy and Reciprocity." Working paper #98-08-074. Santa Fe Institute: Santa Fe, New Mexico.
- Bolton, P. and M. Dewatripont. 1994. "The Firm as a Communication Network." *Quarterly Journal of Economics*, 109: 809-839.
- Brooks, R., ed. 1994. Artificial Life IV. MIT Press: Cambridge, Massachusetts.
- Buchanan, J.M. and Y.J. Yoon. 1994. *The Return to Increasing Returns*. University of Michigan Press: Ann Arbor, Michigan.
- Buldyrev, S.V., L.A.N. Amaral, S. Havin, et al. 1997. "Scaling Behavior in Economics: II. Modeling Company Growth. *Journal de Physique I France*, 7: 635-650.
- Canning, D. 1995. "Evolution of Group Cooperation through Inter-Group Conflict." Working Paper. Department of Economics, Queens University of Belfast, Northern Ireland.
- Canning, D., L.A.N. Amaral, Y. Lee, M. Meyer and H.E. Stanley. 1998. "Scaling the Volatility of GDP Growth Rates." *Economics Letters*, 60: 335-341.
- Carroll, G.R. 1983. "A Stochastic Model of Organizational Mortality: Review and Reanalysis." *Social Science Research*, 12 (4): 303-329.
- Cartwright, N. 1983. How the Laws of Physics Lie. Oxford University Press.
- Castelfranchi, C. and J.-P. Müller, editors. 1993. *From Reaction to Cognition*. Proceedings of the Fifth European Workshop on Modeling Autonomous Agents in a Multi-Agent World (MAAMAW). Springer-Verlag: New York.

- Caves, R.E. 1998. "Industrial Organization and New Findings on the Turnover and Mobility of Firms." *Journal of Economic literature, XXXVI*: 1947-1982.
- Cederman, L-E. 1997. *Emergent Actors in World Politics*. Princeton University Press: Princeton, New Jersey.
- Chen, S-H. and C-H. Yeh. 1997. "Modeling Speculators with Genetic Programming." In Angeline *et al.*, eds.
- Coase, R. 1937. "The Nature of the Firm." Economica, n.s., 4: 386-405.
- Colander, D.C. and H. Landreth. 1999. "Increasing Returns: Who, If Anyone, Deserves Credit for Reintroducing It into Economics?" Paper presented at the American Economic Association meetings, New York.
- Curry, B. and K.D. George. 1983. "Industrial Concentration: A Survey." Journal of Industrial *Economics*, 31: 203-255.
- DeCanio, S.J. and W.E. Watkins. 1998. "Information Processing and Organizational Structure." Journal of Economic Behavior and Organization, 36: 275-294.
- DeVany, A. 1993a. "Edgeworth and the Traveling Salesman: The Complexity of Coalition Structures." Working paper. Department of Economics, University of California, Irvine.
 - ——. 1993b. "Hard Cores and Soft Cores: Evolving Coalitions, Edgeworth and the Boltzmann Machine." Working paper. Department of Economics, University of California, Irvine.
 - ——. 1993c. "Information, Bounded Rationality, and the Complexity of Economic Organization." Working paper. Department of Economics, University of California, Irvine.
- ———. 1996. "The Emergence and Evolution of Self-Organized Coalitions." In Gilli, ed.
- Dong, Xioa-yuan and G. Dow. 1993a. "Does Free Exit Reduce Shirking in Production Teams?" *Journal of Comparative Economics*, 17: 472-484.
- —— and ——. 1993b. "Mutual Monitoring in Production Teams." *Journal of Political Economy*.
- Dow, G. 1990. "The Organization as an Adaptive Network." *Journal of Economic Behavior and Organization*, 14: 159-185.
- Dresher, M., L.S. Shapley, and A.W. Tucker. 1964. *Advances in Game Theory*. Annals of Mathematical Studies No. 52. Princeton University Press: Princeton, New Jersey.
- Durlauf, S. 1993. "Nonergodic Economic Growth." Review of Economic Studies, 60: 349-366.
- Ericson, R. and A. Pakes. 1995. "Markov-Perfect Industry Dynamics: A Framework for Empirical Work." *Review of Economic Studies*, 62 (1): 53-82.
- Encinosa, W.E., M. Gaynor, and J.B. Rebitzer. 1997. "The Sociology of Groups and the Economics of Incentives: Theory and Evidence on Compensation Systems." Working paper.
- Epstein, J.M. and R. Axtell. 1996. *Growing Artificial Societies: Social Science from the Bottom Up.* Brookings Institution Press/MIT Press: Washington, D.C./Cambridge, Massachusetts.
- Farrell, J. and S. Scotchmer. 1988. "Partnerships." Quarterly Journal of Economics, 103 (2): 279-297.
- Friedman, M. 1953. Essays in Positive Economics. University of Chicago Press: Chicago, Illinois.

- Gilboa, I. and A. Matsui. 1991. "Social Stability and Equilibrium." *Econometrica*, 59 (3): 859-867.
- Gilli, M. 1996. Computational Economic Systems: Models, Methods and Econometrics. Kluwer Academic Publishers: Boston, Massachusetts.
- Glance, N.S., T. Hogg, and B.A. Huberman. 1997. "Training and Turnover in the Evolution of Organizations." *Organization Science*, 8: 84-96.
- Harpending, H. and A. Rogers. 1987. "On Wrights's Mechanism for Intergroup Selection." Journal of Theoretical Biology, 127: 51-61.
- Hart, O. 1995. Firms, Contracts and Financial Structure. Oxford University Press: New York.
- Hart, S. and M. Kurz. 1983. "Endogenous Formation of Coalitions." *Econometrica*, 51: 1047-1064.
- Hausman, D. 1992. The Inexact and Separate Science of Economics. Cambridge University Press: New York.
- Heikkinen, T. 1997. "Economic Dynamics on Torus." Working paper. Department of Mathematics, University of Helsinki, Helsinki Finland.
- Hempel, C.G. 1966. *Philosophy of Natural Science*. Prentice-Hall: Englewood Cliffs, New Jersey.
- Hillebrand, E. and J. Stender, eds. 1994. *Many-Agent Simulation and Artificial Life*. IOS Press: Washington, D.C.
- Holler, M.J., ed. 1984. *Coalitions and Collective Action*. Physica-Verlag: Wuerzburg, Vienna, Austria.
- Holmström, B. 1982. "Moral Hazard in Teams." Bell Journal of Economics, 13: 324-340.
- Holmström, B. and J. Tirole. 1988. "The Theory of the Firm." In Schmalansee and Willig, eds.
- Huberman, B.A. and N.S. Glance. 1993. "Evolutionary Games and Computer Simulations." Proceedings of the National Academy of Sciences, USA, 90: 7716-18.
- and —. 1998. "Fluctuating Efforts and Sustainable Cooperation." In Prietula *et al.*, eds.
- Ijiri, Y. and H. Simon. 1977. Skew Distributions and the Sizes of Business Firms. New York: North-Holland.
- Ichiishi, T. 1977. "Coalition Structure in a Labor-Managed Market Economy." *Econometrica*, 45: 341-360.
- ——. 1993. The Cooperative Nature of the Firm. Academic Press: New York.
- Jennings, N.R., K. Sycara, and M. Wooldridge. 1998. "A Roadmap of Agent Research and Development." Autonomous Agents and Multi-Agent Systems, volume 1, number 1: 7-38.
- Johnson, N.L., S. Kotz and N. Balakrishnan. 1995. *Continuous Univariate Distributions*. Volume 2. Second edition. Wiley: New York.
- Joskow, P.L. 1993. "Asset Specificity and the Structure of Vertical Relationships: Empirical Evidence." In Williamson and Winter, eds.
- Jovanovic, B. 1982. "Selection and the Evolution of Industry." Econometrica, 50 (3): 649-70.
- Judd, K.L. 1997. "Computational Economics and Economic Theory: Substitutes or Complements." *Journal of Economic Dynamics and Control*, 21: 907-942.

97

- ——. 1998. Numerical Methods in Economics. MIT Press: Cambridge, Massachusetts.
- Kaldor, N. 1934. "The Equilibrium of the Firm." Economic Journal, 44: 60-76.
- ——. 1972. "The Irrelevance of Equilibrium Economics." *Economic Journal*, 82 (328): 1237-55.
- Kihlstrom, R. and J.-J. Lafont. 1979. "A General Equilibrium Entrepreneurial Theory of Firm Formation Based in Risk Aversion." *Journal of Political Economy*, 87: 719-748.

Kirman, A.P. 1995. "Economies with Interacting Agents." Games and Economic Behavior.

- Kirman, A.P. and N.J. Vriend. 1998. "Evolving Market Structure: A Model of Price Dispersion and Loyalty." Working paper. University of London, London, U.K.
- Kleindorfer, P. and M.R. Sertel. 1979. "Profit-Maximizing Design of Enterprises Through Incentives." *Journal of Economic Theory*, 20 (3).
- and . 1982. "Value-Added Sharing Enterprises." In Sertel, ed.
- Klepper, S. 1996. "Entry, Exit, Growth, and Innovation over the Product Life Cycle." *American Economic Review*,, *86* (3): 562-583.
- Klusch, M. and O. Shehory. 1996a. "Coalition Formation Among Rational Information Agents." In Proceedings of the Seventh European Workshop on Modeling Autonomous Agents in a Multi-Agent World (MAAMAW): 204-217.
- and . 1996b. "A Polynomial Kernel-Oriented Coalition Algorithm for Rational Information Agents." In Tokoro, ed.: 157-164.
- Krugman, P. 1996. The Self-Organizing Economy. Blackwell: Cambridge, Massachusetts.
- Kwasnicki, W. 1998. "Skewed Distribution of Firm Sizes—An Evolutionary Perspective." *Structural Change and Economic Dynamics*, volume 9: 135-158.
- Kwoka, Jr., J.E. 1982. "Regularity and Diversity in Firm Size Distributions in U.S. Industries." *Journal of Economics and Business*, 34: 91-395.
- Laherrère, J. and D. Sornette. 1998. "Stretched Exponential Distributions in Nature and Economy: "Fat Tails" with Characteristic Scales." *The European Physical Journal B*, 2: 525-539.
- Lane, D. 1993. "Artificial Worlds and Economics." Parts 1 and 2. *Journal of Evolutionary Economics*, 3: 89-107, 177-197.
- Lane, S.J. and M. Schary. 1991. "Understanding the Business Failure Rate." *Contemporary Policy Issues*, IX: 93-105.
- Latsis, S., ed. 1976. *Method and Appraisal in Economics*. Cambridge University Press: New York.
- Laussel, D. and M. LeBreton. 1995. "A General Equilibrium Theory of Firm Formation Based on Unobservable Skills." *European Economic Review*, 39: 1303-1319.
- Lazonick, W. 1991. Business Organization and the Myth of the Market Economy. Cambridge University Press: Cambridge, U.K.
- Lee, Y., L.A.N. Amaral, D. Canning, M. Meyer and H.E. Stanley. 1998. "Universal Features in the Growth Dynamics of Complex Organizations." *Physical Review Letters*, 81 (15): 3275-3278.
- Legros, P. and A.F. Newman. 1996. "Wealth Effects, Distribution, and the Theory of Organization." *Journal of Economic Theory*, 70: 312-341.
- Lewin, D.I. 1997. "Growing an a Tree of Knowledge in an Object-Oriented Eden." *Computers in Physics*, 11 (1): 6-7.

Lucas, Jr., R.E. 1978. "On the Size Distribution of Business Firms." *Bell Journal of Economics*, 9: 508-523.

Luenberger, D.G. 1979. Introduction to Dynamic Systems. Wiley: New York.

- Luna, F. 1997. "Induction and Firm Creation." Working Paper. University of Venice Ca' Foscari: Venice, Italy.
- Majumdar, M. 1996. Organizations with Incomplete Information. Cambridge University Press: New York.
- Malone, T.W. and S.A. Smith. 1988. "Modeling the Performance of Organizational Structures." *Operations Research*, 36: 421-436.

Mandelbrot, B. 1997. Fractals and Scaling in Finance. Springer-Verlag: New York.

- ———. 1998. Lecture at the University of Michigan, October 7.
- Marschak, J. and R. Radner. 1972. *Economic Theory of Teams*. Yale University Press: New Haven, Connecticut.
- Marshall, A. 1920 Principles of Economics. Eighth edition. Macmillan Press: London, U.K.
- Mazzucato, M. 1998. "A Computational Model of Economies of Scale and Market Share Instability." Structural Change and Economic Dynamics, 9: 55-83.
- Miller, J. 1998. "Evolving Information Processing Organizations." Working paper #95-06-053. Santa Fe Institute: Santa Fe, New Mexico.
- Morel, B. 1998. "The Distribution of Business Firms Size, Stochasticities, and Self-Organized Criticality." Working Paper. Carnegie-Mellon University.
- Murray, J.D. 1989. Mathematical Biology. Springer-Verlag: New York.
- Nagel, K. and S. Rasmussen. 1994. "Traffic at the Edge of Chaos." In Brooks, ed.
- Nelson, R., and S. Winter. 1982. An Evolutionary Theory of Economic Change. Harvard University Press: Cambridge, Massachusetts.
- Olson, M. 1965. *The Logic of Collective Action*. Harvard University Press: Cambridge, Massachusetts.
- Osborne, M.J. and A. Rubinstein. 1994. A Course in Game Theory. MIT Press: Cambridge, Massachusetts.
- Ostrom, E. 1990. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press: New York.
- Padgett, J.F. 1997. "The Emergence of Simple Ecologies of Skill: A Hypercycle Approach to Economic Organization." In Arthur *et al.*, eds.
- Pakes, A. and P. McGuire. 1994. "Computing Markov-Perfect Nash Equilibria: Implications of a Dynamic Differentiated Product Model." *RAND Journal of Economics*, 79 (2): 555-589.
- and . 1996. "Stochastic Algorithms for Dynamic Models: Markov Perfect Equilibrium and the 'Curse' of Dimensionality." Working paper. Yale University: New Haven, Connecticut.
- Papadimitriou, C. 1993. "Computational Complexity as Bounded Rationality." *Mathematics* of Operations Research.
- Papageorgiou, Y.Y. and T.R. Smith. 1983. Agglomeration as a Local Instability of Spatially Uniform Steady-States." *Econometrica*, 51 (4): 1109-1119.

Penrose, E.T. 1959. The Theory of the Growth of the Firm. Sharpe: White Plains, New York.

- Prietula, M., K. Carley and L. Gasser, eds. 1998. *Simulating Organizations: Computational Models of Institutions and Groups.* MIT Press: Cambridge, Massachusetts.
- Putterman, L. 1998. "The Firm as Association Versus the Firm as Commodity." *Economics and Philosophy*, *4*: 243-266.
- Quandt, R.E. 1966. "On the Size Distribution of Firms." American Economic Review: 416-432.
- Radner, R. 1993. "The Organization of Decentralized Information Processing." *Econometrica*, 61(5): 1109-1146.
- Robertson, D.H. 1930. Control of Industry. Nisbet and Company: London.
- Robinson, E.A.G. 1953. *The Structure of Competitive Industry*. Cambridge University Press: Cambridge, U.K.
- Roth, A. 1984. "Stable Coalition Formation: Aspects of a Dynamic Theory." In Holler, ed.
- Sandholm, T., K. Larsen, M. Andersson, O. Shehory and F. Tohmé. 1998. "Anytime Coalition Structure Generation with Worst Case Guarantees." *Proceedings of the American Association for Artificial Intelligence*.
- Schary, M.A. 1991. "The Probability of Exit." RAND Journal of Economics, 22 (3): 339-353.
- Schelling, T.C. 1978. Micromotives and Macrobehavior. Norton: New York.
- Schmalansee, R. 1989. "Inter-Industry Studies of Structure and Performance." In Schmalansee and Willig, eds.
- Schmalansee, R. and R.D. Willig, eds. 1988. *Handbook of Industrial Organization, Volume I.* Elsevier: New York.
- and —, eds. 1989. Handbook of Industrial Organization, Volume II. Elsevier: New York.
- Sertel, M.R., editor. 1982. Workers and Incentives. North-Holland: New York.
- Shapley, L. 1964. "Some Topics in Two-Person Games." In Dresher et al.
- Shehory, O. and S. Kraus. 1993. "Coalition Formation Among Autonomous Agents: Strategies and Complexity." In Castelfranchi and Müller, eds.
- Sherstyuk, K. 1998. "Efficiency in Partnership Structures." Journal of Economic Behavior and Organization, 36: 331-346.
- Shubik, M. 1997. "Why Equilibrium? A Note on the Noncooperative Equilibria of Some Matrix Games." *Journal of Economic Behavior and Organization*, 29: 537-539.
- Shy, O. 1995. Industrial Organization. MIT Press: Cambridge, Massachusetts.
- Simon, H.A. 1947. Administrative Behavior. Macmillan: New York.
- ———. 1976. "From Substantive to Procedural Rationality." In Latsis, ed.
- -----. 1986. "The Failure of Armchair Economics." Challenge, 29 (5): 18-25.
- ——. 1997a. *The Sciences of the Artificial*. Third edition. MIT Press: Cambridge, Massachusetts.
- ——. 1997b. An Empirically-Based Microeconomics: The Raffaele Mattioli Lectures. Cambridge University Press: Cambridge, United Kingdom.
- Sober, E. and D.S. Wilson. 1998. Unto Others: The Evolution and Psychology of Unselfish Behavior. Harvard University Press: Cambridge, Massachusetts.

- Sraffa, P. 1926. "The Laws of Returns Under Competitive Conditions." *Economic Journal*, XXXVI (144): 535-550.
- Stanley, M.H.R., L.A.N. Amaral, S.V. Buldyrev, S. Havlin, H. Leschhorn, P. Maass, M.A. Salinger, H.E. Stanley. 1996. "Scaling Behavior in the Growth of Companies." *Nature*, 379: 804-806.
- Steindl, J. 1965. Random Processes and the Growth of Firms. Hafner Publishing Company: New York.
- Sutton, J. 1997. "Gibrat's Legacy." Journal of Economic Literature, XXXV: 40-59.
- ——. 1998. "The Size Distribution of Businesses." Seminar Handout.
- Teitelbaum, D. 1998. "Technological Change and Pollution Control: An Adaptive Agent-Based Analysis." Ph.D. dissertation. Carnegie Mellon University.
- Tesfatsion, L. 1998. "Ex Ante Capacity Effects in Evolutionary Labor Markets with Adaptive Search." Economic Report Series No. 48. Iowa State University: Ames, Iowa.
- Tokoro, M., ed. 1996. Proceedings of the Second International Conference on Multi-Agent Systems (ICMAS 96). AAAI Press: Menlo Park, California.
- Uyenoyama, M.K. 1979. "Evolution of Altruism Under Group Selection in Large and Small Populations in Fluctuating Environments." *Theoretical Population Biology*, 15: 58-85.
- Uyenoyama, M.K. and M.W. Feldman. 1980. "Theories of Kin and Group Selction: A Population Genetics Approach." *Theoretical Population Biology*, 17: 380-414.
- Van Zandt, T. 1996. "Organizations with and Endogenous Number of Information Processing Agents." In Majumdar, ed.

——. Forthcoming. "Real-Time Decentralized Information Processing as a Model of Organizations with Boundedly Rational Agents." *Review of Economic Studies.*

Varian, H. 1984. Microeconomic Analysis. Second edition. Norton: New York.

- ------. 1990. "Monitoring Agents with Other Agents." *Journal of Institutional and Theoretical Economics*, 46 (1): 153-174.
- Vernon, R. 1966. "International Investment and International Trade in the Product Cycle." *Quarterly Journal of Economics*, 80 (2): 190-207.
- Williamson, O.E. 1975. Markets and Hierarchies. Macmillan: New York.
- ——. 1985. *The Economic Institutions of Capitalism*. Free Press: New York.
- Williamson, O.E. and S.G. Winter, eds. 1993. *The Nature of the Firm*. New York: Oxford University Press.
- Wilson, R. 1968. "The Theory of Syndicates." Econometrica, 36 (1): 119-132.
- Winter, S.G. 1993. "On Coase, Competence and the Corporation." In Williamson and Winter, eds.

Young, A. 1928. "Increasing Returns and Economic Progress." Economic Journal, 38: 527-542.

Young, H.P. 1993. "The Evolution of Conventions." *Econometrica*, 61 (1): 57-84.

rla/msw/179