ECONOMIC ANALYSIS AND WORKERS' MANAGEMENT, 3, XXII (1988), 139—167

THE EVOLUTION OF ORGANIZATIONAL FORM: SELECTION, EFFICIENCY, AND THE NEW INSTITUTIONAL ECONOMICS

*Gregory K. DOW**

Abstract: Economists frequently claim that organizational forms which have survived a competitive natural selection process will have favorable efficiency properties. Without third-party contract enforcers, however, selection forces will operate upon the strategies of individual factor suppliers rather than upon organizational structures. This process of *dual selection* often generates inefficient but evolutionarily stable equilibria. The Pareto efficiency of equilibrium structures can be assured only in a regime of complete contracts, where selection forces operate upon organizational structures as units. Since contracts among factor suppliers are usually incomplete, surviving organizational forms will often constitute evolutionary market failures.

1. INTRODUCTION

1.1. Evolution and Efficiency

When asked to explain why production activities are organized as they are, economists often suggest that efficiency requirements would not otherwise be satisfied.[1] If pressed for some reason why efficient organizational forms should arise and persist, economists frequently fall back upon a natural selection argument: existing organizations have had to pass competitive survival tests, and would not have passed these tests without possessing some efficiency advantage. Surviving firms will be organized 'as if' they were consciously designed according

* University of Alberta — Department of Economics, Edmonton, Canada.

* Paper prepared for the session on Explanations of the Roles of Capital and Labor within a Firm sponsored by the Union of Radical Political Economics at the Allied Social Science Associations Annual Meeting, Chicago, December 28, 1987.

[1] This functionalist argument parallels attempts in sociology to explain organizational structures via contingency theory (Pfeffer, 1982, Ch. 4). Functionalist social theory is rejected by Elster (1983, Ch. 2), but van Parijs (1981, Ch. 2) takes a more favorable view. For an evolutionary view of organizational sociology, see McKelvey and Aldrich (1983).

to the Pareto criterion, much as Friedman (1953) defended the orthodox axiom of profit maximization by appealing to selection forces.

Transaction cost economics (TCE), in particular, operates from the premise that extant governance structures are explicable primarily in efficiency terms, and relies on selection arguments to buttress this premise (Williamson, 1975, Ch. 8; 1985: 22—23).[2] My procedure in this paper will be to assess TCE's efficiency premise in the context of some specific evolutionary mechanisms. As might be expected by aficionados of the new institutional economics, I conclude that the efficiency of evolved structures hinges on the precise nature of the contracting institutions available to factor suppliers.

Call a world without third-party contract enforces *private ordering,* and call a world where binding contracts are feasible *legal centralism* (these terms are borrowed from Williamson, 1985: 20—21; 1986). Distinct selection processes will operate in these two institutional environments:

(a) *Dual Selection.* If the variation on which selection forces operate is variation in the individual *strategies* used by workers or capitalists, as would be expected in a world of *private ordering,* evolved governance structures will often prove inefficient.

(b) *Joint Selection.* If the variation on which selection forces operate is variation in *strategy bundles* — that is, variation in the joint strategies of workers and capitalists — as would be expected in a world of *legal centralism,* then evolved governance structures will be Pareto efficient.

Thus, if TCE wishes to rely on »the efficacy of competition to perform a sort between more and less efficient modes (of organization) and to shift resources in favor of the former« (Williamson, 1985: 22), then it must grant a prominent role to legal centralism. If this is regarded as a violation of the stylized facts,[3] then (i) TCE could adopt a more intentionalistic justification for its efficiency premise, along the lines of the Coase Theorem, and forego the evolutionary rationale; or (ii) TCE could embrace evolutionary selection as its core explanatory principle, while refraining from general efficiency claims.

The remainder of this section argues that the new institutional economics requires invisible-hand theories of organizational form. Section 2 introduces some game-theoretic notation and formalizes the

[2] Reliance on evolutionary arguments to support efficiency or optimality claims is by no means confined to transaction cost theorists. For an optimal contracting example, see Stiglitz (1975); an agency theory view is provided by Jensen (1983). Selection arguments are particularly prominent in TCE, because TCE's emphasis on bounded rationality undercuts more intentionalistic rationales for the efficiency premise (Dow, 1987a).

[3] We note that »transaction cost economics poses the problem of economic organization as a problem of contracting,« but »most disputes ... are resolved by avoidance, self-help, and the like« (Williamson, 1985: 20). The first remark points toward legal centralism, while the second points toward private ordering.

concept of organizational structure as a strategy bundle. The dual selection process, which I associate with a world of private ordering, is modeled in section 3. I show that stable evolutionary equilibria for this process are frequently inefficient. The joint selection process of section 4, which is plausible for a world of legal centralism, is then shown to yield efficient organizational outcomes. In section 5, I close with some comments on the implications of these results for transaction cost and Marxian views of the firm.

1.2. *Institutions, Rules, and Behavior*

There is an intimate connection between invisible-hand explanations, of which evolutionary models are one variant,[4] and the new institutional economics. Langlois offers the following definition of an institution:

> Institutions ... are orderly and more or less persistent behavior patterns. At a more abstract level, they are the rules or sets of rules that constrain or govern organized patterns of behavior. In either case, institutions are structures. (1986: 247).

In particular, firms are institutions whose organizational structures are defined by the persistent behavior patterns of factor suppliers engaged in production activities.

Langlois' reference to rules that constrain or govern behavior is amplified by Schotter:

> There are basically two views of institutions. In the first, which I shall call the *rules* view, social or economic institutions are seen as sets of rules that *constrain* individual behavior ... [S]ocial institutions are planned and designed mechanisms given exogenously to or imposed upon a society of agents. Institutional change is a process of social engineering that takes place through the manipulation of the rules.
>
> The other view of social institutions, which I shall call the *behavioral* view, [interprets] social institutions not as sets of predesigned rules, but rather as unplanned and unintended regularities of social behavior ... that emerge »organically« ... Institutions are outcomes of human action that no single individual intended to occur. (1986: 117—118; emphasis added).

For a parallel dichotomization of the organization theory literature, see Dow (1988a).

Given the assumed absence of intentionality, institutions in the behavioral sense call out for invisible-hand explanations:

---

[4] Unless otherwise indicated, I use the term 'evolutionary' in the broad sense of van Parijs (1981). This includes both Darwinian selection models involving differential mortality or reproduction rates for actors of distinct types, and also models where alternative behavioral traits of given actors are differentially reinforced.

> An invisible-hand explanation explains a well-structured social pattern or institution. It typically replaces an easily forthcoming and initially plausible explanation according to which the explanandum phenomenon is the product of intentional design with a rival account according to which it is brought about through a process involving the separate actions of many individuals who are supposed to be minding their own business unaware of and a fortiori not intending to produce the ultimate overall outcome. (Ullmann-Margalit, 1978: 267).

Such explanations will often treat *other* background institutions as exogenously given rules of the game. For example, extant contracts may bind certain actions taken by factor suppliers. But if these contracts are incomplete, then we might want to construct an invisible hand model which explains how the contractual gaps are filled in. For this purpose, the specific provisions of prevailing contracts play a role as background institutions of the 'rules' variety, alongside the more general rules imposed by statutes and judicial precedents (Masten, 1986).

2. EVOLUTION, GAME THEORY, AND ORGANIZATIONAL FORM

Evolutionary biology and game theory have recently engaged in considerable cross-fertilization (Maynard Smith, 1982; Axelrod, 1984; Friedman and Rosenthal, 1986). This section exploits these links by defining an organizational structure to be a vector of strategies in a non-cooperative game played among factor suppliers. I then sketch out a selection process by which these strategies might evolve over time in a world of private ordering. This process is subsequently formalized in section 3.

2.1. *Organizational Structures as Strategy Bundles*

Consider a firm consisting of one capital supplier (K) and one labor supplier (L). Assume that these agents engage in collaborative production activities over some significant time interval due to asset specificity. The firm's organizational structure will have two components: those elements of K and L's behavior which are contractually pre-specified, and any residual patterns of behavior not governed by legally binding contracts. Call these the *contractual* and *strategic* components of organizational structure.[5] When contracts are incomplete, the contractual lacunae will define some non-cooperative game in extensive form to be played by K and L. The strategic component of

[5] This resembles the distinction drawn between formal and informal structure in organization theory (Scott, 1981). The distinction made here is not identical, because some organizational attributes typically regarded as aspects of formal structure, such as authority relations, may have a strategic rather than a contractual basis (Dow, 1987b).

structure will be determined by the strategies describing K and L's behavior in this game whenever prevailing contracts are silent or unenforceable.[6]

Assume that we are given a formal description of the extensive-form game played by K and L as they fill in the gaps left by the prevailing contractual framework. [An example is provided by the authority game in Dow (1987b)]. We can then derive a game in normal form, as illustrated in Figure 1.

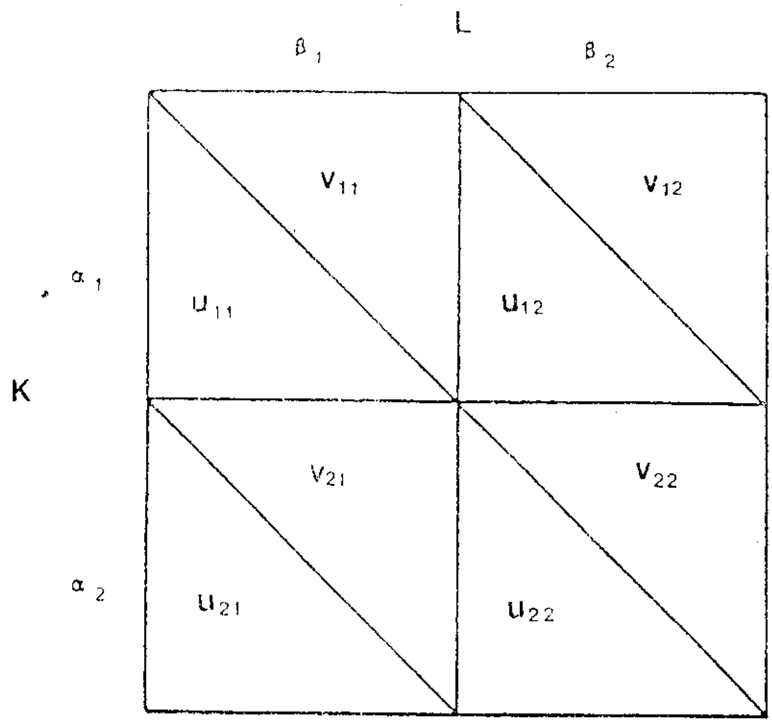

FIGURE 1

Organizational Structures in Normal Form

Denote this normal-form game by Γ, where the payoffs in Figure 1 are von Neumann-Morgenstern expected utilities. Since all essential points can be made in the context of 2×2 games, I confine attention to games of this type. Each of the four pure strategy pairs in Γ generates a distinctive pattern of behavioral interaction in the original extensive-form game. Pure-strategy vectors of the form (α_i, β_j) therefore specify alternative organizational structures and will be called *strategy bundles*.[7]

---

[6] The strategic choices of K and L in this game represent an exercise of residual ownership rights (Grossman and Hart, 1986), and therefore presuppose some set of background institutions defining the property rights of K and L.

[7] Production technology, market prices, extant contracts, and all relevant background institutions are summarized by the elements of the payoff matrix for Γ. Strictly speaking, the analysis in the text deals only with the strategic component of organizational structure, but since the terms of existing contracts are treated as exogenously fixed, I suppress this distinction for the sake of brevity.

Let $\sigma_K \in [0, 1]$ be the probability that K plays α_1, and let $\sigma_L \in [0, 1]$ be the probability that L plays β_1. Then K and L's payoffs from the mixed strategy pair $\sigma = (\sigma_K, \sigma_L) \in [0, 1]^2$ are

$$U(\sigma) = \sigma_K \sigma_L u_{11} + \sigma_K (1 - \sigma_L) u_{12} + (1 - \sigma_K) \sigma_L u_{21} + (1 - \sigma_K)(1 - \sigma_L) u_{22} \quad (1)$$

$$V(\sigma) = \sigma_K \sigma_L v_{11} + \sigma_K (1 - \sigma_L) v_{12} + (1 - \sigma_K) \sigma_L v_{21} + (1 - \sigma_K)(1 - \sigma_L) v_{22}$$

Definition: A *Nash equilibrium* for Γ is a probability pair $\sigma^* = (\sigma_K^*, \sigma_L^*)$ such that

$$U(\sigma_K^*, \sigma_L^*) \geq U(\sigma_K, \sigma_L^*), \text{ all } \sigma_K \in [0, 1]; \quad \text{and}$$

$$V(\sigma_K^*, \sigma_L^*) \geq V(\sigma^*, \sigma_L), \text{ all } \sigma_L \in [0, 1].$$

Nash equilibria may involve mixed strategies unless otherwise stated. As is well known (Owen, 1982: 126—129), the game Γ always has at least one Nash equilibrium in mixed strategies.

A simple example may help clarify this abstract framework. Suppose that K and L disagree from time to time over how the production process should be modified in response to unforeseen environmental shocks, and existing contracts are silent on the subject. On these occasions, if K insists on a production plan unfavorable to L, then L can adopt a strategy of striking (β_1) or not striking (β_2). Similarly, if L demands a response regarded as unfavorable by K, then K can use a lock-out (α_1) or refrain from a lock-out (α_2). Assume that if L is willing to strike but K will not use lock-outs, then L can decide how contractual gaps will be filled in. Conversely, assume that if K is prepared to lock out workers but L is unwilling to strike, then K can modify the production process unilaterally.[8]

For this model, cells (α_1, β_2) and (α_2, β_1) of Figure 1 assign de facto managerial authority to capital and labor, respectively. We might imagine that in cell (α_2, β_2) some neutral arbitrator is called in, while cell (α_1, β_1) involves open class warfare. Because the strategies in Γ are chosen non-cooperatively by assumption, Pareto inefficient outcomes are quite possible. For instance, Γ could be a prisoners' dilemma, as in the 'negotiation' equilibrium described by Dow (1987b).

2.2. *A Selection Mechanism for Organizational Forms*

Now consider an evolutionary model of the strategy choices made in the game Γ, with the following properties.[9]

---

[8] For a related model of strategic behavior, see Dow (1985).

[9] This evolutionary process is closely akin to Friedman and Rosenthal's (1986) model of emigration from and immigration to various pure strategies, and at a sufficiently abstract level becomes a special case of their approach. However, Friedman and Rosenthal are interested in *alternatives* to Nash equilibrium, whereas I am exploring the organizational *implications* of Nash equilibrium in an evolutionary context (see section 2.3).

1. There exist two large, internally homogeneous populations, the class of all capitalists and the class of all workers. Capitalists and workers are equally numerous.
2. At the start of round $t \in T = \{0, 1, \ldots\}$, each worker is randomly paired with some capitalist. The proportion of capitalists using strategy α_1 in round is $\sigma_K(t)$, and the proportion of workers using strategy β_1 in round t is $\sigma_L(t)$. The probability that any specific worker and capitalist are paired against each other is independent of the strategies used by each. Each pair of factor suppliers then plays the game Γ using their current strategies.
3. After round t, each capitalist observes the average payoffs obtained in round t by capitalists using the strategies α_1 and α_2. If one strategy yields a higher average payoff than the other, the fraction of capitalists using the high-payoff strategy will increase in round $t+1$.
4. Each worker observes the average payoffs obtained in round t by workers using the strategies β_1 and β_2, and the relative frequencies of worker strategies adjust in the same manner.
5. Steps 2—4 repeat in round $t+1$.

Notice that σ_K and σ_L in step 2 are now the *proportions* of the two classes using α_1 and β_1, rather than probabilities that given *individuals* will choose these strategies. No mixing among pure strategies by individual agents is assumed.

This evolutionary mechanism need not depict a Darwinian process where the differential mortality or reproduction of fit and unfit players is responsible for changes in the relative frequency of a given strategy. Instead, the selection mechanism rests upon the differing propensities of factor suppliers to adopt successful and unsuccessful strategies. In the scheme of van Parijs (1981), this is an R-evolutionary mechanism, involving differential reinforcement of strategies, rather than an NS-evolutionary mechanism involving Darwinian selection among the players themselves.[10]

The successive rounds of this process do nothing more than give the individual members of each class an opportunity to change their minds, in response to feedback on the aggregate performance of other agents like themselves. The folk theorem for repeated games (Fudenberg and Maskin, 1986) does not apply, because the probability of meeting the same opponent twice is negligible. We can thus ignore reciprocal altruism (Axelrod, 1984: 118—120) and reputation effects (Kreps, 1984). Indeed, to highlight the minimal cognitive capacities

[10] A Darwinian approach might be appropriate if the agents in the model are themselves coalitions or organizations. For example, failing corporations might be eliminated by bankruptcy or merger, while ineffective labor unions might be decertified by a vote of the relevant bargaining unit. Economic models of R-evolutionary and NS-evolutionary processes are provided by Schotter (1981) and Nelson and Winter (1982), respectively.

demanded of the players, we could dispense with the game-theoretic notion of 'strategies' entirely, substituting instead the evolutionary concept of 'routines' (Nelson and Winter, 1982).

2.3. *Why Use Evolutionary Explanations?*

Section 3 will show that for the dual selection mechanism sketched above, a given relative frequency pair $\sigma^* = (\sigma_K^*, \sigma_L^*)$ is an *evolutionary* equilibrium if and only if σ^* is a Nash equilibrium of the static game Γ in Figure 1. The central message of the paper then reduces to the following: in a world of private ordering, organizational forms *are* non-cooperative Nash equilibria, and since Nash equilibria need not be Pareto efficient, the same holds true for organizational forms. We thus need to forestall the following objection: why not just use the concept of Nash equilibrium directly, discarding the evolutionary trappings? Ignoring the possibility that disequilibrium dynamics may be intrinsically interesting, several responses can be given.

1. Although many writers believe that Nash equilibrium has strong a priori claims as a solution for non-cooperative games (Binmore and Dasgupta, 1986: 1—45), some skepticism lingers (Mirowski, 1986). Friedman and Rosenthal (1986), for example, propose a concept of populational steady states as an alternative to Nash equilibrium. It may therefore be more persuasive to derive the Nash outcome from a plausible evolutionary adjustment process than to postulate it as an axiom. We are, after all, concerned with the validity of *selection* arguments as a basis for efficiency claims.

2. The set of dual selection equilibria is never larger than the set of Nash equilibria, and is sometimes smaller when unstable equilibria are excluded. Stability criteria can thus be used to refine the Nash solution. Maynard Smith's (1982) concept of an evolutionarily stable strategy (ESS) already serves as a Nash equilibrium refinement for certain normal-form games (Binmore and Dasgupta, 1986: 18—19), and the dual selection mechanism extends the ESS approach.

3. Normal-form games often have many Nash equilibria. These non-uniqueness problems are resolved in the dual selection model by noting that from any given initial state, there is a unique Nash equilibrium toward which the selection process gravitates.[11] An observed equilibrium can be distinguished from a hypothetical alternative by the fact that the former is compatible with history, while the latter is not.

4. Nash equilibrium provides an invisible-hand model in the limited sense that no single player intends to bring about the equilibrium

[11] This follows from the deterministic nature of the selection process in section 3. In a stochastic process model (see Cross, 1983, for example), only a unique probability distribution over strategy bundles, perhaps dependent upon initial conditions, would be obtainable.

strategy vector. Nonetheless, the nature of the game being played is transparent to each player. By eliminating redundant rationality, evolutionary models provide 'true' invisible-hand explanations, where individual actors need not have any understanding of the forces acting to stabilize existing institutions. In Marxian language, these institutions will be opaque to the very actors whose behavior continually reproduces them. If invisible-hand explanations are deemed to be desirable in the social sciences (Ullmann-Margalit, 1978; Elster, 1983), then the evolutionary approach will have some methodological appeal.

3. PRIVATE ORDERING AND DUAL SELECTION

This section formalizes the selection process outlined in section 2.2. In a world of private ordering, where binding contracts do not exist, factor suppliers can be expected to look to the past experience of similar agents when choosing a strategy for the current round of production activities. Selection forces will therefore operate separately on each element of the strategy bundle, rather than upon strategy bundles as units. Stable evolutionary equilibria for this process will often be Pareto inefficient, for the same reasons that a Nash equilibrium in the original static game Γ may be inefficient.

3.1. *Profit-Responsive Dual Selection Mechanisms*

If the stochastic features of the underlying extensive form game are independent across firms in each round of the selection process, then an appeal to large numbers allows us to replace the *expected* payoffs in Figure 1 with *realized averages* across the capitalist and worker populations.[12] Capitalists who use the strategy α_1 in round t will obtain the average payoff

$$\bar{u}_1 [\sigma_L(t)] = \sigma_L(t)u_{11} + [1 - \sigma_L(t)]u_{12} \tag{2a}$$

and capitalists using α_2 will obtain the average payoff

$$\bar{u}_2[\sigma_L(t)] = \sigma_L(t)u_{21} + [1 - \sigma_L(t)]u_{22} \tag{2b}$$

In order to avoid lengthy payoff expressions, I shall adopt the following notation:

$$\hat{u}_1 = u_{11} - u_{21} \tag{3}$$

[12] We might also consider a single (K, L) pair who play the same game Γ many times. If the rounds are treated as independent of one another, then the law of large numbers will come into play for strategy revisions that reflect average performance over a large number of past trials (Crawford, 1985). Repeated game phenomena become more troublesome in this case because there is no random remixing of players.

$$\hat{u}_2 = u_{12} - u_{22}$$

$$\hat{v}_1 = v_{11} - v_{12}$$

$$\hat{v}_2 = v_{21} - v_{22}$$

$$\hat{u} = \hat{u}_1 - \hat{u}_2 = u_{11} + u_{22} - u_{12} - u_{21}$$

$$\hat{v} = \hat{v}_1 - \hat{v}_2 = v_{11} + v_{22} - v_{12} - v_{21}$$

The *average realized profit* obtained by the capitalists using strategy α_1 in round t can then be defined by

$$\pi_K[\sigma_L(t)] \equiv \bar{u}_1[\sigma_L(t)] - \bar{u}_2[\sigma_L(t)] \qquad (4)$$
$$= \sigma_L(t)\,\hat{u}_1 + [1 - \sigma_L(t)]\,\hat{u}_2$$

The *capitalist selection mechanism* will be written as

$$\sigma_K(t+1) - \sigma_K(t) = f_K\{\pi_K[\sigma_L(t)],\ \sigma(t)\}, \qquad \text{all } t \in T \qquad (5)$$

The inclusion of $\sigma(t)$ as an argument of f_K reflects the fact that σ_K is restricted to the interval [0, 1] regardless of the value taken on by π_K. I shall say that f_K is *profit-responsive* if

$$f_K(\pi_K, \sigma) > 0 \qquad \text{whenever } \pi_K > 0 \text{ and } \sigma_K < 1;$$
$$f_K(\pi_K, \sigma) < 0 \qquad \text{whenever } \pi_K < 0 \text{ and } \sigma_K > 0;$$
$$f_K(\pi_K, \sigma) = 0 \qquad \text{otherwise.} \qquad (6)$$

A mechanism of this type increases the relative frequency of the capitalist strategy having the higher average payoff in the preceding period, whenever this is feasible.

We define the *worker selection mechanism* similarly:

$$\pi_L[\sigma_K(t)] \equiv \sigma_K(t)\,\hat{v}_1 + [1 - \sigma_K(t)]\,\hat{v}_2; \qquad (7)$$

$$\sigma_L(t+1) - \sigma_L(t) = f_L\{\pi_L[\sigma_K(t)],\ \sigma(t)\}, \qquad \text{all } t \in T; \qquad (8)$$

$$f_L(\pi_L, \sigma) > 0 \qquad \text{whenever } \pi_L > 0 \text{ and } \sigma_L < 1;$$
$$f_L(\pi_L, \sigma) < 0 \qquad \text{whenever } \pi_L < 0 \text{ and } \sigma_L > 0;$$
$$f_L(\pi_L, \sigma) = 0 \qquad \text{otherwise.} \qquad (9)$$

For stability analysis, it is more convenient to work with a continuous time version of (5) and (8). Setting the time interval between rounds equal to Δt and taking limits, we obtain

$$\dot{\sigma}_K(t) = f_K\{\pi_K[\sigma_L(t)],\ \sigma(t)\}, \qquad \text{all } t \geq 0; \qquad \sigma_K(0) = \sigma_{K0}; \qquad (10a)$$

$$\dot{\sigma}_L(t) = f_L\{\pi_L[\sigma_K(t)],\ \sigma(t)\}, \qquad \text{all } t \geq 0; \qquad \sigma_L(0) = \sigma_{L0};$$

or more compactly,

$$\dot{\sigma}(t) = F[\sigma(t)], \qquad \text{all } t \geq 0; \qquad \sigma(0) = \sigma_0, \qquad (10b)$$

where $\sigma_0(0) = \sigma_0 = (\sigma_{K0}, \sigma_{L0})$ is the initial state of the system. We may think of the underlying extensive form game described in section 2 as being played in 'game time', which is of negligible duration relative to the 'selection time' to which (10) refers. The set of stationary points for (10) is

$$E = \{\sigma \in [0,\ 1]^2 : F(\sigma) = 0\} \qquad (11)$$

If $\sigma \in E$, then I shall call σ an *evolutionary equilibrium*.

I refer to this as a *dual* selection process because σ_K and σ_L are modified according to separate selection criteria; there is no selection process operating on strategy bundles per se. This is appropriate when strategies are chosen non-cooperatively. Casual attempts to transfer selection models from evolutionary biology to economics prove misleading in the absence of binding contracts, because Darwinian fitness provides a scalar selection criterion, rather than a vector-valued criterion as in (10).

We note certain important properties of the system in (10):

(i) The state space $[0, 1]^2$ is compact and convex. If the functions f_K and f_L are continuously differentiable, then for any initial condition $\sigma(0) = \sigma_0$ there exists a unique solution $\phi(t, \sigma_0)$ which satisfies (10), and this solution is continuous in t and σ_0 (Brauer and Nohel, 1969, Ch. 3).

(ii) A relative frequency pair σ satisfies $\sigma \in E$ if and only if σ is a Nash equilibrium of the original static game Γ.

Recall that in the dual selection model, no individual player randomizes among strategies. Hence, in every round at most four distinct organizational structures exist, corresponding to the four cells of Figure 1. The relative frequency of each of these pure-strategy bundles in round t is uniquely determined by $\sigma(t)$.

3.2. *Integrable Games*

Section 3.3. will show that for a large class of 2×2 games, stability criteria provide a refinement of Nash equilibrium (for technical references on stability issues, see Brauer and Nohel, 1969, Ch. 5; Hirsch and Smale, 1974). It is first necessary, however, to identify a Liapunov function $W\colon [0, 1]^2 \to R$, mapping relative frequency pairs σ into the real numbers, whose time derivative is positive whenever the process (10) is out of equilibrium. Formally, we require

$$\dot{W}[\sigma(t)] = \frac{\partial W[\sigma(t)]}{\partial \sigma_K} \dot{\sigma}_K(t) + \frac{\partial W[\sigma(t)]}{\partial \sigma_L} \dot{\sigma}_L(t) =$$

$$= \frac{\partial W[\sigma(t)]}{\partial \sigma_K} f_K\{\pi_K[\sigma_L(t)],\ \sigma(t)\} +$$

$$+ \frac{\partial W[\sigma(t)]}{\partial \sigma_L} f_L\{\pi_L[\sigma_K(t)],\ \sigma(t)\} > 0, \tag{12}$$

whenever $\sigma(t) \notin E$;

$\dot{W}[\sigma(t)] = 0$, whenever $\sigma(t) \in E$.

Using the profit-responsiveness assumptions (6) and (9), it is easy to see that $W(\sigma)$ satisfies this requirement if

$$\frac{\partial W(\sigma)}{\partial \sigma_K} = \pi_K(\sigma_L), \qquad \text{and} \qquad \frac{\partial W(\sigma)}{\partial \sigma_L} = \pi_L(\sigma_K). \tag{13}$$

In this case, if $\dot{\sigma}_K(t) \neq 0$, then the first term of (12) must be positive, and if $\dot{\sigma}_L(t) \neq 0$, then the second term of (12) must be positive. The time derivative $\dot{W}[\sigma(t)]$ is thus strictly positive whenever the system is out of equilibrium, as required.

A function $W(\sigma)$ satisfying (13) exists if and only if

$$\frac{\partial^2 W(\sigma)}{\partial \sigma_K \partial \sigma_L} = \frac{\partial}{\partial \sigma_L} \pi_K(\sigma_L) = \frac{\partial}{\partial \sigma_K} \pi_L(\sigma_K) = \frac{\partial^2 W(\sigma)}{\partial \sigma_L \partial \sigma_K} \tag{14}$$

Using the definitions of π_K and π_L from (4) and (7), the middle equality implies

$$u_{11} + u_{22} - u_{12} - u_{21} = v_{11} + v_{22} - v_{12} - v_{21}. \tag{15}$$

or in the more compact notation of (3), $\hat{u} = \hat{v}$. There is no reason why an arbitrarily chosen 2×2 game Γ should satisfy (15). Whenever $\hat{u} \cdot \hat{v} > 0$, however, there exists a scalar $\mu > 0$ such that $\hat{u} = \mu \hat{v}$. We can therefore replace L's expected utility function $V(\sigma)$ in section 2.1. by the rescaled function $\tilde{V}(\sigma) = \mu V(\sigma)$, so that (15) holds when $\tilde{V}(\sigma)$ is used. The preferences of player L are not altered by this transformation.

Definition: The 2×2 game Γ is *integrable* if $\hat{u} \cdot \hat{v} > 0$, or if $\hat{u} = \hat{v} = 0$. I assume throughout the remainder of this section that Γ is integrable, and that L's utility function has already been rescaled if necessary.[13] Whenever Γ is integrable, a Liapunov function with the desired property (13) is:

$$W(\sigma) = \sigma_K \sigma_L \hat{u} + \sigma_K \hat{u}_2 + \sigma_L \hat{v}_2. \tag{16}$$

The integrability condition has the following intuitive interpretation. Suppose that an impartial referee offers each player a choice between two lotteries:

(A) With equal probabilities, one of the two *diagonal* cells in the payoff matrix for Γ will be chosen; or

(B) With equal probabilities, one of the two *off-diagonal* cells will be chosen.

The preference orderings of players K and L will rank lotteries (A) and (B) in the same way if and only if Γ is integrable. This suggests that integrable games require some minimal commonality of interests between K and L. Indeed, while integrability holds for prisoners' dilemma games and games of coordination (see Appendix B), it does not hold for zero-sum games, or games of 'partiality' (Ullmann-Margalit, 1977).

3.3. *Stable Equilibria for the Dual Selection Mechanism*

Assume that Γ is integrable, and let W be defined by (16). Then consider any evolutionary equilibrium $\sigma^* \in E$. The following results are stated without proof (see Appendix A for formal definitions of the stability concepts).

(a) $\sigma^* \in E$ if and only if σ^* satisfies the Kuhn-Tucker (KT) necessary conditions for a local maximum of $W(\sigma)$.
(b) $\sigma^* \in E$ is *asymptotically stable* if and only if σ^* is a *strict local maximum* of W.
(c) $\sigma^* \in E$ is *stable* if and only if σ^* is a *local maximum* of W.
(d) $\sigma^* \in E$ is *unstable* if and only if σ^* is *not* a local maximum of W.

The maximizing character of a stable equilibrium is evident from the fact that $W[\sigma(t)]$ is strictly increasing out of equilibrium; for

---

[13] The method used in the text can be extended to all 2×2 games along the following lines. If $\hat{u} \cdot \hat{v} < 0$, define $W(\sigma)$ so that its partial derivatives with respect to σ_K and σ_L are $\pi_K(\sigma_L)$ and $-\pi_L(\sigma_K)$ respectively, rescaling one of the utility functions if necessary. Nash equilibria will still correspond to evolutionary equilibria, but *stable* equilibria will now be characterized by saddle points of W rather than maxima as in the text. The details of this generalization are unimportant for present purposes and are omitted.

methods of formal proof, see Dow (1986b). We observe that a stable evolutionary equilibrium exists for every integrable game, since W must reach a maximum at some $\sigma^* \in [0, 1]^2$.

While each Nash equilibrium of Γ corresponds to a selection equilibrium, the set of *stable* evolutionary equilibria may be smaller than the set of Nash equilibria for the static game Γ. Each Nash equilibrium point satisfies the KT conditions for a local maximum of W, but the KT conditions are only necessary, not sufficient, for a maximum. Any Nash equilibrium *not* associated with a local maximum of W is evolutionarily unstable. An example is given in Appendix B of an unstable Nash equilibrium which corresponds to a saddle point of W rather than a local maximum.[14] It can be shown that if the number of Nash equilibria for Γ is finite, then each corresponds to a selection equilibrium which is either asymptotically stable (if a local maximum of W) or unstable (otherwise).

The examples in Appendix B show that evolutionarily stable equilibria are often *Pareto inefficient* Nash equilibria in the original static game Γ. If Γ has a unique Nash equilibrium, for instance, this point will be the unique evolutionary equilibrium regardless of its efficiency properties: the prisoners' dilemma provides a stark illustration. Appendix B also shows that the game Γ may have both efficient and inefficient stable equilibria, rendering the efficiency of the evolutionary outcome dependent upon history. Selection forces provide no automatic escape from problems of contractual incompleteness, and a world of private ordering may well be littered with evolutionary market failures.

4. LEGAL CENTRALISM AND JOINT SELECTION

We now assume, contrary to the assumption of section 3, that capitalists and workers can sign binding contracts which will govern their strategy choices in the game Γ. In this world of legal centralism, the selection process will operate at the level of strategy bundles or contracts, rather than upon the individual elements (strategies) of these bundles. I call this a process of joint selection, and show below that such a process yields Pareto efficient evolutionary equilibria.

4.1. Profit-Responsive Joint Selection Mechanisms

Call the pure-strategy bundle (α_i, β_j) a *contract* of type (ij), where i, j = 1, 2. Define the *weighted average payoff* of contract type (ij) by

$$\rho_{ij} = \lambda u_{ij} + (1 - \lambda) v_{ij}, \qquad \text{with } 0 < \lambda < 1. \tag{17}$$

---

[14] It can be shown that no integrable game has a stable mixed-strategy equilibrium; Crawford (1985) reached a similar conclusion in a related dynamic model. Since every integrable game has an evolutionarily stable equilibrium, it follows that every such game has a stable pure-strategy equilibrium. This result also implies that two distinct organizational structures cannot co-exist in a stable dual selection equilibrium.

I interpret λ as a fixed parameter, independent of contract type, describing the relative bargaining power of the two classes. The selection process operates as follows.

1. As before, capitalists and workers are drawn from large, internally homogeneous populations.
2. At the start of round $t \in T = \{0, 1, \ldots\}$, each worker is randomly paired with some capitalist, and each (K, L) pair negotiates a contract to govern production activities in round t. The proportion of (K, L) pairs adopting contract type (ij) in round t is $\theta_{ij}(t)$, where $i, j = 1, 2$. Each pair then plays the game Γ using the pure-strategy bundle (α_i, β_j) specified by its contract.
3. After round t, all capitalists and workers observe the average payoffs u_{ij} and v_{ij} obtained by the members of each class from participation in the four contract types. (This is a more stringent informational requirement than in section 3, where players of type K only observed the average payoffs of the strategies α_1 and α_2, and players of type L observed the average payoffs of the strategies β_1 and β_2.).
4. The frequency θ_{ij} of contract type (ij) expands in round $t+1$ relative to the alternative contract (kl) if $\rho_{ij} > \rho_{kl}$, and declines relative to (kl) if $\rho_{ij} < \rho_{kl}$.
5. Steps 2—4 repeat in round $t+1$.

At step 3, we again exploit the large numbers assumption in order to equate the expected payoffs in Figure 1 with ex post averages. The parameter λ used to compute the ρ_{ij} in step 4 weights the payoffs of the two classes in evaluating past contracts, and can be regarded as a measure of relative bargaining strength in the negotiations leading up to contract choices in each round.

We now formalize this selection process. Let the state of the system be given by

$$\theta(t) \equiv [\theta_{11}(t), \theta_{12}(t), \theta_{21}(t), \theta_{22}(t)] \in [0, 1]^4$$

where $\theta(t)$ is the vector of relative frequencies for the four contract types. The *joint selection mechanism* is defined by

$$\theta_{ij}(t+1) - \theta_{ij}(t) = g[\rho_{ij}, \theta(t)], \qquad i, j = 1, 2; \tag{18}$$

$$\sum_{i=1}^{2} \sum_{j=1}^{2} g[\rho_{ij}, \theta(t)] = 0.$$

By contrast with the variables π_K and π_L in section 3, which were functions of σ, fixed parameters ρ_{ij} do not depend upon the state variable θ. I shall ignore contract-specific aspects of the selection process other than ρ_{ij} for notational simplicity.

Now define the *maximizing* set of contract types:

$$\Omega = \{(ij) : i,\ j \in \{1,\ 2\},\quad \rho_{ij} = \max_{(kl)} (\rho_{kl})\} \tag{19}$$

This is the set of optimal contracts according to criterion (17). We now define profit-responsiveness for mechanism (18):

$$g\,[\rho_{ij},\ \theta(t)] > g\,[\rho_{kl},\ \theta(t)], \tag{20}$$

$$\text{whenever } \rho_{ij} > \rho_{kl} \quad \text{and} \quad \sum_{(ij)\in\Omega} \theta_{ij}(t) < 1;$$

$$g\,[\rho_{ij},\ \theta(t)] = 0, \quad \text{all (ij)}, \quad \text{whenever} \quad \sum_{(ij)\in\Omega} \theta_{ij}(t) = 1.$$

The first property states that if some suboptimal contract is in use, inferior contracts will be displaced over time by contracts with higher weighted average payoffs. The second property states that when all contracts in use belong to the maximizing set Ω, no further changes in relative frequencies will occur.

Moving to continuous time, we obtain the system

$$\dot{\theta}_{ij}(t) = g\,[\rho_{ij},\ \theta(t)], \quad \text{all } t \geq 0; \quad \theta_{ij}(0) = \theta_{ijo}; \quad i,\ j = 1,\ 2; \tag{21a}$$

or more compactly,

$$\dot{\theta}(t) = G\,[\theta(t)], \quad \text{all } t \geq 0; \quad \theta(0) = \theta_o. \tag{21b}$$

The set of joint-selection equilibria will be denoted by

$$J = \{\theta \in [0,\ 1]^4 : G(\theta) = 0\}. \tag{22}$$

4.2. *Contractual Evolution with Joint Selection*

To allow for the feasibility of correlated strategies in a regime of legal centralism, we use a straightforward extension of the expected utility functions defined in (1):

$$U(\theta) = \sum_{i=1}^{2} \sum_{j=1}^{2} \theta_{ij}\, u_{ij}$$

$$V(\theta) = \sum_{i=1}^{2} \sum_{j=1}^{2} \theta_{ij}\, v_{ij} \tag{23}$$

Now define a Liapunov function for the joint selection process by

$$Z(\theta) \equiv \lambda U(\theta) + (1-\lambda) V(\theta) \tag{24}$$

$$= \sum_{i=1}^{2} \sum_{j=1}^{2} \theta_{ij}\, \rho_{ij}$$

The profit-responsiveness assumptions in (20) imply

$$\dot{Z}[\theta(t)] = \sum_{i=1}^{2} \sum_{j=1}^{2} \frac{\partial Z[\theta(t)]}{\partial \theta_{ij}} \dot{\theta}_{ij}(t) \tag{25}$$

$$= \sum_{i=1}^{2} \sum_{j=1}^{2} \rho_{ij}\, g[\rho_{ij}, \theta(t)] > 0,$$

whenever $\theta(t) \notin J$; and

$\dot{Z}[\theta(t)] = 0,$ whenever $\theta(t) \in J$.

We note that $\theta^* \in$ J if and only if θ^* maximizes Z. This implies that (in contrast to the dual selection model in section 3) every evolutionary equilibrium of the joint selection process is stable, because every such point maximizes the Liapunov function Z. Some stable evolutionary equilibrium must exist, because Z (θ) must reach a maximum somewhere on the set $[0, 1]^4$.

4.3. Joint Selection and Pareto Efficiency

Jointly randomized strategies of the form $\theta \in [0, 1]^4$ yield a set φ of *feasible* payoff vectors (u, v) for the game Γ in Figure 1 which is the convex hull of the pure-strategy payoffs.

$$\varphi \equiv \{(u, v) \in R^2 : \text{for some } \theta \in [0, 1]^4, \tag{24}$$

$$u = \sum_{i=1}^{2} \sum_{j=1}^{2} \theta_{ij} u_{ij} \quad \text{and}$$

$$v = \sum_{i=1}^{2} \sum_{j=1}^{2} \theta_{ij} v_{ij} \}$$

The *Pareto efficient* subset φ^* of φ is the set of undominated payoff vectors in φ.

$$\varphi^* \equiv \{(u, v) \in \varphi : \text{if } (u', v') \geq (u, v) \tag{27}$$

$$\text{and } (u', v') \neq (u, v), \text{ then } (u', v') \notin \varphi\}$$

Figure 2 illustrates these concepts.

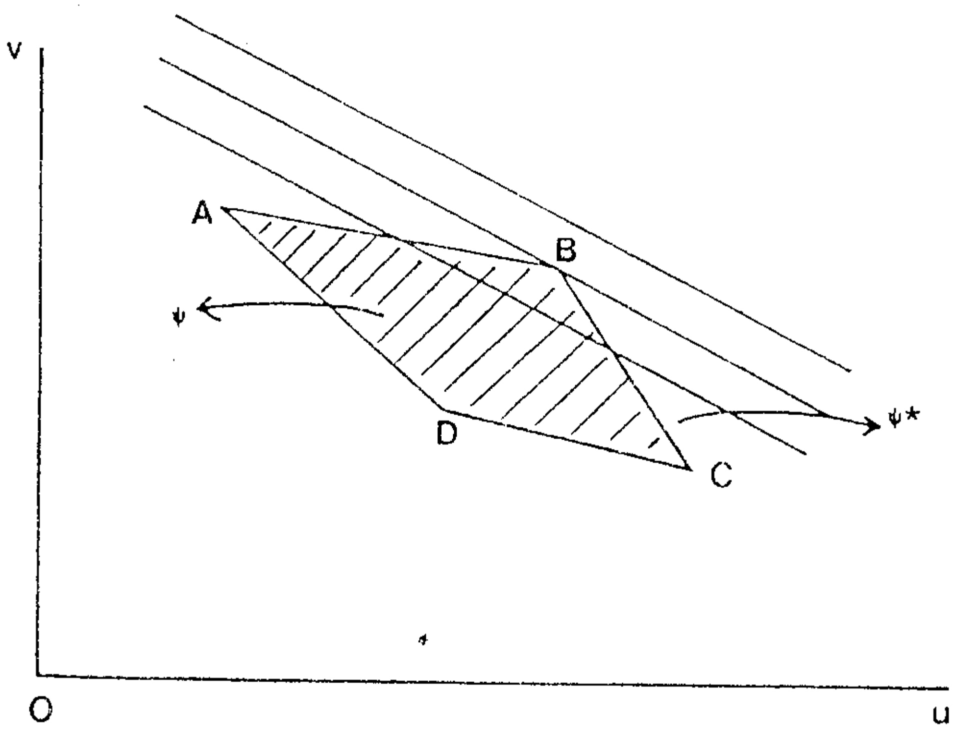

FIGURE 2

Joint Selection and Efficient Contracts

Points A, B, C, and D represent arbitrarily chosen values for the four pure-strategy payoff vectors from Figure 1. The Pareto frontier φ^* is the set of points in (u, v) space lying on the extreme northeasterly boundary of φ—that is, the union of the line segments AB and BC. More generally, φ^* could consist of one, two or three such segments.

The vector $\theta^* \in [0, 1]^4$ is a joint-selection equilibrium if and only if it achieves a maximum of the Liapunov function $Z(\theta)$. When $\lambda \in (0, 1)$, so that each player's payoff has strictly positive weight in the evaluation of alternative contracts, then a maximum of Z can occur only on the Pareto frontier; that is, $\theta^* \in J$ only if $[U(\theta^*), V(\theta^*)] \in \varphi^*$. This is illustrated in Figure 2, where the parallel lines indicate level sets of the expression $\lambda u + (1 - \lambda) v$, and point B is the unique equilibrium point.

The relative bargaining power of the two classes in the contract negotiation process will determine which of the various points on the Pareto frontier is actually reached.[15] The common slope of the level

[15] A point maximizing the linear expression Z is not to be confused with the Nash bargaining solution for cooperative games (for discussions of the latter, see Dow, 1985; 1988b; Binmore and Dasgupta, 1987). The joint selection model has no counterpart to the threat point concept used in the Nash bargaining solution. I simply assume that each capitalist and worker pair will somehow agree to adopt one of the four contract types in each round, and model the aggregate result of these choices at the populational level by (18) or (21).

sets in Figure 2 is $-\lambda/(1-\lambda) \in (-\infty, 0)$. Any point on the Pareto frontier φ^* can be a maximum of $\lambda u + (1-\lambda) v$, and thus a stable joint selection equilibrium, for some value of the parameter λ. Only pure-strategy payoff vectors on the frontier can be sustained as *asymptotically* stable equilibria, since only at these points is it possible for the maximization of $\lambda u + (1-\lambda) v$ to be strict.

A point in the interior of a line segment on the Pareto frontier, implying the co-existence of two distinct contract types, is sustainable as an evolutionary equilibrium only at a unique value of λ (that value which equates $-\lambda/(1-\lambda)$ with the slope of the segment). Each pure-strategy payoff vector on the Pareto frontier is associated with the universal adoption of some corresponding contract type and is sustainable as an equilibrium for a range of λ values. Thus, only pure-strategy payoff vectors will persist as stable equilibria for small perturbations in λ.

5. MARX REVISITS DARWIN: THE FIRM AS A MARKET FAILURE

A century ago, Menger posed what he regarded as the central question for the social sciences:

> How can it be that institutions which serve the common welfare and are extremely significant for its development come into being without a common will directed toward establishing them? (1985: 146, emphasis deleted)

In Menger's view, a satisfactory answer to this question must involve an invisible-hand explanation of the type described in section 1.2. It is intriguing that non-conspiratorial Marxian theories encounter a parallel intellectual problem. To see this, we can »stand Menger on his head« by reversing the value premises of the preceding quote:

> How can it be that institutions which are injurious to the common welfare persist in the face of opposition, despite the absence of a class-wide will directed toward maintaining them?

Once the socially beneficent overtones of Adam Smith's metaphor are discarded, Marxian theorists may find that evolutionary invisible-hand models can provide satisfying answers to questions of this latter sort. Indeed, one could argue that theories of class conflict *must* supply an invisible-hand explanation of existing institutions, since the struggle between contending classes excludes any explanation where these institutions are the intended result of some cooperative social contract.

Few orthodox economists have been willing to accept the Marxian contention that capitalist firms necessarily adopt inefficient organizational forms (Edwards, 1979; Marglin, 1982; Bowles, 1985). Mainstream skepticism derives in large part from faith in the efficacy of market selection processes: how could inefficient organizations survive in competition with efficient ones?[16] There is irony in this, because Marx offerred to dedicate volume 2 of *Das Kapital* to Charles Darwin (Gould, 1977: 26—27).

Dual selection mechanisms provide Marxian theory with an invisible-hand explanation of organizational form which does not compel assent to mainstream efficiency claims. This view of the market selection process is also congenial to the Marxian perspective in other ways. As noted in section 2.3, there is no need to assume that social institutions are transparent to their participants. One could thus view workers and capitalists as engaging in trial-and-error experimentation with new strategies, and retaining those strategies that are found to advance their respective class interests. Indeed, an analysis of the forces acting to stabilize prevailing institutions may itself suggest novel strategies for the transformation of these institutions.

The selection models sketched here also shed new light on the relation between transaction cost economics and the Marxian framework.[17] The joint selection model partially vindicates the efficiency premise of TCE, but only to the extent that factor suppliers can contract on all significant decisions to be made during the production process. If the players cannot easily negotiate enforceable contracts prior to each round, then we are back to a world of private ordering and dual selection.

Dual selection constitutes a form of market failure, where neither factor supplier fully appropriates the social costs and benefits of a given governance structure (Dow, 1987a). In a world of incomplete contracts, evolutionarily stable governance structures will reflect the failure of each agent to internalize all relevant consequences of his or her organizational behavior. Of course, sometimes a Pareto efficient strategy bundle will happen fortuitously to be a Nash equilibrium of the underlying non-cooperative game played among factor suppliers, but this cannot be expected in general (Dow, 1987b). To put the point a bit more colorfully: existing economic institutions need not *solve* some prisoners' dilemma (Ullmann-Margalit, 1977; Schotter, 1981); they may simply *be* prisoners' dilemmas from which the prisoners have failed to escape. The distinction matters, because the second view implies much more scope for deliberate policy interventions aimed at altering background institutions or contracting possibilities in ways which make escape possible.

Real production processes involve both contractual and strategic elements. One might say that for TCE the contractual glass is half

[16] The converse of the mainstream tendency to impute efficiency to existing forms is a tendency to impute inefficiency to proposed alternatives. For instance, various writers (Nozick, 1974, Ch. 8; Jensen and Meckling, 1979; Williamson, 1985: 322—324) have inferred the inefficiency of the labor-managed firm from its failure to emerge and remain viable in competition with capitalist rivals. Proponents of labor management have struggled to rebut this inference (Putterman, 1984; Dow, 1986a; 1987a), but implicit appeals to processes of joint selection are undeniably persuasive to most economists.

[17] For further explorations into the relationship between the transaction cost and Marxian approaches, see Goldberg (1980) and Dow (1987a).

full, while for Marxian economists it is half empty. TCE acknowledges opportunistic behavior, but nonetheless models organizational forms as contractual agreements. Marxians recognize the relevance of legal institutions, but interpret the firm as pre-eminently an arena for strategic behavior and class conflict. The contracting perspective of TCE is hospitable to the efficiency implications of joint selection, while the conflict perspective of Marxian economics is hospitable to the probable inefficiency of dual selection.

As I have stressed repeatedly, the relative importance of the joint and dual selection mechanisms depends, among other things, on the availability of background institutions which facilitate contracting. It is therefore tempting to seek an evolutionary synthesis: the provisions of formal contracts are jointly selected, while contractual gaps are filled in by non-cooperative processes of dual selection. The value of this formulation is unclear, but some misunderstandings may at least be averted by recognizing that transaction cost and Marxian economists implicitly operate with different selection models, motivated by divergent assumptions about the contractual framework of capitalist production.

APPENDIX A
STABILITY CONCEPTS

Stability. The equilibrium point σ^* is *stable* if for every number $\epsilon > 0$, there is a $\delta > 0$ such that when $|\sigma_o - \sigma^*| < \delta$, the solution $\phi(t, \sigma_o)$ exists for all $t \geq 0$, and $|\phi(t, \sigma_o) - \sigma^*| < \epsilon$ for all $t \geq 0$.

Asymptotic Stability. The equilibrium point σ^* is *asymptotically stable* if it is stable, and if there exists a number $\delta_o > 0$ such that whenever $|\sigma_o - \sigma^*| < \delta_o$, then $\lim_{t \to \infty} |\phi(t, \sigma_o) - \sigma^*| = 0$.

Instability. The equilibrium point σ^* is *unstable* if it is not stable.

APPENDIX B
EXAMPLES OF DUAL-SELECTION EQUILIBRIA

Figures 3, 4, and 5 provide examples of integrable games. The evolutionary equilibria of each game correspond to static Nash equilibria, as explained in section 3.3. Nash equilibria need not be Pareto efficient; the same is true of stable evolutionary equilibria, as is demonstrated by the examples in Figures 3 and 4. A game can have both efficient and inefficient stable equilibria, as shown in Figure 5. The same example illustrates an interesting general result: for all integrable 2×2 games, mixed-strategy equilibria must be evolutionarily unstable.

Figure 3: The Symmetric Prisoners' Dilemma

Symmetric games are always integrable. For the specific game shown in Figure 3a, we have $\hat{u} = \hat{v} = 0$, $\hat{u}_2 = \hat{v}_2 = -1$. Hence, $W(\sigma) = -(\sigma_K + \sigma_L)$. This function achieves a strict global maximum at $\sigma^* = (0, 0)$, which is the standard (Pareto inefficient) prisoners' dilemma outcome. This evolutionary equilibrium is unique and asymptotically stable, as indicated by the phase diagram in Figure 3b, which shows the direction of movement for σ_K and σ_L when the system is out

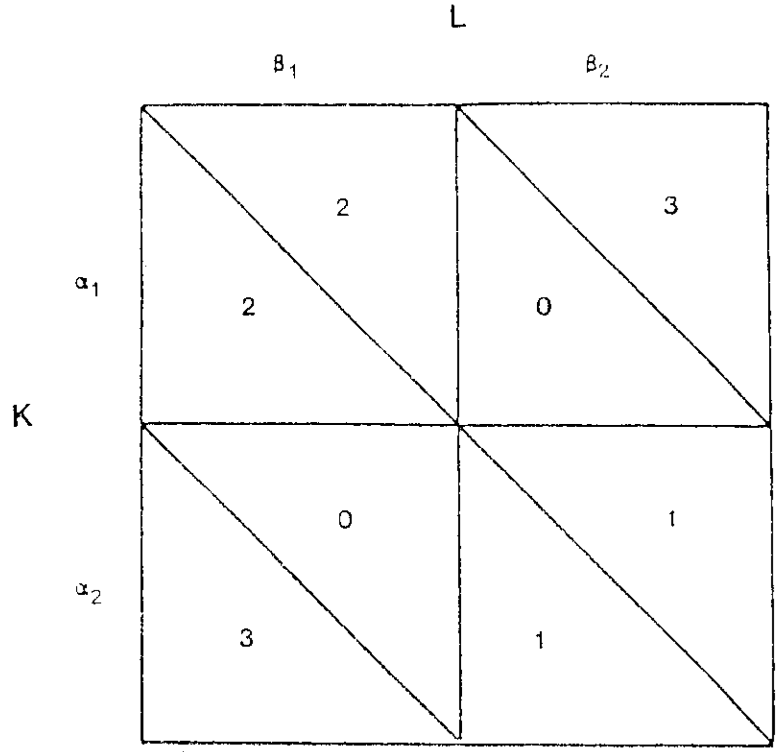

FIGURE 3a
The Symmetric Prisoners' Dilemma

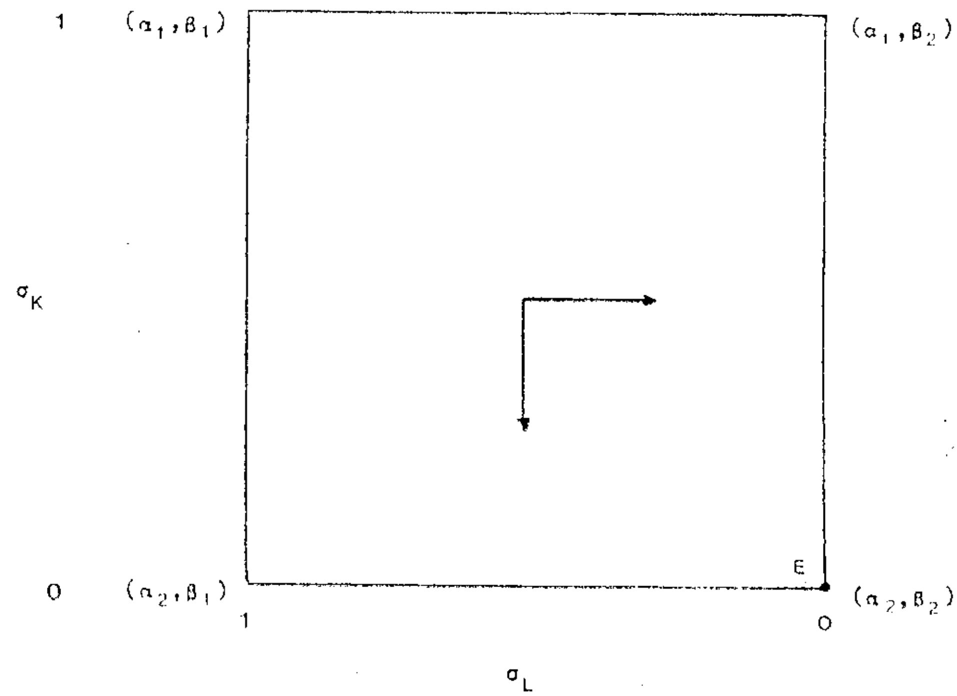

FIGURE 3b
Phase Diagram for the Symmetric Prisoners' Dilemma

of equilibrium. The axes of the phase diagram are defined so that the corners of Figure 3b correspond to the pure-strategy cells in Figure 3a. The equilibrium point (α_2, β_2) is indicated by E in Figure 3b.

Figure 4: The Asymmetric Prisoners' Dilemma

Example 4a is given primarily to emphasize that symmetry is not essential for integrability. In this game, L has the same incentive to defect from the cooperative outcome (α_1, β_1) as in Figure 3a, but K will not defect. Nonetheless, the unique Nash equilibrium is still the inefficient outcome (α_2, β_2).

For this game, $\hat{u} = \hat{v} = 2$, $\hat{u}_2 = -1$, and $\hat{v}_2 = -3$, so that $W(\sigma) = 2\sigma_K\sigma_L - \sigma_K - 3\sigma_L$. As in Figure 3, this function achieves a strict global maximum at $\sigma^* = (0, 0)$. Notice that

$$\frac{\partial W(\sigma)}{\partial \sigma_K} = \pi_K(\sigma_L) = 2\sigma_L - 1 \gtreqless 0, \quad \text{for} \quad \sigma_L \gtreqless 1/2.$$

This change in the gradient of $W(\sigma)$ is indicated by the vertical line at $\sigma_L = 1/2$ in the phase diagram Figure 4b. Since $\dot{\sigma}_L < 0$ for all σ_K, no point on this vertical line can be an evolutionary equilibrium. The outcome (α_2, β_2), indicated by point E, is asymptotically stable as before.

Figure 5: A Coordination Game

It is evident from Figure 5a that (α_1, β_1) and (α_2, β_2) are both pure-strategy Nash equilibria. The former is Pareto efficient while the latter is not. There is also a mixed-strategy Nash equilibrium $\sigma = (1/6, 1/6)$, whose payoff vector $(5/6, 5/6)$ is Pareto dominated by both pure-strategy equilibria. For this game, $\hat{u} = \hat{v} = 6$ and $\hat{u}_2 = \hat{v}_2 = -1$, so that $W(\sigma) = 6\sigma_K\sigma_L - \sigma_K - \sigma_L$. Both pure-strategy equilibria are strict local maxima of W and are therefore (locally) asymptotically stable.

We have

$$\frac{\partial W(\sigma)}{\partial \sigma_K} = \pi_K(\sigma_L) = 6\sigma_L - 1 \gtreqless 0 \qquad \text{for } \sigma_L \gtreqless 1/6;$$

$$\frac{\partial W(\sigma)}{\partial \sigma_L} = \pi_L(\sigma_K) = 6\sigma_K - 1 \gtreqless 0 \qquad \text{for } \sigma_K \gtreqless 1/6.$$

This sign change in the gradient of W is indicated in Figure 5b by the horizontal and vertical lines at $\sigma_K = 1/6$ and $\sigma_L = 1/6$, respectively. The intersection of these lines yields the mixed equilibrium

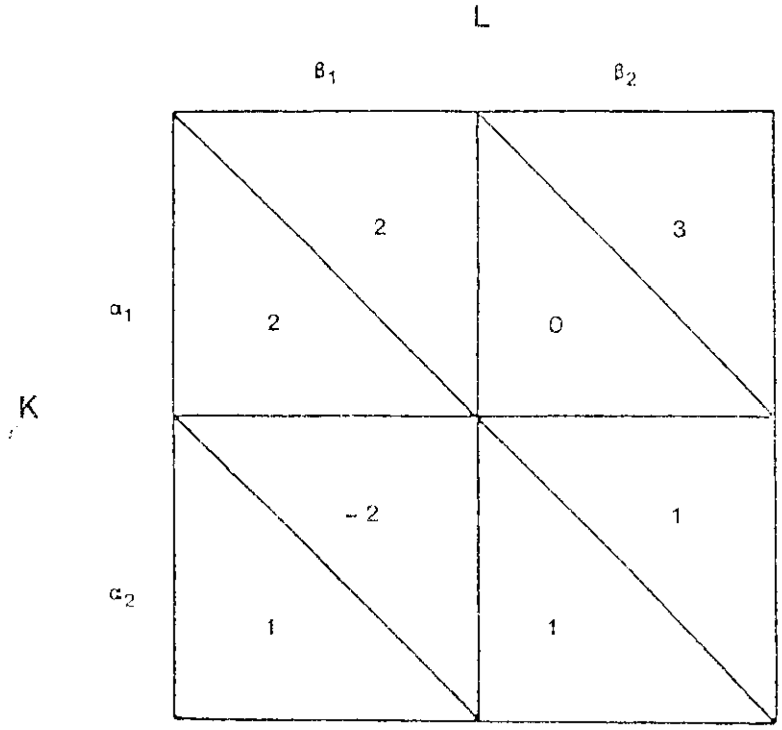

FIGURE 4a
The Asymmetric Prisoners' Dilemma

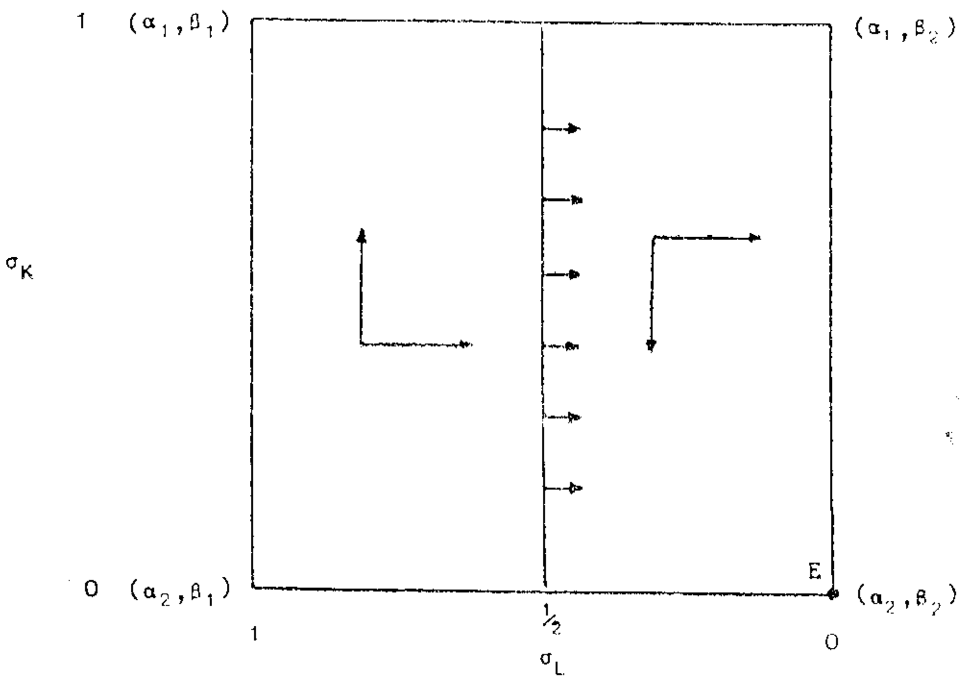

FIGURE 4b
Phase Diagram for the Asymmetric Prisoners' Dilemma

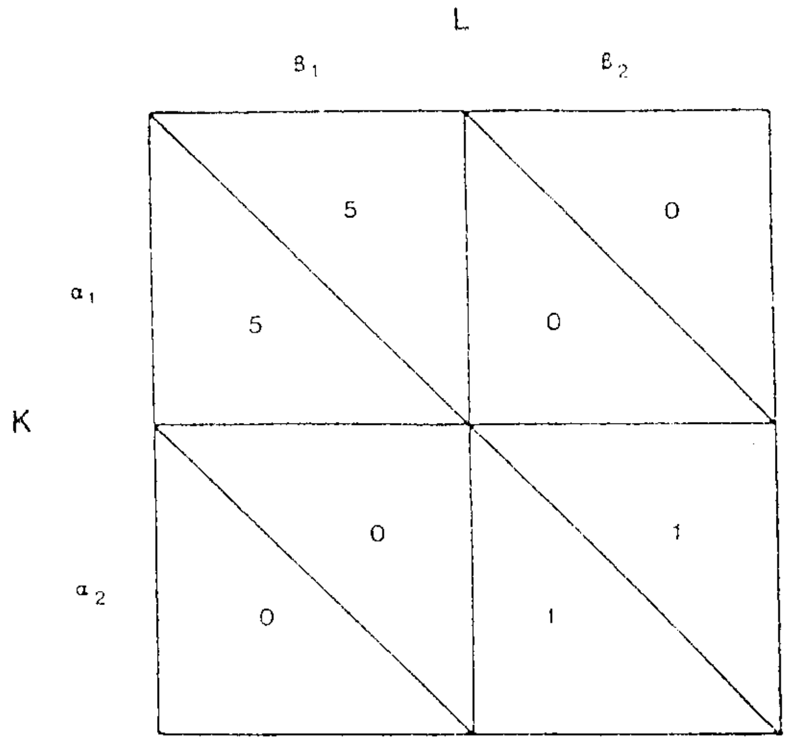

FIGURE 5a
The Coordination Game

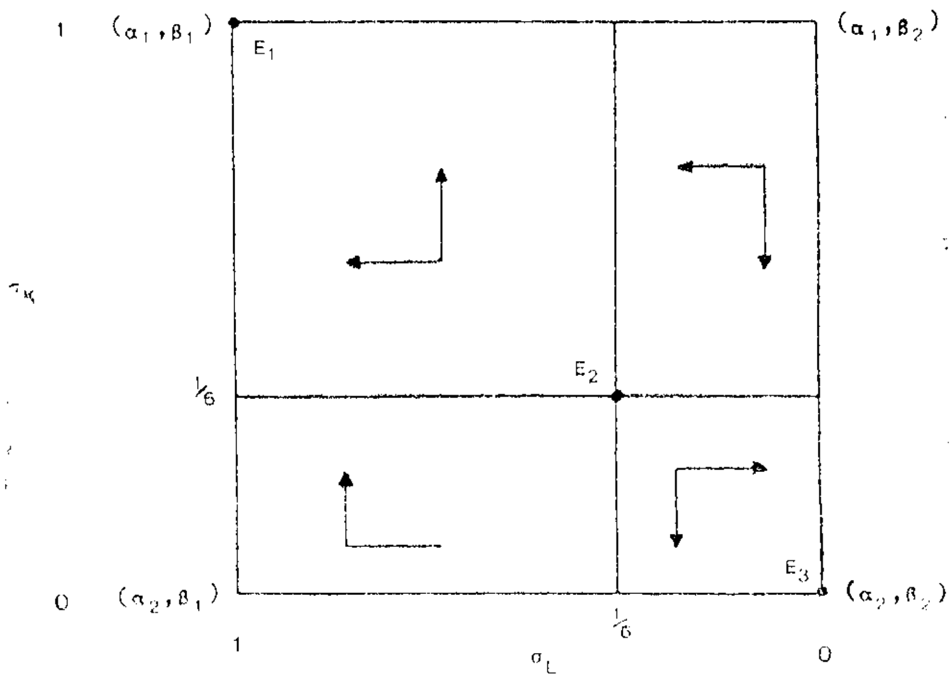

FIGURE 5b
Phase Diagram for the Coordination Game

(1/6, 1/6). This equilibrium is evolutionarily unstable, as indicated by the arrows of the phase diagram. This instability reflects the fact that the equilibrium ((1/6, 1/6) is a saddle-point of W rather than a local maximum. Indeed, *any* mixed-strategy equilibrium of an integrable game must be unstable for this reason. (It is also possible to construct examples of evolutionarily unstable pure-strategy equilibria.) The specific stable equilibrium toward which the system gravitates depends upon the initial relative frequency pair σ_0; that is, upon the history of the system.

Received: 16.05. 1988.
Revised: 06. 06. 1988.

REFERENCES

Axelrod, R., 1984, The Evolution of Cooperation (Basic Books, New York).

Binmore, K. and P. Dasgupta, eds., 1986, Economic Organizations as Games (Basil Blackwell, New York).

Binmore, K. and P. Dasgupta, eds.. 1987, The Economics of Bargaining (Basil Blackwell, New York).

Bowles, S., 1985, The production process in a competitive economy, American Economic Review 75: 16—36.

Brauer, F. and J. A. Nohel, 1969, Qualitative Theory of Ordinary Differential Equations (W. A. Benjamin, Inc., New York).

Crawford, V, 1985, Learning behavior and mixed-strategy Nash equilibria, Journal of Economic Behavior and Organization 6: 69—78.

Cross, J., 1983, A Theory of Adaptive Economic Behavior (Cambridge University Press, New York).

Dow, G., 1985, Internal bargaining and strategic innovation in the theory of the firm, Journal of Economic Behavior and Organization 6, 301—320.

Dow, G., 1986a, Control rights, competitive markets, and the labor management debate, Journal of Comparative Economics 10: 48—61.

Dow, G., 1986b, Stability analysis for profit-responsive selection mechanisms, Mathematical Social Sciences 12: 169—183.

Dow, G., 1987a, The function of authority in transaction cost economics, Journal of Economic Behavior and Organization 8: 13—38.

Dow, G., 1987b, Self-enforcing authority structures: A stochastic game approach, Working paper 87—11, Department of Economics, University of Alberta.

Dow, G., 1988a, Configurational and coactivational views of organizational structure, Academy of Management Review 13: 53—64.

Dow, G., 1988b, Information, production decisions, and intra-firm bargaining, International Economic Review 29: 57—79.

Edwards, R., 1979, Contested Terrain (Basic Books, New York).

Elster, J., 1983, Explaining Technical Change (Cambridge University Press, New York).

Friedman, J. and R. Rosenthal, 1986, A positive approach to noncooperative games, Journal of Economic Behavior and Organization 7: 235—251.

Friedman, M., 1953, Essays in Positive Economics (University of Chicago Press, Chicago, IL), Ch. 1.

Fudenborg, D. and E. Maskin, 1986, The folk theorem in repeated games with discounting or with incomplete information, Econometrica 54: 533—554.

Goldberg, V., 1980, Bridges over contested terrain: Exploring the radical account of the employment relationship, Journal of Economic Behavior and Organization 1: 249—274.

Gould, S., 1977, Ever Since Darwin (Norton, New York).

Grossman, S. and O. Hart, 1986, The costs and benefits of ownership: A theory of vertical and lateral integration, Journal of Political Economy 94: 691—719.

Hirsch, M. and S. Smale, 1974, Differential Equations, Dynamical Systems, and Linear Algebra (Academic Press, New York).

Jensen, M., 1983, Organization theory and methodology, The Accounting Review 58: 319—339.

Jensen, M. and W. Meckling, 1979, Rights and production functions: An application to labor-managed firms and codetermination, Journal of Business 52: 469—506.

Kreps, D., 1984, Corporate culture and economic theory, mimeo, Stanford University.

Langlois, R., 1986, Rationality, institutions, and explanation, Ch. 10 in R. Langlois, ed., Economics as a Process: Essays in the New Institutional Economics (Cambridge University Press, New York): 225—255.

Marglin, S., 1982, Knowledge and power, in F. Stephen, ed., Firms, Organization and Labour (St. Martin's Press, New York): 146—164.

Masten, S., 1986, The institutional basis for the firm, mimeo, University of Michigan.

Maynard Smith, J., 1982, Evolution and the Theory of Games (Cambridge University Press, New York).

McKelvey, B. and H. Aldrich, 1983, Populations, natural selection, and applied organizational science, Administrative Science Quarterly 28: 101—128.

Menger, C., 1985, Investigations into the Method of the Social Sciences with Special Reference to Economics, ed. L. Schneider, trans. F. Nock (New York University Press, New York).

Mirowski, P., 1986, Institutions as a solution concept in a game theory context, in P. Mirowski, ed., The Reconstruction of Economic Theory (Kluwer-Nijhoff, Boston): 241—263.

Nelson, R. and S. Winter, 1982, An Evolutionary Theory of Economic Change (Belknap, Cambridge, MA).

Nozick, R., 1974, Anarchy, State, and Utopia (Basic Books, New York).

Owen, G., 1982, Game Theory, 2nd ed. (Academic Press, New York).

Pfeffer, J., 1982, Organizations and Organization Theory (Pitman, Marshfield, MA).

Putterman, L., 1984, On some recent explanations of why capital hires labor, Economic Inquiry 22: 171—187.

Schotter, A., 1981, The Economic Theory of Social Institutions (Cambridge University Press, New York).

Schotter, A., 1986, The evolution of rules, Ch. 5 in R. Langlois, ed., Economics as a Process: Essays in the New Institutional Economics (Cambridge University Press, New York): 117—133.

Scott, W. R., 1981, Organizations: Rational, Natural, and Open Systems (Prentice-Hall, Englewood Cliffs, NJ).

Stiglitz, J., 1975, Incentives, risk, and information: Notes towards a theory of hierarchy, Bell Journal of Economics 6: 552—579.

Ullmann-Margalit, E., 1977, The Emergence of Norms (Oxford University Press, New York).

Ullmann-Margalit, E., 1978, Invisible hand explanations, Synthese 39: 263—291.

Van Parijs, P., 1981, Evolutionary Explanation in the Social Sciences (Rowman and Littlefield, Totowa, NJ).

Williamson, O., 1975, Markets and Hierarchies (Free Press, New York).

Williamson, O., 1985, The Economic Institutions of Capitalism (Free Press, New York).

Williamson, O., 1986, The economics of governance: Framework and implications, Ch. 8 in R. Langlois, ed., Economics as a Process: Essays in the New Institutional Economics (Cambridge University Press, New York): 171—202.