Entrepreneurs: Self-Confident Agents Busting Self-Confirming Equilibria

Michael D. Ryall

SFI WORKING PAPER: 1996-08-068

Entrepreneurs: self-confident agents busting self-confirming equilibria

Michael D. Ryall[1]

Department of Economics
University of California, Los Angeles

January 29, 1996

[1]E-mail: mryall@ucla.edu

1 Hero with a Thousand Faces

The purpose of this project is to develop a model to compare and contrast several competing economic views of the entrepreneur. At issue are the dimensions along which agents bearing the moniker *entrepreneur* differ from the rest of society and the results to which these differences lead. Based upon recent work in game theory, a natural definition of the entrepreneur arises: those individuals who move the rest of society away from self-confirming behaviors rooted in incorrect beliefs. We wish to explore the conditions under which such change can be effected and contrast them with traditional economic views of entrepreneurial activity. Agents in this model interact strategically while attempting to learn the true structure of their environment from observational data.

Since the advent of the industrial revolution, we have often held the entrepreneur before us as an exemplar of social virtue. In recent times, our image of the entrepreneur remains positive even in the face of the fashionable contempt often heaped upon the business community in general. From a cultural point of view, it is interesting that the image of the entrepreneur takes on such a wide range of forms depending upon the context of our invocation: from underdog economic paladin gallantly battling the corporate Goliath, to street urchin made prince by way of high character and hard work; from shrewd financial impresario to wacky inventive genius; from speculative idea merchant to resolute commercial visionary. In light of this diversity of popular meaning, it is not too surprising that Economics, though explicitly tasked with bringing precision to such notions, has itself attached a number of concepts to the word *entrepreneur.* These differences are not trivial, particularly when they represent mutually exclusive, competing views of the entrepreneur's economic function.

Economists, from the earliest neoclassical thinkers, certainly agree that entrepreneurs are crucial elements of a smoothly functioning capitalist system. The issue comes in describing just exactly *what* it is they do that elevates them to this status. Jevons [1871] (who, along with Menger and Walras, co-founded the neoclassical revolution) explicitly identified a special niche for the entrepreneur:

> The fact that workmen are not their own capitalists introduces complexity into our problem. The capitalists, or entrepreneurs, enter as a distinct interest. It is they who manage a branch of production, and form estimates as to the expected produce.

This early definition of the entrepreneur is more akin to our modern notion of a business manager than it is to that of an innovator.[1] It was not until the earlier part of this century that economists seriously began to view entrepreneurs as a special subset of managerial

[1] This attitude appears to be shared by Jevons' co-revolutionaries. Menger [1871], also the originator of the Austrian school, explicitly dismisses the suggestion of a contemporary writer (Mangoldt) that risk-bearing is the essential function of entrepreneurship.

agents. Clark [1918] differentiates between inventors, or "physical discoverers," and innovators, or "aspiring adopters," of new technologies. Following him, Knight [1933] explores the role of specialized risk-bearers in the coordination of economic activity. Joseph Schumpeter [1942], bestows upon entrepreneurs the prime, almost mystical, role of agents-of-change in his process of "creative destruction." In more recent frameworks, for example Fudenberg and Levine [1995], we envision entrepreneurs as those who experiment against the tide of prevailing activity, either as random "mutants" in an evolutionary learning model, or as more thoughtful decision makers solving a difficult bandit problem.

This literature can be grouped according to which of three essential driving forces of entrepreneurial behavior is being espoused: 1) risk taking (the Knightian view); 2) vision (the Austrian view); and, 3) experimentation (the game theoretic view).[2] Each of these motivations should correlate with specific kinds of observable behavior. In order to explore these issues, we can imagine a model in which agents may differ along each of three dimensions: 1) risk preferences; 2) world-view; and 3) uncertainty with respect to world-view. Although one might imagine several of these driving forces at work within a single individual, their effects can be isolated by attributing entrepreneurship to one of three corresponding types: 1) Risk-Takers: individuals with lower levels of risk aversion relative to the population at large; 2) Visionaries: agents whose world-view differs significantly from those of their peers; and, 3) Experimentalists: people with greater uncertainty regarding their world-view versus the corresponding uncertainty of those around them. We can proceed by examining one entrepreneurial type per economy, thereby developing a theory with respect to the expected salient features associated with each.

In the next section we discuss the notion of a world-view in greater depth and explain why it may be crucial to our understanding of entrepreneurial activity. In Section 3, the Bayesian network is used as a tool to formalizes the idea of world-view. Following this, in Section 4, we present the details of a specific model designed to examine the issues raised above. In Section 5, we return to the Visionary versus Experimentalist motivations and explain how they are represented in this model. We then discuss expected results in Section 6 and conclude with a few brief comments in Section 7.

2 What's in a World-view?

As we consider which attributes motivate people to innovate, we must necessarily consider the dual question of what it is we mean by innovation. When we speak loosely of innovation, we usually refer to individuals whose economic behavior *successfully* departs in significant ways from the majority of those around them. Innovation also connotes a sense of leadership.

[2]Note that we do not concern ourselves with the managerial coordination or technological invention roles sometimes attributed to entrepreneurship.

That is, once successful departure from the norm is publically observed, the norm tends to shift toward the newly proven behavior. These associations seem most consistent with the Austrian view of entrepreneurs as agents-of-change who consciously drive the economy to new locations. They appear most inconsistent with the Knightian view of entrepreneurs as risk-takers. In the latter case, for example, we would expect equilibria to appear in which those individuals typified by lower levels of risk aversion consistently conduct the high-risk operations, while their counterparts continue to happily pursue less dangerous activities; the risk-takers do not drive the cohabitants of their model world to new behavior by dint of their example. The game theoretic position lies somewhere between these two views, but is probably closer to the Austrian view. The key difference between the Visionary and the Experimentalist is that the Visionary departs from the rest of the pack due to a strong conviction that her deviant world-view is correct, while the Experimentalist tries new activities due to uncertainty, or a lack of conviction, in a world-view that may actually be identical to those held by everyone else. While the end result, discovery of errant beliefs in the general population, may be the same, the motivations are quite distinct.

For the purpose of this discussion, we now focus only on the Experimentalist versus Visionary explanations of entrepreneurial activity. Toward the end of the paper, we will return to issues of risk aversion and suggest how they might be incorporated into the framework under consideration. In order to proceed, we must somehow capture in our model the feel of the Visionary explanation – the sense that entrepreneurs act out of a conviction that they, and they alone, actually "get it." They must equate the fact that everyone else appears to disagree more to profit opportunities than to evidence that they might be wrong. Implicit in this task is coming to grips with the associated issues of knowledge and learning. We must imbue the qualitative notion of "world-view" with some quantitative content. Recent advances in game theory appear to show strong promise toward this end.

Several game theorists have explored the results of loosening the common knowledge of rationality assumption implicit in the definition of Nash equilibria. Examples of this literature include Battigalli and Guatoli's "conjectural" equilibrium [1988], Rubinstein and Wolinsky's "rationalizable conjectural" equilibrium [1990], Fudenberg and Levine's "self-confirming" equilibrium [1993], and Kalai and Lehrer's "subjective" equilibrium [1993]. These analyses probe what happens in strategic situations in which agents are allowed to maintain erroneous beliefs about the strategies being implemented by their opponents with the proviso that the results they observe do not contradict these beliefs. Thus, player actions can be self-confirming in the sense that by expecting dire consequences to result from the play of certain actions, these actions are never played and knowledge of what would have actually happened had they been is never gained.

For the most part, these theories assume a high degree of common knowledge amongst the players. For example, players are generally assumed to know the complete structure of the game in which they are involved. Kalai and Lehrer's framework puts the least demands on the prior knowledge of the players, assuming that players may be ignorant about the structure of the game, its payoffs, and the actions available to other players (not to mention opponent strategies). From a game theoretic perspective, this latter approach is problematic because it completely lacks predictive value: *every* possible outcome is supportable as an equilibrium by simply investing players with whatever crazy beliefs allow it to be maintained. Yet, in developing a theory of entrepreneurial activity, particularly along the lines of the Austrian school, Kalai and Lehrer seem to come closest to the mark – presenting a framework that allows us to incorporate agents struggling primarily to figure out the way the world works not just the behavior of their opponents. This distinction is not trivial. Drucker [1985], for example, relates the following story:

> More than thirty years ago, I was told by the chairman of New York's largest department store, R. H. Macy, "We don't know how to stop the growth of appliance sales."
>
> "Why do you want to stop them?" I asked, quite mystified. "Are you losing money on them?"
>
> "On the contrary," the chairman said, "profit margins are better than on fashion goods; there are no returns, and practically no pilferage."
>
> "Do the appliance customers keep away the fashion customers?" I asked.
>
> "Oh no," was the answer. "Where we used to sell appliances primarily to people who came in to buy fashions, we now sell fashions very often to people who come in to buy appliances. But," the chairman continued, "in this kind of store, it is normal and healthy for fashion to produce seventy percent of sales. Appliance sales have grown so fast that they now account for three-fifths. And that's abnormal. We've tried everything we know to make fashion grow to restore the normal ratio, but nothing works. The only thing left now is to push appliance sales down to where they should be."

Drucker goes on to point out that, viewing the same trend, Bloomingdales, then a weak number four in the New York department store lineup, inferred that these observations were an indication of deeper changes in the marketplace. As a result, Bloomingdales built a new market position around its Housewares department, quickly rising to the number two position. Macy's was eventually forced to come to grips with the new reality and the number two and three stores of that age have long since disappeared altogether.

Note well that the problem in this story is not confusion over opponent's strategies, but the more fundamental question of comprehending the underlying structure of one's environment.[3] Still, the general game theoretic idea of self-confirming activity is clearly a useful one in this context. We submit that entrepreneurial innovation may be well described

[3]In general, we support the notion that the kind of innovation highlighted by Drucker's anecdotal tale is actually key to what drives economic efficiency and is indeed at the heart of the entrepreneurial function.

as situations in which deviant behavior, from whatever cause, leads to the collapse of erroneous beliefs on the part of the non-innovating agents, with the stipulation that the key beliefs at issue are those regarding environmental structure.

3 Formalizing "World-view"

In this section, we present a specific tool, the Bayesian network, that ultimately forms the core of our analytical approach. In order to motivate the choice of this particular apparatus, we begin by discussing the ideal set of characteristics sought in a framework designed to meet the broad objectives outlined in the preceding sections. We then present the specifics of the Bayesian network approach and highlight the points of congruence with our desired theoretical attributes.

3.1 Theoretical Requirements

We have identified our primary goal as exploring the opportunity presented by connecting the game-theoretic notion of self-confirming equilibria with the economic function of entrepreneurship.[4] Obviously, learning dynamics must play a key role in this theory. Thus, our model borrows much from the literature on non-cooperative, game theoretic learning. We present these details in a later section. Here, we wish to focus on those aspects of our approach that depart from standard methodologies.

Much of the literature in game-related learning involves agents competing in relatively simple environments, such as the Prisoner's Dilemma. While it is true that the *strategy spaces* associated with such games are indeed quite complex (thus presenting the players therein a non-trivial learning exercise), we do not view strategy-enlightenment as the primary benefit afforded a society by entrepreneurial activity. Our desire is to aim attention at an agent's attempt to construct a workable theory regarding his or her environment in the face of tangled and potentially ambiguous observational data. It is hard to imagine agents, even those at the low end of the bell curve, finding much difficulty in figuring out the structure inherent in the Prisoner's Dilemma game. Thus, we require: 1) enough complexity in the structure of the game to make our agents' environmental inference task interesting; and 2) a methodology for the representation of environmental knowledge (i.e. *world-views*). Assuming, for the moment, that the construction of an appropriately complex environment is not the significant stumbling block here, we center our attention instead on the knottier issue of how to depict environmental knowledge for the purpose of analyzing the role of the entrepreneur.

[4]Throughout this paper, the term "self-confirming" is used in its broad, non-technical, sense.

Accepting that the representation of structural knowledge is at the heart of our undertaking, we now identify the following desired attributes and explain why they are essential to our objectives:

Structural Theories from Data A good metaphor for what we seek is the agent-as-econometrician – a model in which everyone is out there acquiring data and forming theories with respect to how their actions influence the direction of the system (or key aspects of the system). Entrepreneurs are either agents with different theories, or are agents with a greater motivation than their peers to verify their theories. We might imagine these theories manifesting themselves as sets of structural equations, with parameters being estimated from the observed data. The problem with a specific implementation along the structural equations line is that the space of possible theories is huge (infinite). This is not workable. We need a framework that maintains the spirit of this metaphor on the one hand, while allowing us to work with a relatively small number of parameters on the other.

Bayesian Updating Game-theoretic learning models generally depend upon some form of Bayesian updating. Thus, in order to avail ourselves of results in this important area of game theory, our framework must be consistent with Bayesian learning.

Visionaries versus Experimentalists We wish to differentiate our models of entrepreneurial motivation along two dimensions: 1) specific beliefs; and 2) uncertainty in these beliefs. The Visionary model is one in which the entrepreneurial impulse results from beliefs that are at odds with those of the population around them. The Experimentalist model is one in which entrepreneurs may have similar beliefs, but experience greater uncertainty than their views. Specifically, we require a framework that allows us to depict one case (Visionary)in which agents have heterogeneous mean beliefs regarding potential structural theories but homogeneous mean variance with respect to these theories, and another (Experimentalist) in which agents have homogeneous mean beliefs, but heterogeneous variances in those beliefs.

Simultaneous Theories We wish to assume common knowledge of the space of all possible theories. As suggested above, entrepreneurs will differ either in the weights they put on these theories or in the confidence they place in them. Thus, our approach must also allow agents to sustain and evaluate several potential theories simultaneously.

Causality The search for causal explanations of real world phenomena appears to be an important characteristic of human nature.[5] This is consistent with our view. As we discuss above, we wish to create agents that develop theories with respect to how their actions influence the overall system, i.e. agents who search for causal relationships. Entrepreneurs may actually be individuals who see different causal implications in

[5]See Tversky and Kahneman [1977].

the statistical correlations they observe and who, therefore, choose to take different actions. Furthermore, as Pearl [1988] points out, "... people can easily and confidently detect dependencies, even though they may not be able to provide precise numerical estimates of probabilities. Evidently, the notions of relevance and dependence are far more basic to human reasoning than the numerical values attached to probability judgments." Although our agents will, in fact, provide precise numerical estimates of probabilities, we would not only like to build a theory consistent with these comments, but will actually find it helpful to utilize a framework that facilitates the identification of different causal theories on the part of the agents.

Extant Results We prefer utilizing existing frameworks to inventing our own, and, obviously, prefer existing frameworks attached to a greater volume useful literature and results.

3.2 Bayesian Networks

One solution that appears to meet all of the requirements laid out above is a tool developed in the decision sciences literature called the *Bayesian network*. The Bayesian network consists of two interconnected objects: 1) a set of dependency assertions that provide a qualitative summary of the causal relationships between the variables in the system; and, 2) a joint probability distribution over these variables. The dependency assertions are assumed to identify all relevant restrictions on the probability distribution. Suppose that an agent observes a set of N_V discrete variables each period defined as $V := \{C \cup X\}$ where

i) $C := \{C_1, \cdots, C_{N_C}\}$ is the set of N_C control variables available to agents,

ii) $X := \{X_1, \cdots, X_{N_X}\}$ is the set of N_X elements outside anyone's direct control, with the payoff set $\pi := X_{N_X} \in X$, defined as a special element of X, containing all payoff results.

Based upon these observations, the agent updates priors on a theory (or theories) regarding the causal structure assumed to be generating this data. The theories are defined as follows:

Definition 1 *Given the set of observed elements V, the set of* **direct causal relations** *$\vec{E}$ on the elements in V, denoted by $(V_q \rightarrow V_r)$ for some $V_q, V_r \in V$, means "V_q directly influences (or directly causes) V_r" such that,*

i) the relations are asymmetric, i.e. if $(V_i \rightarrow V_j)$ then $\neg(V_j \rightarrow V_i)$, *ii) if* V_r^+ *denotes the set of immediate causal predecessors or* **parents** *of* V_r*, then* $V_q \in V_r^+ \Longrightarrow (V_q \rightarrow V_r)$*,*

iii) cycles of the form $V_i \rightarrow V_j \rightarrow V_k \rightarrow V_i$ *are ruled out*

iv) $C_q^+ = \{\emptyset\} \quad \forall C_q \in C$ *(i.e. agents have free will to choose their actions)*

Definition 2 *The* k^{th} **structural hypothesis** *of an agent, denoted* $B^k = \{\vec{E}^k, P^k\}$*, consists of a set of causal relations* $\vec{E}_k$ *on the observed variables* V *and a probability measure* P^k *in the probability space* $(\Omega, \mathcal{F}, P^k)$ *where* $\Omega := C_1 \times \ldots \times C_{N_C} \times X_1 \times \ldots \times X_{N_X}$ *and* σ*-algebra* $\mathcal{F} := 2^{\Omega}$*. The set of all structural hypotheses is denoted by* $B := \{B^1, \ldots, B^K\}$*.*[6]

It is assumed that if a hypothesis is indeed true, then the set of causal relations corresponds to restrictions on the joint probability distribution P^k as follows:

Assumption 1 *Given the structural hypothesis* $B^k = \{\vec{E}^k, P^k\}$ *regarding a set of variables* V*, it is assumed that* P^k *obeys:*

$$P(v_1, \ldots, v_{N_V}) \;=\; P(v_1 \mid v_1^+) \times \ldots \times P(v_{N_V} \mid v_{N_V}^+),$$

where the v_i^+ *are determined by the set of causal relations* $\vec{E}^k$ *on* V*.*

This framework is appealing for several reasons. First, Bayesian networks were explicitly developed as data-driven causal inference mechanisms.[7] Thus, they are consistent both with our need to create structural theories from data (and come quite close to the metaphor of agent-as-econometrician described above), as well as with the observation that causality plays an important role in such efforts. Second, this is a Bayesian framework and can, therefore, be incorporated directly into a game-theoretic learning environment in which players must weigh a number of possible theories and act accordingly. Third, because the network provides an environmental theory based upon Bayesian reasoning, we can readily define Visionaries and Experimentalists in terms amenable to game-theoretic learning. Fourth, there is already a large extant literature on these objects, many results of which may be incorporated directly into our analysis.[8] Fifth, as we will see in a moment, the Bayesian network's causal ordering component provides a natural qualitative interpretation of the causal relationships present in the environment.[9] Finally, as an added bonus, this theory is highly compatible to empirical testing.[10]

---

[6]Note that these definitions allow us to easily represent a structural hypothesis B_k with a qualitative diagram, namely the graph with vertices V and edges $\vec{E}_k$.

[7]See, for example, Heckerman and Shachter [1995].

[8]See Shafer and Pearl [1990] and Edwards [1995].

[9]For an extensive treatment of this approach, see Pearl [1988].

[10]See Galles and Pearl [1995].

4 An Environmentally Challenging Example

In order to illustrate the setup, we now present an example of the type of model we have in mind. Suppose that we have a situation in which F firms are competing in a particular market producing a good X. To keep things as simple as possible, assume that prices are fixed at some level $\bar{p}$, forcing firms to compete along quality dimensions.[11] We assume a population of I consumers, each of whom purchases one X per period. Now, suppose that firms can add some feature to the product, where $\lambda_j = 1$ indicates the decision by firm $j \in F$ to offer the feature, and $\lambda_j = 0$ indicates the decision not to. Thus, products $X \in \{X_0, X_1\}$ are differentiated by the feature (but not by firm). We assume a positive cost C_λ of adding the feature. Let consumers' preferences by represented by a random utilities model such that each period consumers either: 1) are indifferent toward the feature or 2) strictly prefer it (with random noise leading a consumer to switch from time to time). Consumer preferences are indicated each period by $\delta_i \in \{0, 1\}$ for each $i \in I$ where $\delta_i = 0$ indicates indifference to the feature and $\delta_i = 1$ indicates a preference for it. Finally, suppose each firm can choose to advertise the feature $\alpha_f = 1$ in an attempt to increase consumers' preferences for it (or choose not to advertise, $\alpha_f = 0$) at a cost of C_α. Finally, suppose that, each period, firm j observes nothing prior to making the joint feature/advertising decision. Assume this is a repeated game and that, after each period, the firm observes the following ex-post data (we drop time subscripts for clarity):

(1) $S = \sum_{f \in F} \lambda_f$, the number of firms offering the feature,

(2) $A = \sum_{f \in F} \alpha_f$, the number of firms advertising,

(3) $M = \sum_{i \in I} \delta_i$, the number of consumers with preferences for the feature prior to purchase,[12]

(4) $q_j = \begin{cases} \frac{I-M}{F} & \text{if } \lambda_f = 0, S > 0 \\ \frac{I}{F} & \text{if } \lambda_f = 0, S = 0 \\ \frac{I-M}{F} + \frac{M}{S} & \text{if } \lambda_f = 1 \end{cases}$

(5) $\pi_j = \bar{p} q_j - \lambda_j C_\lambda - \alpha_j C_\alpha$, the firm's own profitability.

The key is the quantity variable, which depends upon whether the firm decides to add the feature, how many other firms produce the enhanced version, and the share of consumers who strongly prefer it. The incentive to produce X_1 goes down as S increases and/or M decreases. Suppose these equations are common knowledge. Given the (unknown) strategies of other

[11] We might imagine a government-regulated industry or a retail market in which vertical price restraints are in effect as described by Klein and Murphy [1988].

[12] i.e. assume each firm has excellent ex-post market research surveys

firms and the random element in consumer demand, firms must estimate the expected value of their profits conditional upon their choice of control variables.

4.1 Causal Theories

We assume that firms construct causal theories in the form of Bayesian networks in order to estimate the effects of their actions. Using the bayesian network notation, C represents the set of control variables and X represents the set of observed variables. In our example, $C = \{\lambda_j \cup \lambda_{-j} \cup \alpha_j \cup \alpha_{-j}\}$ and $X = \{\delta \cup S \cup A \cup M \cup q_j \cup \pi_j\}$.[13] The firms' k^{th} theory consists of the causal ordering $\vec{E}^k$ and the corresponding probability distribution P^k.

For example, one potential hypothesis, which we'll call the *Marketeer's View*, is illustrated in Figure 1.[14] This graph represents the hypothesis that advertising actually influences preferences. Note that in this theory other player's actions are treated as if they are random variables chosen from some stationary distribution. The arc from λ_{-j} to α_{-j} indicates the rational knowledge that firms don't advertise the added feature in periods they don't produce it. For comparison, the graph corresponding to what we'll call the *Economist's Hypothesis* – that advertising has no effect on preferences – is shown in Figure 2. This theory is constructed by simply removing the arc $A \longrightarrow \delta$.

The firm wishes to calculate the expected value of its profit given its choice of actions. Let V_j^+ represent the set of direct parents for one of the system variables $V_j \in V$. Under Assumption 1, node separation in the graph corresponds to independencies in the joint probability distribution such that the joint distribution can be written down as the product of the conditional probability distributions $\prod_{V_j \in V} P(V_j \mid V_j^+)$. This turns out to be especially tractable in our discrete variable case (although combinations of discrete and continuous variables can be included within this framework[15]). The total number of hypotheses possible, given a *complete* lack of structural knowledge, is large (the number of variables in the system raised to the power of the number of values they each assume). However, reasonable assumptions with respect to the prior information possessed by the agents can often reduce these numbers to a manageable level.

[13] Note that, following the usual notation,

$$\begin{aligned}\lambda_{-j} &= (\lambda_1, \ldots, \lambda_{j-1}, \lambda_{j+1}, \ldots, \lambda_F)\\ \alpha_{-j} &= (\alpha_1, \ldots, \alpha_{j-1}, \alpha_{j+1}, \ldots, \alpha_F)\\ \delta &= (\delta_1, \ldots, \delta_I)\end{aligned}$$

[14] In keeping with standard influence diagram notation: the square nodes indicate the firm's decision variables, circular nodes indicate random variables (from the firm's perspective), and circular nodes with a double line indicate deterministic variables.

[15] See Heckerman [1995].

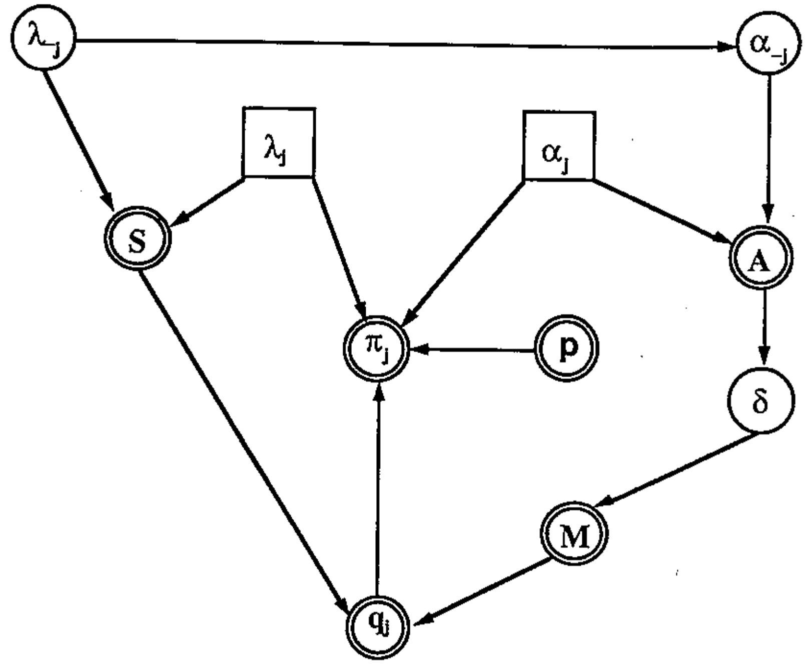

Figure 1: The Marketeer's View

Thus, the joint probability distribution corresponding to Figure 1 can be written in the form:

(4.1.1) $$P(V) = P(\lambda_{-j})P(\alpha_{-j} \mid \lambda_{-j})P(A \mid \alpha)P(S \mid \lambda)P(\delta \mid A)P(M \mid \delta)P(q_j \mid M, S)P(\pi_j \mid q_j, \lambda_j, \alpha_j, \bar{p}),$$

Where $\alpha = \{\alpha_j, \alpha_{-j}\}$ and $\lambda = \{\lambda_j, \lambda_{-j}\}$. By eliminating the deterministic relationships, we can write down the equation of primary interest to the firm:

(4.1.2) $$P(\pi_j, \delta, \lambda, \alpha) = P(\lambda_j)P(\alpha_j)P(\lambda_{-j})P(\alpha_{-j} \mid \lambda_{-j})P(\delta \mid \alpha)P(\pi_j \mid \delta, \lambda, \alpha).$$

We will return to illustrate how the firm evaluates these quantities in a moment. First, though, we describe the learning assumptions.

4.2 Learning Dynamics

The example described above is a simultaneous stage game, with each firm making the joint decision regarding advertising and product features at the same time. There are a fixed number of F players who, we assume, all observe the same variables (in this case all the variables in the system). Later, we imagine introducing informational asymmetries by allowing firms to observe only their own move and some aggregate statistics. Furthermore, we assume a form of fictitious play: each period firms play a best response to the joint distribution implied by their posteriors as updated by the historical frequency of play. Thus,

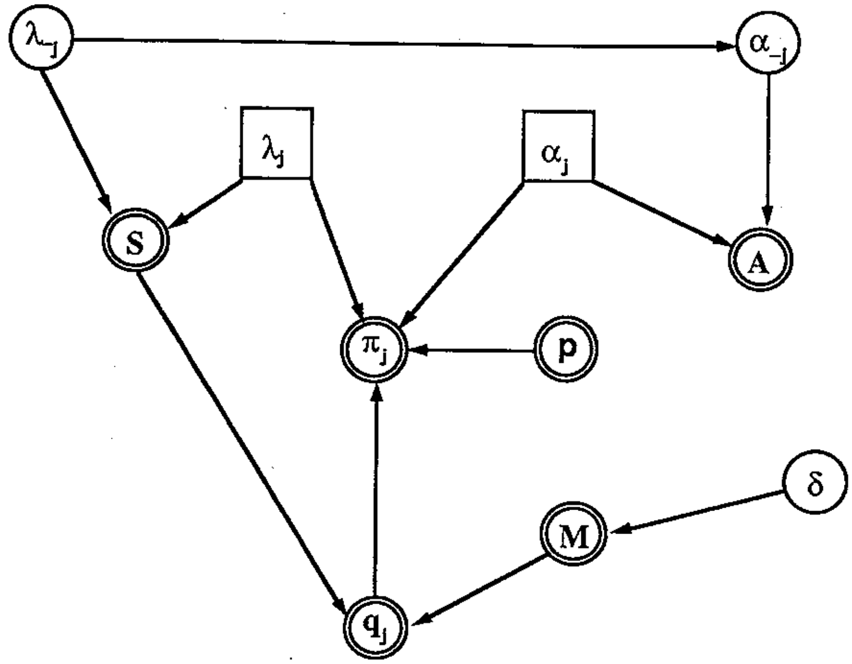

Figure 2: The Economist's Hypothesis

each firm treats the actions of its opponents as if they result from a stationary distribution, even though these frequencies actual change as their opponents also learn. In later work involving information asymmetries, we will employ partial best response dynamics in which only a fixed percentage of the population is allowed to switch from their current action to a best response to the joint distribution implied by the aggregate statistics.

4.3 Discovering Probabilities and Structure

Finally, we come to the issue of how firms evaluate their environments. To facilitate our discussion, we now further simplify the model by assuming that there are only *two* firms in the market. Furthermore, we assume that it is common knowledge that the universe of possible hypotheses is limited to the two depicted in Figures 1 & 2. That is, both firms know the fact that either the Marketeer's View *or* the Economist's Hypothesis is true (to the exclusion of all other possible theories). However, neither firm knows for certain *which* theory, respectively referred to as B^M and B^E, is actually correct. Thus, firm i places probability $0 < \beta_i < 1$ on the proposition $(B^M = T)$ and probability $1-\beta_i$ on the proposition $(B^M = F)$, where $(B^M = F) \Longrightarrow (B^E = T)$. Keep in mind that we are assuming $(B^M = T)$ describes the true state of the world.

From the preceding discussion, we know that associated with each of these hypotheses are the probability distributions P^M and P^E. From the definition of a Bayesian network, we also know that these distributions can be described by the product of the conditional distributions given in equation 4.1.1. Each firm k is concerned with the random variables

δ, λ_j and α_j, where the latter two variables are binomial and δ is the sequence of I binomial outcomes resulting from each of the δ_i's. Thus, the joint probability distribution is simply determined by a set of binomial probability parameters which we now define as $\Theta := \{\sigma \cup \omega\}$, where $\sigma := \Delta(C)$ is the set of probabilities on control variables and $\omega := \Delta(X)$ is the set of probabilities on non-control variables. Finally, we assume that the consumers are homogeneous.

The knowledge of Firm 1 is summarized by the parameters $\{\beta_1, \Theta_1\}$. So, for example, the parameters associated with the Marketeer's View can be depicted as the following set of conditional probability tables (with the sigma's referring to the probabilities associated with his *opponent's* strategies):

$\mathbf{B_1^M}$

λ_2

0	1
$1-\sigma^M$	σ^M

α_2

λ_2	0	1
0	$1-\sigma_0^M$	σ_0^M
1	$1-\sigma_1^M$	σ_1^M

δ_i

A	0	1
0	$1-\omega_0^M$	ω_0^M
1	$1-\omega_1^M$	ω_1^M
2	$1-\omega_2^M$	ω_2^M

Thus, σ is Firm 1's estimate of the probability that Firm 2 will play $\lambda = 1$, while σ_λ is her estimate of the probability that $\alpha_2 = 1$ given her opponent's choice of λ. As a result, for example, $\sigma(1,1) = \sigma\sigma_1$. We'll assume that everyone knows firms don't advertise features their products don't have and set $\sigma_0 = 0$ accordingly. Furthermore, we assume that consumers care only about whether or not advertising is done, not about the actual level of advertising. So, we set $\omega_1 = \omega_2 = \omega_A$. We can write these simplified parameters as:

$\mathbf{B_1^M}$

λ_2

0	1
$1-\sigma^M$	σ^M

α_2

λ_2	0	1
0	1	0
1	$1-\sigma_1^M$	σ_1^M

δ_i

A	0	1
$A=0$	$1-\omega_0^M$	ω_0^M
$A>0$	$1-\omega_A^M$	ω_A^M

Similarly, the competing hypothesis B^E for Firm 1 (i.e. the one shown in Figure 2) has parameters:

$\mathbf{B_1^E}$

λ_2

0	1
$1-\sigma^E$	σ^E

α_2

λ_2	0	1
0	1	0
1	$1-\sigma_1^E$	σ_1^E

δ_i

0	1
$1-\omega^E$	ω^E

It is worth highlighting the fact that firms can have different parameters under each of the two possible distributions. For example, a firm may believe that in the Economist's world firms may be less likely to produce the feature than in the Marketeer's world (i.e. $\sigma^M > \sigma^E$).

Firms must estimate both the β's and the parameters corresponding to each of the competing hypotheses. We assume that the unknown probabilities $\rho_i \in \{\beta \cup \Theta\}$ are positive, independent, and each have the beta distribution:

$$f(\rho_i) = c \cdot \rho_i^{N_1^i - 1}(1-\rho_i)^{N_2^i - 1}$$

where c is a normalization constant. The nice feature about this distribution is that the expected value is given by:

$$\bar{\rho}_i = \frac{N_1^i}{N^i}$$

and variance given by:

$$Var(\rho_i) = \frac{\bar{\rho}_i(1-\bar{\rho}_i)}{N^i + 1}$$

where $N^i = N_1^i + N_2^i$ is a reflection of the firm's confidence in the expected value of theta. One interesting interpretation of this formula is that N^i is the estimated number of observations a player would have to see, starting from complete ignorance, in order to attain the present level of confidence in the parameter. We assume the firms begin the game with priors (N_1^i, N_2^i) for all parameters, including β_1.

If $v_{kt} \in \{C \times X\}$ denotes firm k's observation of V following period t, then the database available to the firm for decision making in period $t+1$ is given by $D := \{v_{i1}, \ldots, v_{it}\}$. Following Heckerman [1995], let r_i be the number of states of variable $V_i \in V$. Let V_i^{k+} be the parents of V_i under the structural theory B^k. Then, $q_i = \prod_{V_l \in V_i^{k+}} r_l$ is the number of states of V_i^{k+}. Let $\theta_{ij}^k \in \Theta$ denote the probability that $V_i = 1$ given $V_i^{k+} = j$. Let N_{ij0}^k and N_{ij1}^k represent the firm's priors on θ_{ij}^k. Then, if N'_{ij1} is the number of cases in D in which $V_i = 1$ and $V_i^{k+} = j$, we obtain the updating formula

$$f(\theta_{ij}^k \mid D, B^k) = c \cdot (\theta_{ij}^k)^{N_{ij1}^k + N'_{ij1} - 1}(1-\theta_{ij}^k)^{N_{ij0}^k + N'_{ij0} - 1}$$

which is actually quite straightforward. Furthermore, it can be shown that

$$f(B^k \mid D) = c_k \cdot \beta_k \cdot \prod_{i=1}^{N_V} \prod_{j=1}^{q_i} \frac{\Gamma(N_{ij}^k)}{\Gamma(N_{ij}^k + N'_{ij})} \cdot \prod_{l=0}^{1} \frac{\Gamma(N_{ijl}^k + N'_{ijl})}{\Gamma(N_{ijl}^k)},$$

where $\Gamma(N)$ is the gamma function with parameter N.

4.4 Nash Equilibria

We close this section by mentioning initial results regarding the potential stage game Nash equilibria that arise when the Marketeer's View reflects the true state of the world. Nash equilibria are determined by the cost parameters $\{C_\lambda, C_\alpha\}$ and the true underlying conditional utility probabilities $P(\delta \mid A)$. Depending upon these parameters, one of four stage game outcomes is possible:

1. One pure strategy Nash equilibrium in which both firms produce the feature, but neither advertises.
2. One pure strategy Nash equilibrium in which neither firm produces the feature nor advertises.
3. Three pure strategy Nash equilibria: a) Firm 1 produces the feature and advertises while Firm 2 does neither; b) Firm 2 produces the feature and advertises while Firm 1 does neither; and c) both firms produce the feature, but neither advertises.
4. There is also a mixed equilibria in which each firm produces the feature and advertises with the same frequency as the other.

The equilibria for $(B^E = T)$ are a subset of these. Also, it is not yet determined with which of the first three cases, if any, the mixed equilibrium is consistent.

5 Entrepreneurs

These details now allow us to describe the specifics of the Visionary versus Experimentalist models mentioned at the beginning of the paper. In the Visionary case, we propose populating the model with firms heterogeneous in mean beliefs on β, but homogeneous in the confidence in those beliefs (i.e. equivalent N^β's). The Experimentalist model, on the other hand, takes the opposite position, with firms homogeneous in mean beliefs and heterogeneous in variances on those beliefs.

For example, assume that the Marketeer's View is correct ($\beta = 1$) with $\omega_0 \simeq 0$ and $\omega_1 \simeq 1$ (in other words, advertising has a big effect). Suppose, however, that $F - 1$ firms believe that $\beta \simeq 0$ with $\omega^E \simeq 0$ (advertising has no effect and people don't care about the feature). Suppose that, in contrast to this, one firm holds different beliefs. This firm believes

$\beta \simeq 1$ with $\omega_0^M \simeq 1$ and $\omega_1^M > \omega_0^M$ (advertising has an effect, but there is no point in doing it since people want the feature anyway). Initially, the $F-1$ firms will neither produce the feature nor advertise. The potential innovator will produce the feature but not advertise. As time passes, the $F-1$ firms strengthen their self-confirming beliefs. The deviant firm, however, will begin to see that its priors on ω_0^M are erroneous. Eventually a trigger point will be reached at which time the deviant firm will innovate, both producing the feature and advertising. Then, and for some time following, a period of confusion will evolve in which the $F-1$ firms try to figure out what it is about their world-view that they've got wrong, with the innovator raking in quasi-rents all the while. Eventually, the industry will settle down at a new equilibrium – having been driven there by the deviant beliefs and "courageous" action of the entrepreneur.

Note that as the model has been structured, the firms all face a multi-armed bandit problem. The setup provides all the information necessary to determine optimal experimentation rules in the face of uncertainty per a Gittens index framework. This side of the model still requires a lot of attention.

6 Anticipated Results

We believe that this framework provides a a very rich model with which to study the implications of the different entrepreneurial assumptions. Of course, it is difficult to anticipate all of the results without having worked through the details. Still some potential results do spring to mind.

For example, persistence is an important element in the mix of entrepreneurial attributes. Innovators must often undergo long periods of sacrifice and even public ridicule before achieving success. What qualities allow them to maintain their activities in the face of these significant personal costs? We posit that, given their strong convictions, Visionaries will persist much longer than their Experimentalist counterparts, who (by definition) have negative expected returns to their activity.

In a related issue, we would like to examine the relative numbers of each type of entrepreneur required to move society away from erroneous beliefs. With reasoning similar to the comments above, we believe that, other things equal, it is likely that fewer Visionaries are required to provide a conclusive societal correction.

On the other hand, it is possible to consider those situations in which one type is preferable, from an efficiency standpoint, to the other. For example, in the case where the majority view is actually correct, relentless pursuit of entrepreneurial quasi-rents based on erroneous vision could be costly to everyone. In the case of finitely-lived agents, brief public experiments may serve to maintain the accuracy of everyone's beliefs. Or, we might imagine a model in which private information accrues to firms within the industry. Although superior beliefs may be the basis of superior profits, at some point it may actually behoove incumbents to make their knowledge public in order to dissuade certain individuals (those who would enter on the basis of erroneous beliefs) from disrupting the market.

There are also institutional implications to these ideas. For example, it is difficult to imagine Experimentalists being very successful in their ability to raise venture capital funding from third parties. On the other hand, corporations may have a significant incentive to experiment given a long-term commitment to their own industry. This may help explain the difference between entrepreneurial and "intrapreneurial" sources of venture funding.

7 Conclusion

Hopefully, the preceding section provides a glimpse of the many avenues that we expect to arise from this framework. Although the proposed model joins two separate technical areas – Bayesian networks and evolutionary game theory[16] – we feel that the two are quite compatible and that sufficient results exist within each body of literature to make this a manageable proposition.

[16]Three, if you count the bandit literature.

Bibliography

Banks, J. S. and R. K. Sundaram (1992). "A Class of Bandit Problems Yielding Myopic Optimal Strategies." Journal of Applied Probability 29: 625-632.

Banks, J. S. and R. K. Sundaram (1992). "Denumerable-Armed Bandits." Econometrica 60(5): 1071-1096.

Battigalli, P. and D. Guatoli (1988). Conjectural Equilibria and Rationalizability in a Macroeconomic Game with Incomplete Information. Milan, Istituto Economia Politica.

Clarke, J. B. (1918). Essentials of Economic Theory. New York, Macmillan.

Edwards, D. (1995). Introduction to Graphical Modelling. New York, NY, Springer-Verlag.

Fudenberg, D. and D. K. Levine (1993). "Self-confirming Equilibrium." Econometrica 61(3): 523-45.

Fudenberg, D. and D. K. Levine (1995). Theory of Learning in Games, Unpublished Manuscript.

Gittins, J. C. and D. M. Jones (1972). A Dynamic Allocation Index for the Sequential Design of Experiments. Progress in Statistics, Budapest, North-Holland Publishing Company.

Heckerman, D. (1995). A Bayesian Approach to Learning Causal Networks. Redmond, Microsoft.

Heckerman, D., D. Geiger, et al. (1995). Learning Bayesian Networks: the combination of knowledge and statistical data. Redmond, Microsoft.

Jevons, H. S. (1871). The Theory of Political Economy. New York, Sentry Press.

Kalai, E. and E. Lehrer (1995). "Subjective games and equilibria." Games and economic behavior 8: 123-63.

Klein, B. and K. M. Murphy (1988). "Vertical Restraints as Contract Enforcement Mechanisms." Journal of Law and Economics XXXI.

Knight, F. H. (1933). The Economic Organization. New York, H. Wolff.

Menger, C. (1871). Principles of Economics. Grove City, Pa., Libertarian Press.

Pearl, J. (1988). Probabilisitic reasoning in intelligent systems: networks of plausible inference. San Francisco, Morgan Kaufmann.

Rubinstein, A. and A. Wolinsky (1990). Rationalizable Conjectural Equilibrium: Between Nash and Rationalizability. Ramat Aviv, Tel-Aviv University.

Schumpeter, J. A. (1942). Capitalism, Socialism and Democracy. New York, Harper & Bros.

Shafer, G. and J. Pearl, Eds. (1990). Readings in uncertain reasoning. San Mateo, Morgan Kaufmann.