To appear in ***Proceedings of the International Conference on Autonomous Agents and Multi-Agent Systems***
Bologna, Italy, 15-19 July 2002

Non-Cooperative Dynamics of Multi-Agent Teams

Robert L. Axtell
The Brookings Institution
1775 Massachusetts Avenue, NW
Washington, D.C. 20036 USA
202-797-6020
raxtell@brookings.edu

ABSTRACT[1]

Results on the formation of multi-agent teams are reviewed and extended. Conditions are specified under which it is individually rational for agents to spontaneously form coalitions in order to engage in collective action. In a cooperative setting the formation of such groups is to be expected. Here we show that in non-cooperative environments—presumably a more realistic context for a variety of both human and software agents—self-organized coalitions are capable of extracting welfare improvements. The Nash equilibria of these coalitional formation games are demonstrated to always exist and be unique. Certain free rider problems in such group formation dynamics lead to the possibility of dynamically unstable Nash equilibria, depending on the nature of intra-group compensation and coalition size. Yet coherent groups can still form, if only temporarily, as demonstrated by computational experiments. Such groups of agents can be either long-lived or transient. The macroscopic structure of these emergent 'bands' of agents is stationary in sufficiently large populations, despite constant adaptation at the agent level. It is argued that assumptions concerning attainment of agent-level (Nash) equilibrium, so ubiquitous in conventional economics and game theory, are difficult to justify behaviorally and highly restrictive theoretically, and are thus unlikely to serve either as fertile design objectives or robust operating principles for realistic multi-agent systems.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence – *intelligent agents, multiagent systems*

I.6.8 [**Simulation and Modeling**]: Types of Simulation – *distributed, Monte Carlo*

J.4 [**Social and Behavioral Sciences**]: Economics

K.4.4 [**Computers and Society**]: Electronic Commerce – *distributed commercial transactions*

General Terms

Performance, Design, Economics, Experimentation, Theory.

Keywords

Multi-agent coalitions, team formation, increasing returns, proportional reward, equal division, unstable Nash equilibria, transient groups, stationary distribution of group sizes

1. INTRODUCTION

Coalition formation among self-interested agents is a subject of long-standing of interest among researchers in the social sciences, particularly economics and game theory, and more recently in computer science and multi-agent systems (MAS). Here we will explore a class of models that relax conventional assumptions concerning rationality and equilibrium, focusing instead on adaptive behavior that may or may not lead to stable equilibria. From the perspective of the social sciences a primary motivation for this work stems from the observation that .real people do not behave in ways that comport closely to the tenets of rationality. This, combined with the fact that game theoretic results on coalition formation seem to have little empirical relevance, leads one to hope that more realistic specifications of individual behavior will yield salient empirical modelss. Alternatively, a motivation closer to MAS originates with the desire to assess the possibility that autonomous software agents could organize themselves—self-organize—into coalitions ('roving bands') that would then attempt to bargain collectively in e-commerce markets and exchanges. Consider the case of a price-discriminating supplier selling a homogeneous good with near zero marginal cost (e.g., airline seats) to agents having heterogeneous valuations. Perfect price discrimination results in agents with high values paying more than agents with lower values. We investigate circumstances in which it is welfare-improving for agents, whether heterogeneous or not, to engage in collective bargaining with the supplier.

Conventional game theory has both positive and normative functions in the social sciences. To the extent that it employs unrealistic assumptions about human behavior, game theory is primarily normative, describing how a perfectly rational agent should behave so as to extremize its welfare (e.g., maximize profit). Insofar as the synthesis of well-functioning multi-agent systems utilizes highly rational agents, conventional normative-style game theoretic results are potentially applicable in the design of such systems. However, the positive (descriptive) utility of such results is very ambiguous in the context of human agents, due to the unrealistic behavioral axioms on which the theory is based. This state of affairs has led in recent years to the development of a more evolutionary, adaptive-agent game theory [70], providing more plausible mechanisms by which results based on perfectly rational agents (e.g., attainment of Nash equilibrium) might be achieved by boundedly rational actors. But this work also demonstrates that many of the conventional results are deeply problematical—e.g., impossibility of learning certain equilibria, exponential complexity of designing efficient mechanisms. Increasingly, these evolutionary approaches provide novel empirical explanations of multi-agent institutions [3].

There is no little irony in the adoption of game theoretic methods by the distributed AI (DAI) and multi-agent systems

(MAS) communities. For within the social sciences there exists a group of 'early adopters' of MAS technology whose main goal is to use such systems precisely to circumvent the kinds of highly idealized and unrealistic assumptions that are conventional in game theory specifically and the formal social sciences generally (see [4, 26, 44] for additional background on the use of MAS in the social sciences).

The motivations for using game theory in MAS are clear enough, particularly as a tool for analysis of the performance of systems of self-interested agents [65]. What seems much more dubious is the adoption of the cognitively, behaviorally and socially narrow definition of 'agent' implicit in conventional game theory. This is especially ironic insofar as the structure and specification of software agents permits the kinds of adaptive, non-equilibrium behavior that are outside the norms of the theory.

In this paper we take a stand for a more evolutionary game theoretic formulation of coalition formation.[2] In particular, agents are postulated to be heterogeneous, boundedly rational, and capable of engaging in interactions outside of equilibrium. It will be argued that assuming micro-level (agent) equilibrium, as is the norm in conventional game theory [7, 8, 24], is highly restrictive theoretically and often false empirically. Similarly with rationality—while the purposiveness of human subjects is usually clear in experimental settings of both the cognitive science and behavioral economics varieties, there is also great evidence that people systematically depart from purely rational behavior. Of course, it is the restrictive nature of the equilibrium and rationality assumptions that are the primary basis for their widespread use, not their realism. Highly restrictive assumptions make it possible to deduce the properties of a model formally and to prove theorems. But the empirical relevance of such theorems are only as good as the assumptions upon which they are based.

Faced with models featuring full blown heterogeneous, boundedly rational actors interacting out of equilibrium, we invoke the technology of multi-agent systems to bypass the limits on pure analytical deduction, in effect using the computer as an engine of mechanical deduction. We answer a host of questions in this way, such as 'What kinds of aggregate phenomena exist when agents have such and such behavior?' and 'Which specifications of agent behaviors yield statistically indistinguishable behavior at the population level?'

In particular, this paper explores the dynamical formation of multi-agent teams, theoretically, computationally, and to a lesser extent empirically, while presenting some thoughts on the ultimate relevance of the results for the design and synthesis of well-functioning systems of agents in the service of practical ends. We find that heterogeneous agents who are purposive but sub-rational and who interact directly with one another away from equilibrium are capable of arranging themselves into groups or teams that have population-level statistics that closely resemble firms in actual industrial economies. It is demonstrated that each of the features of the micro-specification is necessary in order to obtain the main results—homogeneous agents produce empirically unreasonable group size distributions, while the highly dynamic nature of the economic environment precludes fully rational agents from ever being able to deduce very much about the economy in which they live. It is then argued from sufficiency that since real human agents organize themselves into groups having such properties, and since we have discovered how to make artificial agents behave in this way, that this agent specification may constitute a useful starting point for understanding multi-agent teams that may spontaneously appear in real and artificial economic environments, and perhaps provide useful design guidelines for multi-agent organizations.

In the next section the general structure of this model is analyzed game theoretically. It is demonstrated that, depending on the intra-group compensation mechanism, for groups of any composition there is always a size above which the Nash equilibrium in the group is unstable to infinitessimal perturbations. This leads to a non-equilibrium theory of agent groups, studied computationally in section 3 computationally. Section 4 summarizes the main results and draws conclusions concerning general questions of the endogenous formation multi-agent groups.

2. A DYNAMICAL THEORY OF TEAMS

Consider a group of agents who are collaborating to produce or buy a good or service. Each agent i has an endowment, e_i, and a utility function, U^i, and contributes some amount, x_i, of its endowment to the team. Total contributions by team members amount to $X \equiv \sum x_i$; $X_{\sim i} \equiv X - x_i$. The team uses these inputs to create or acquire output having value $Q(X)$. In return for their contributions, agents receive compensation, either as some fraction of the output or as income derived from its sale. Compensation is a non-decreasing function of agent contribution, x_i, meaning simply that increased contributions to production are not penalized. Agent utility is increasing in the amount of compensation and the quantity of endowment not contributed. The agents are heterogeneous with respect to their relative interest in compensation vs. maintenance of their endowment.

2.1 Compensation Proportional to Input

Specifically, if each agent has Cobb-Douglas preferences for compensation and endowment, and if compensation is proportional to individual input, then utility can be written as

$$U^i\left(x_i;\theta_i,e_i,X_{\sim i}\right) = \left(\frac{x_i}{x_i + X_{\sim i}} Q\left(x_i;X_{\sim i}\right)\right)^{\theta_i} \left(e_i - x_i\right)^{1-\theta_i} \qquad (1)$$

where θ_i we shall term, heuristically, agent i's 'preference for income.'

In order to understand how agents, whether fully rational or merely purposive, might behave in this economic environment, it is next necessary to specify the structure of production, i.e., how the various inputs are turned into output. Here we shall give production the character of increasing returns to scale—to total input—which can be motivated in at least two different ways.

The first argument for increasing returns simply rehearses this as the primary reason for the existence of teams or firms in the first place, i.e., two agents working together can achieve more

[2] In this [16-20, 23, 43, 46, 47, 68] are closest in spirit, while more conventional approaches include [31, 37-41, 48, 51-53, 55, 57, 62, 63, 67].

than the sum of their efforts working alone [1, 11, 14, 15, 49, 50, 66, 69]. Formally, this amounts to $Q(x_i + x_j) \geq Q(x_i) + Q(x_j)$. A simple parameterization of this is $Q(X) = aX^{\beta}$, where $\beta \geq 1$ and a is a constant of proportionality; the limiting case of $\beta = 1$ amounts to constant returns to scale.

A different motivation for increasing returns gets more inside the black box of technology and pricing and is closer to the motivation of our agents as software objects deployed in an e-commerce environment. Consider a homogeneous good being supplied in large quantity by a price-discriminating seller at heterogeneous prices. The seller is willing to give quantity discounts in order to increase revenue and keep its large capacity supply process in high utilization. Individual agents show up and bid to purchase the good, usually in relatively small quantities. Individuals bids are either accepted or not, meaning low bids may never transact, providing the supplier with a way to distinguish the high value bidders. Specifically, say that the seller prices the good inversely proportional to the quantity bid on, q, like $q^{-(\beta-1)/\beta}$. This is shown in figure 1 for various values of the parameter.

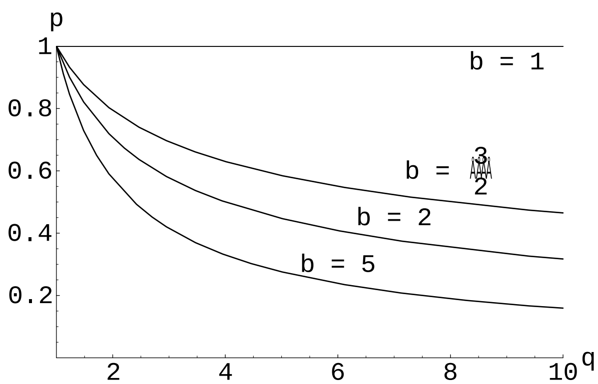

Figure 1: Dependence of price on quantity for various values of the parameter β

The seller's revenue is then $kq^{1/\beta}$, where k is a constant of proportionality. If a group of agents were to bid $X = kq^{1/\beta}$. they would receive quantity $Q(X) = aX^{\beta}$, where $a(k)$ is a constant. The is the bidding group's 'production function,' although it is perhaps somewhat unconventional to call it so.

These two distinct motivations yield the same function $Q(X)$. Substituting this relation into the expression for utility yields

$$U^i(x_i;\theta_i,e_i,X_{\sim i},a,\beta) = \left(ax_i(x_i + X_{\sim i})^{\beta-1}\right)^{\theta_i}(e_i - x_i)^{1-\theta_i} \quad (2)$$

Purposive agents will seek to make contributions to production that increase their utilities. Rational agents are presumed to contribute exactly the amount that maximizes utility. Our assumption of non-cooperative behavior means that no binding side agreements are made between agents, that each is free to act in its own self-interest; for a cooperative approach to group formation see [35, 36].

1.1.1 Nash Equilibrium Within a Team

We will model agent behavior as intermediate between mere purposiveness, e.g., gradient-like groping for utility improvements, and full-blown rationality in which each agent must have an internal model of every other agent, etc. Rather, we will invoke a 'best reply' type of adjustment dynamic, as is common in evolutionary game theory. Each agent knows its preferences and endowment, and the total team input in the previous period. It then uses this information to compute the best input in the next period, assuming that the overall input by its teammates will not change. That is, it computes

$$x_i^*(\theta_i,e_i,X_{\sim i},a,\beta) = \arg\max_{x_i} U^i(x_i;\theta_i,e_i,X_{\sim i},a,\beta). \quad (3)$$

In the case of $\beta = 2$, use of elementary calculus reveals the optimal contribution to be

$$x_i^*(\theta_i,e_i,X_{\sim i}) = \frac{2\theta_i e_i - X_{\sim i} + \sqrt{4\theta_i e_i X_{\sim i}(1+\theta_i) + (X_{\sim i} - 2\theta_i e_i)^2}}{2(1+\theta_i)} \quad (4)$$

This expression is increasing in endowment and preference for income, and decreasing in the contributions of the other agents. That is, as other members of the team contribute more, the best response is to contribute less; this is shown graphically in figure 2.

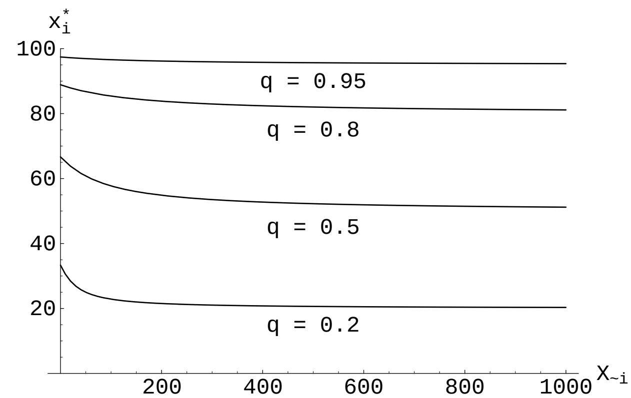

Figure 2: Dependence of optimal input on the contributions of other team members; agent i's endowment is 100

The agent doing so takes advantage of the others' efforts and by virtue of increasing returns can increase its individual welfare by saving some of its input. Thus there is an essential tension between cooperating for the good of the team and free-riding [12, 34]. In the limit as $X_{\sim i} \to \infty$, the optimal input level is simply $\theta_i e_i$.

1.1.1.1 Existence and Uniqueness of Equilibrium

For the input adjustment mechanism given by the previous expression, it can be shown that a unique Nash equilibrium exists. This is a direct consequence of the contraction mapping theorem. Since individual input levels are monotone decreasing in the team input[3], and thus decreasing in the inputs of all other agents, iterating equation (4) converges geometrically to a fixed point, which is necessarily unique. Thus, any group of agents having any composition will always have a Nash equilibrium in inputs from its members. No agent must know, *a priori*, what the Nash equilibrium of its group is, but by each agent making stepwise improvements in its input there results a Nash equilibrium.

[3] This will be made more precise in the next section.

It can also be shown that agents under-contribute input at the Nash equilibrium in comparison to the Pareto optimal inputs [33]. It is easiest to see this with homogeneous agents.

1.1.1.2 Identical Agents

Here we consider that each agent has the same endowment, e, and preferences, θ, and that there are N agents in the group. In this situation we substitute $(N-1)x$ for $X_{\sim i}$ into (4) and solve for x, yielding

$$x^{Nash}(\theta, e, N) = \frac{\theta e(N+1)}{N+\theta}. \qquad (5)$$

This Nash equilibrium is symmetric, i.e., like agents behave identically. Now it is easy to see that in the limit of infinite group size the Nash effort level is θe.

If instead we substitute $(N-1)x = X_{\sim i}$ into (3) and then maximize utility we obtain the symmetric[4] Pareto optimal input contribution levels of

$$x^{Pareto}(\theta, e) = \frac{2\theta e}{1+\theta}. \qquad (6)$$

It is not hard to show that the Pareto optimal inputs are always larger than the Nash levels, and tedious but demonstrable that agent welfare is greater for inputs given by (6) versus those produced by (5). However, Pareto input levels are not individually-rational—if a group is placed in any Pareto optimal configuration, each agent finds it welfare-improving to unilaterally reduce its contribution until the Nash level obtains.

1.1.2 Stability of Nash Equilibrium

Despite the ostensibly strong result regarding the existence and uniqueness of Nash equilibria in any group, and that this can be achieved via best reply type dynamics, it remains to be shown that such results are robust to a variety of perturbations, e.g., noise. This is possible in the present case due to the continuous nature of the action space involved.

Let us first be explicit about how input is adjusted:

$$x_i^{t+1}\left(\theta_i, e_i, X_{\sim i}^t\right) = \frac{2\theta_i e_i - X_{\sim i}^t + \sqrt{4\theta_i e_i X_{\sim i}^t\left(1+\theta_i\right) + \left(X_{\sim i}^t - 2\theta_i e_i\right)^2}}{2\left(1+\theta_i\right)}$$

To assess the stability of this system the Jacobian matrix must be computed and its eigenvalues determined. This calculation is straightforward, yielding for $i \neq j$,

$$\left.\frac{\partial x_i^*}{\partial x_j}\right|_{x_j = x_j^*} = \frac{-1 + \dfrac{X_{\sim i} + 2\theta_i^2 e_i}{\sqrt{4\theta_i e_i X_{\sim i}\left(1+\theta_i\right) + \left(X_{\sim i} - 2\theta_i e_i\right)^2}}}{2\left(1+\theta_i\right)},$$

while the entry is 0 for $i = j$. Clearly each of these is negative. Now, given that the rows of the Jacobian are identical, the dominant eigenvalue is just (N-1) the smallest (most negative) entry. It can be shown that

$$\lim_{N \to \infty} (N-1)\frac{\partial x_i^*}{\partial x_j} = 0$$

meaning that the Nash equilibria are always stable.

1.2 Proportional Reward + Equal Shares

In most team production environments it is usually either difficult or impossible to accurately assess individual contributions to production in a timely manner [22, 25, 27, 58, 64]. Imagine a baseball or soccer (football) team, for which certain statistics are kept—runs batted in, batting average with men in scoring position, or goals scored/game—but for which it is not generally possible to disaggregate the final outcome—games won—into purely individual contributions.

Even in the case of multi-agent bidding on a homogeneous good, there may be costs at the group level that serve to reduce overall consumable output, leading to non-trivial cost allocation problems, e.g., should the overhead be allocated according to equal shares or in proportion to inputs?

Such considerations lead us to view the previous subsections compensation function to be a very special case, and one wonders how the introduction of more realistic features will alter the results. To accomplish this we specify that overall compensation in any particular group will be a linear combination of proportional reward and equal shares. This results in agent utility having the form

$$U^i\left(x_i; \theta_i, e_i, X_{\sim i}, N, \alpha\right) = \left[\left(\frac{\alpha x_i}{x_i + X_{\sim i}} + \frac{1-\alpha}{N}\right)Q\left(x_i; X_{\sim i}\right)\right]^{\theta_i}\left(e_i - x_i\right)^{1-\theta_i}$$

where α specifies the specific mixture of compensation. Substituting the increasing returns production function gives

$$U^i\left(x_i; \theta_i, e_i, X_{\sim i}, N, a\right) = \left[\alpha x_i\left(x_i + X_{\sim i}\right) + (1-\alpha)\frac{\left(x_i + X_{\sim i}\right)^2}{N}\right]^{\theta_i}\left(e_i - x_i\right)^{1-\theta_i}$$

1.1.1 Nash Equilibrium Within a Team

We solve for Nash equilibrium as before, with each agent taking the other agents' inputs as given. This yields the following rather unwieldly expression:

$$x_i^*\left(\theta_i, e_i, X_{\sim i}, \alpha\right) = \frac{1}{2\lambda\left(1+\theta_i\right)}\left\{2\theta_i e_i \lambda - X_{\sim i}\left(\lambda + 1 - \alpha\right) + \sqrt{\left[X_{\sim i}\mu - 2\lambda\theta_i e_i\right]^2 + 4\lambda X_{\sim i}\left(1+\theta_i\right)\left[\theta_i e_i \mu + X_{\sim i}(1-\alpha)\left(1-\theta_i\right)\right]}\right\}$$

where $\lambda = 1 + \alpha(N - 1)$ and $\mu = \lambda + 1 - \alpha$. Next we show graphically how the optimal effort level depends on other agent effort. Figure 3 is directly analogous to figure 2 above, with an equal mixture of the two compensation systems, i.e., $\alpha = 1/2$.

[4] There are other, asymmetric Pareto optimal configurations as well.

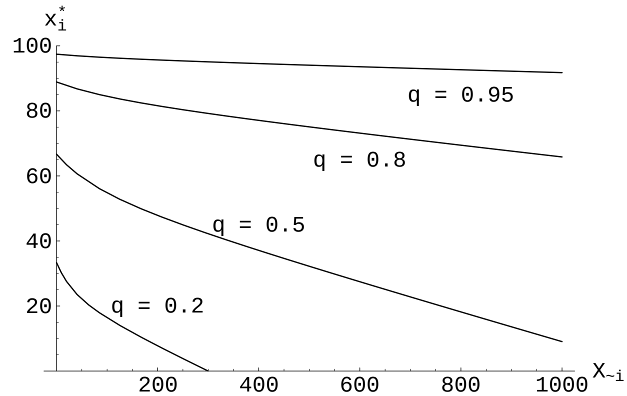

Figure 3: Dependence of optimal input on the contributions of other team members; agent *i*'s endowment is 100 and N = 10

The fact that each of the lines in figure 3 has negative slope, while the analogous lines plateau in figure 2, means that agents with this compensation will always have a tendency to free ride as their group size grows.

1.1.1.1 Existence and Uniqueness of Equilibrium

As in the previous case, the general character of equilibrium can be easily obtained through the contraction mapping theorem, guaranteeing that in any group the adjustment dynamics will quickly (geometrically) converge to the unique Nash equilibrium configuration.

1.1.1.2 Identical Agents

For a homogeneous agent population, substituting $(N-1)x$ for $X_{\sim i}$ leads to the following symmetric Nash equilibrium:

$$x^{Nash}(\theta, e, N, \alpha) = \frac{\theta e(\lambda + 1)}{N + \theta(\lambda + 1 - N)}$$

Note that in the case of α = 1, which implies $\lambda = N$, the previous result on identical agents is recovered. The symmetric Pareto optimal solutions are the same as before, and again involve more effort than the Nash configuration while not being individually-rational:

1.1.2 Stability of Nash Equilibrium

It can be shown that for any distribution of agent types in a group there exists a maximum group size beyond which individual agent contribution adjustment dynamics are dynamically unstable. That is, if a group is larger than its maximum stable size, any perturbation of the Nash level of contributions leads to a series of individual adjustments on the part of all the agents that has no rest point. To say this yet another way, the Nash equilibrium exists in such circumstances but it is unstable.

As before, compute the eigenvalues of the Jacobian matrix by differentiating the expression for Nash input level with respect to other agent input. This gives a very complicated result but one for which it can be shown that

$$\lim_{N \to \infty} \frac{\partial x_i^*}{\partial x_j} < 0 .$$

Therefore, for large enough N the dominant eigenvalue will be less than –1 and the Nash equilibrium unstable. So there exists a maximum stable group size, N^{max}, such that for $N > N^{max}$ the group is dynamically unstable to infinitesimal perturbations.[5] The stability boundary is shown in figure 4 for various α.

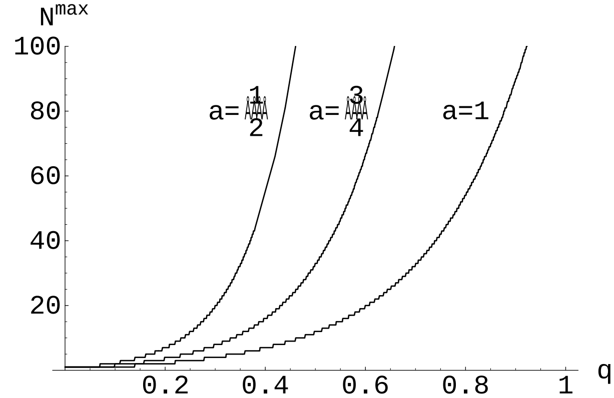

Figure 4: Stability diagram for mixed compensation, various α; region below and to curves represent stable group sizes

Conceptually, the existence of unstable Nash equilibria means simply that the behavior of the agents in the model is not static. Rather, the agents rearrange themselves into some configuration that may possess stationary aggregate behavior, even though the individual agents are perpetually adjusting their behavior, adapting to new circumstances [30]. Unfortunately, it is very difficult to analyze such intrinsically transient circumstances analytically. At least one reason for the pervasive focus of conventional game theory on static (Nash) equilibria is certainly due to the mathematical difficulties surrounding the systematic analysis of dynamical models. While game theoretic models having non-static solutions—sometimes called equilibria, although this terminology is certainly problematical—have been around for a long time (e.g., [56], [28]), today there is no generally applicable solution concept for them [59, 60].

3. REALIZATIONS WITH AGENTS

Since the Nash equilibria in this model are dynamically unstable, explication of the comparative statics of such equilibria is not informative. However, because the overall system of agents is highly nonlinear, it is also not the case that the unstable group dynamics 'blows up' in the sense that some quantity or another accelerates out toward infinity. Rather, in the same way that a fluid flow arranges itself into eddies as it moves into the turbulent regime, becoming irregular at the microscopic level while possessing stationary macroscopic statistics, the population of agents will arrange itself into self-organized groups that persist for some time before declining and ultimately exiting. But what are the characteristics of these endogenously formed groups?

Mathematical analysis of these transiently-lived groups is surely very difficult. Instead, we study this with a multi-agent system. Each agent has a Cobb-Douglas utility function and starts life working alone, as a singleton. Then each period each agent has an equal probability of awakening, in which case it recomputes its best input level for the following period, taking last period's behavior as given. It also selects a team at

[5] Another model of organizations on the edge of stability is [42].

random and considers how much utility it would derive from working there at its best effort level, and then chooses the option yielding the most utility overall [21]. The base case parameterization of the model is shown in Table 1.

Table 1. Base case configuration of the MAS implementation of the team formation model

Parameter	Value
a	*1*
β	*2*
A	10^6
θ_i	*Uniform (0, 1)*
e_i	*1*

After an initial transient epoch, the model enters into a stationary configuration in which, at the population level there result stationary distributions of agent input levels and utilities, group sizes and lifetimes, as well as group growth rates and related statistical measures. This aggregate level stationarity occurs despite constant flux and adaptation at the agent level. Some of these statistics are summarized presently.

No group is stable in this model, since agents can join any team where they can derive more utility, and successful teams grow beyond their stable size, However, there is a relatively smooth distribution of group lifetimes, as shown in figure 5.

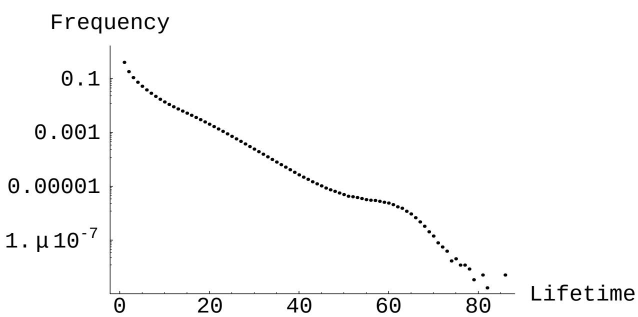

Figure 5: Stationary distribution of group lifetimes realized

The distribution of groups by size in the model is highly skewed, closely following a power law distribution having slope of –1 in doubly log coordinates, as depicted in figure 6.

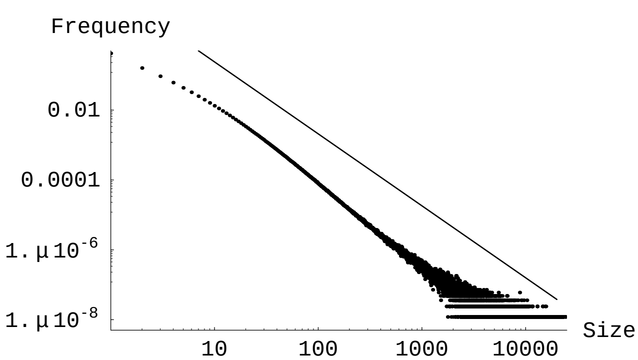

Figure 6: Stationary distribution of group size realized

The *qualitative* nature of these results is robust to alternative parametric specifications of the model, although the *quantitative* character does depend on specific parameters.

The fact that this model closely reproduces many of the features of business firm data is perhaps its best defense. However, having said this, the fact that the model utilizes boundedly rational agents and is intrinsically out of equilibrium places it at odds with conventional game theory. If the model yielded only theoretical propositions having unknown empirical relevance then it would seem that the 'burden of proof' would be to further demonstrate the value of departing from convention. However, insofar as the model is grounded empirically the 'burden' seems shifted onto those who would defend static equilibria and hyper-rationality.

4. SUMMARY AND CONCLUSIONS

A dynamical theory of endogenous team formation and evolution has been described. Conventional game theory is ill-suited to studying the kinds of meta-stable structures that emerge and transiently survive in this model. MAS realizations of the full transient structure of this model yield interesting stationary structures and behavioral regularities at the aggregate level. Given that these aggregate properties are closely related to empirical properties of firms, we have argued that to limit the focus of one's analysis to equilibria, while certainly augmenting the mathematical tractability, is both highly restrictive and unrealistic, and likely to render the resulting models empirically false and operationally brittle.

The model of team formation has recently been extended to the context of city sizes [6], another topic for which there is extensive empirical data to use as a target of our model-building. Giving each firm a location, we stipulate that agents who team up must all occupy the same location. Then, when an agent moves to another location it must migrate physically to the new firm's location. Lastly, when an agent starts up a new firm it stays with high probability in its present location, but with some small probability it selects a new location at random. This simple re-specification of the firm formation model is sufficient to yield a Zipf distribution of agent and firm agglomerations over space. As is well known, the Zipf distribution describes the size distribution of cities in most industrial countries. This model is the first microeconomic model to plausibly yield the Zipf distribution. Once again, an intrinsically dynamical model succeeds where microeconomic equilibrium models have failed.

While there is no coherent solution 'concept' for dynamical situations in conventional game theory, it is also true that there is no prohibition from investigating games in which equilibrium solutions are never realized. For the purposes of positive political economy studied above, we use a MAS to resolve the dynamical behavior, in which there is perpetual flux and adjustment at the micro level while stationary aggregate performance obtains. To limit a game theoretic analysis to merely Nash equilibrium is to severely circumscribe the range of applicability of this otherwise fertile theory. To base the analysis of multi-agent systems on conventional game theory is to so restrict the possible vocabulary of analysis as to impoverish the design of operational systems. Why cannot groups of agents spontaneously band together on the Internet to form a multi-agent coalition for purposes of, for example, extracting better pricing from a product supplier? Or would not a self-organized team of autonomous agents make a far more robust line of attack against invaders than would a single strong but potentially brittle agent? In the end, as Simon's critique [61] of Rubinstein makes clear, progress in the positive social sciences can only be had, ultimately, from empirical work, no matter how beautiful is one's mathematical theory.

5. ACKNOWLEDGMENTS

Thanks are due the John D. and Catherine T. MacArthur Foundation, the National Science Foundation (IIS-9820872), and the Alex C. Walker Foundation for financial support. Comments from Blake LeBaron and H. Peyton Young are gratefully appreciated.

6. REFERENCES

[1] Alchian, A.A. and H. Demsetz. 1972. "Production, Information Costs, and Economic Organization." *Am. Econ. Rev., 62*: 777-795.

[2] Arthur, W.B., S. Durlauf, D. Lane, eds. 1997. *The Economy as an Evolving Complex System, Volume II*: Addison-Wesley: N.Y.

[3] Axtell, R.L. 1999. "The Emergence of Firms in a Population of Agents." Working paper 99-03-019. Santa Fe Institute: Santa Fe, New Mexico. Available at www.brook.edu/es/dynamics/papers.

[4] Axtell, R.L. 2000. "Why Agents?" In Sallach et al., eds. Argonne National Laboratory: Chicago, Illinois.

[5] Axtell, R.L. 2001. Zipf Distribution of U.S. Firm Sizes. *Science, 293*: 1818-1820.

[6] Axtell, R.L. and R. Florida. 2000. "The Evolution of Cities: A Microeconomic Explanation of Zipf's Law." Working paper. The Brookings Institution and Carnegie Mellon University.

[7] Binmore, K.G. 1987. "Modeling Rational Players, Part I." *Economics and Philosophy*, *3*: 179-214.

[8] ———. 1988. "Modeling Rational Players, Part II." *Economics and Philosophy*, *4*: 9-55.

[9] Blume, L.E. and D. Easley. 1998. "Optimality and Natural Selection in Markets." Working paper #98-09-082. Santa Fe Institute: Santa Fe, N.M.

[10] Bolton, P. and M. Dewatripont. 1994. "The Firm as a Communication Network." *Q.J.E.*, *109*: 809-839.

[11] Buchanan, J.M. and Y.J. Yoon. 1994. *The Return to Increasing Returns*. U. Michigan Press: Ann Arbor, Mich.

[12] Canning, D. 1995. "Evolution of Group Cooperation through Inter-Group Conflict." Working Paper. Department of Economics, Queens University of Belfast, Northern Ireland.

[13] 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: N.Y.

[14] Coase, R. 1937. "The Nature of the Firm." *Economica, n.s., 4:* 386-405.

[15] 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.

[16] DeCanio, S.J. and W.E. Watkins. 1998. "Information Processing and Organizational Structure." *JEBO, 36*: 275-294.

[17] DeVany, A. 1993a. "Edgeworth and the Traveling Salesman: The Complexity of Coalition Structures." Working paper. Department of Economics, University of California, Irvine.

[18] ———. 1993b. "Hard Cores and Soft Cores: Evolving Coalitions, Edgeworth and the Boltzmann Machine." Working paper. Department of Economics, University of California, Irvine.

[19] ———. 1993c. "Information, Bounded Rationality, and the Complexity of Economic Organization." Working paper. Department of Economics, University of California, Irvine.

[20] ———. 1996. "The Emergence and Evolution of Self-Organized Coalitions." In Gilli, ed.

[21] Dong, Xioa-yuan and G. Dow. 1993a. "Does Free Exit Reduce Shirking in Production Teams?" *Journal of Comparative Economics, 17*: 472-484.

[22] ——— and ———. 1993b. "Mutual Monitoring in Production Teams." *Journal of Political Economy*.

[23] Dow, G. 1990. "The Organization as an Adaptive Network." *J. Econ. Beh. and Org.*, *14*: 159-185.

[24] Dresher, M., L.S. Shapley, and A.W. Tucker. 1964. Advances in Game theory. Annals of Mathematical Studies, No.52. Princeton University Press: Princeton, N.J.

[25] 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.

[26] Epstein, J.M. and R. Axtell. 1996. Growing Artificial Societies: Social Science from the Bottom Up. The Brookings Institution Press/MIT Press: Washington, D.C./Cambridge, Mass.

[27] Farrell, J. and S. Scotchmer. 1988. "Partnerships." *Quarterly Journal of Economics*, *103* (2): 279-297.

[28] Gilboa, I. and A. Matsui. 1991. "Social Stability and Equilibrium." *Econometrica*, *59* (3): 859-867.

[29] Gilli, M. 1996. *Computational Economic Systems: Models, Methods and Econometrics*. Kluwer Pub.: Boston, Mass.

[30] Glance, N.S., T. Hogg, and B.A. Huberman. 1997. "Training and Turnover in the Evolution of Organizations." *Org. Sci., 8*: 84-96.

[31] Hart, S. and M. Kurz. 1983. "Endogenous Formation of Coalitions." *Econometrica*, *51*: 1047-1064.

[32] Holler, M.J., ed. 1984. *Coalitions and Collective Action*. Physica-Verlag: Wuerzburg, Vienna, Austria.

[33] Holmstrom, B. 1982. "Moral Hazard in Teams." Bell Journal of Economics, 13: 324-340.

[34] Huberman, B.A. and N.S. Glance. 1998. "Fluctuating Efforts and Sustainable Cooperation." In Prietula *et al*., eds.

[35] Ichiishi, T. 1977. "Coalition Structure in a Labor-Managed Market Economy." *Econometrica*, *45*: 341-360.

[36] ———. 1993. *The Cooperative Nature of the Firm*. Academic Press: New York.

[37] 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.

[38] Klusch, M. and O. Shehory. 1996a. "Coalition Formation Among Rational Information Agents." In *Proc. of the Seventh European Wkshp. on Modeling Autonomous Agents in a Multi-Agent World* (MAAMAW): 204-217.

[39] —— and ——. 1996b. "A Polynomial Kernel-Oriented Coalition Algorithm for Rational Information Agents." In Tokoro, ed.: 157-164.

[40] Laussel, D. and M. LeBreton. 1995. "A General Equilibrium Theory of Firm Formation Based on Unobservable Skills." *European Economic Review*, *39*: 1303-1319.

[41] Legros, P. and A.F. Newman. 1996. "Wealth Effects, Distribution, and the Theory of Organization." *J.Ec.Th, 70*: 312-341.

[42] Levitan, B., J. Lobo, S. Kauffman, and R. Schuler. 1999. "Optimal Orgnaizational Size in a Stochastic Envirnoment with Externalities." Working paper 99-04-024. Santa Fe Institute: Santa Fe, New Mexico.

[43] Luna, F. 2000. "Induction and Firm Creation." In Luna and Steffanson, eds.

[44] Luna, F. and B. Steffanson, eds. 2000. *Multi-Agent Simulation in Economics and Finance Using SWARM*. Kluwer.

[45] Majumdar, M. 1996. *Organizations with Incomplete Information*. Cambridge University Press: New York.

[46] Mazzucato, M. 1998. "A Computational Model of Economies of Scale and Market Share Instability." *Structural Change and Economic Dynamics*, *9*: 55-83.

[47] Miller, J. 1998. "Evolving Information Processing Organizations." Working paper #95-06-053. SFI: Santa Fe, N.M.

[48] Olson, M. 1965. *The Logic of Collective Action*. Harvard University Press: Cambridge, Massachusetts.

[49] Ostrom, E. 1990. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge U. Pr: N.Y.

[50] Padgett, J.F. 1997. "The Emergence of Simple Ecologies of Skill: A Hypercycle Approach to Economic Organization." In Arthur *et al*., eds.

[51] Putterman, L. 1998. "The Firm as Association Versus the Firm as Commodity." *Econ. and Phil., 4*: 243-266.

[52] Radner, R. 1993. "The Organization of Decentralized Information Processing." *Econometrica*, *61*(5): 1109-1146.

[53] Roth, A. 1984. "Stable Coalition Formation: Aspects of a Dynamic Theory." In Holler, ed.

[54] Rubinstein, A. 1998. Modeling Bounded Rationality. MIT Press: Cambridge, Mass.

[55] 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*.

[56] Shapley, L. 1964. "Some Topics in Two-Person Games." In Dresher et al., ed.

[57] Shehory, O. and S. Kraus. 1993. "Coalition Formation Among Autonomous Agents: Strategies and Complexity." In Castelfranchi and Müller, eds. [13].

[58] Sherstyuk, K. 1998. "Efficiency in Partnership Structures." *J. Economic Behavior and Organization, 36*: 331-346.

[59] Shubik, M. 1999. Personal communication.

[60] ———. 1997. "Why Equilibrium? A Note on the Noncooperative Equilibria of Some Matrix Games." *JEBO*, *29*: 537-539.

[61] Simon, H.A. 1998. In Rubinstein [54].

[62] Van Zandt, T. 1996. "Organizations with and Endogenous Number of Information Processing Agents." In Majumdar, ed. [45].

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

[64] Varian, H. 1990. "Monitoring Agents with Other Agents." *J. of Inst. and Theo. Econ., 46* (1): 153-174.

[65] Weiß, G. 1999. *Multi-Agent Systems*. MIT Press: Cambridge, Mass.

[66] Williamson, O.E. and S. Winter. 1997. *The Nature of the Firm*. Cambridge University Press: New York, NY.

[67] Wilson, R. 1968. "The Theory of Syndicates." *Econometrica*, *36* (1): 119-132.

[68] Winter, S.G. 1993. "On Coase, Competence and the Corporation." In Williamson and Winter, eds.

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

[70] Young, H.P. 1998. *Individual Strategy and Social Structure*. Princeton University Press: Princeton, New Jersey.