PERCIPIENT CONSULTING WORKING PAPER:

A STATE SPACE MODEL OF THE ECONOMIC AGENT

R. David Seabrook*

The fundamental importance of information in economics demands a new model of the economic agent. Dynamic state space analysis can model important phenomena such as information filtering, experience and expertise, learning, dynamic rent-seeking behavior, the interaction of information, knowledge and physical work and the build versus buy decision. In this model, information and knowledge are not the same thing. The set of possible interactions between two agents provides a framework for understanding sales and marketing, and provides insights into business strategy as dynamic rent-seeking behavior. The most important benefit of the state space approach may be in placing information, knowledge, innovation and learning at the center of economic analysis.

I. Introduction

This paper proposes a general state space model of an economic agent. We now know that many results derived using the neoclassical model of people as perfectly informed, utility-maximizing decision-makers do not hold when information is not perfect (for an overview, see Joseph E. Stiglitz [2000]). The problem of "irrational" behavior has further challenged the neoclassical model (see for example Daniel Kahneman and Amos Tversky [1979]). The neoclassical model turns out to be a special case, unsuitable as a foundation for general conclusions. Yet – although it demonstrates the usefulness of ability to "signal" and "screen"-information economics does not derive a new model of the economic agent.

Game theory provides an important alternative to the neoclassical model. The model of "people and firms as game-players" is dynamic where the standard competitive equilibrium model is static. Games can model strategies under imperfect information. Models of learning provide an explanation for how equilibria come about (see for example Drew Fudenberg and David K. Levine [1993]), yet the nature of equilibrium is still the primary focus of the analysis. Non-steady-state phenomena have largely been ignored, and many results are not general, but dependent on the rules of the game.

Models of firm behavior are even more varied than models of individual behavior. As Oliver E. Williamson [2002] points out, the "lens of contract" offers a fundamentally different focus than the "lens of choice" in describing firm behavior. Yet, firms and individuals seem to have much in common, and the neoclassical model provides the benefit of emphasizing these similarities.

Apart from imperfect information, irrational behavior, non-steady-state phenomena, rule-dependence and the importance of contracts, other issues challenge current models:

- If equilibrium doesn't happen under imperfect information, what *does* happen?
- Are information and knowledge the same thing, or does "perfect knowledge" represent a different set of hidden assumptions in the standard model?

This paper develops a general model of an economic agent, using state space analysis and assuming only that the agent can interact with its environment. Information is central – it is the only input to the model. The purpose of this approach is to establish firm foundations for more specific models by eliminating all but the most fundamental assumptions.

The model suggests some useful conclusions in its own right: information and knowledge are not the same thing; the fundamental “factors of production” are information, knowledge and work; learning is a fundamental aspect of economic behavior; knowledge is a source of rents; the optimization problem is to allocate work across information, knowledge and physical work activities; equilibria and stability depend on the state space concepts of controllability and observability.

The model provides insights into sales and marketing, business strategy and organizational design.

The resulting model, perhaps surprisingly, is intuitive: “Economic agents perform information, knowledge and physical work”. “They have a limited amount of energy, and have to decide how to divide their energy across information, knowledge and work activities”.

II. REQUIREMENTS FOR A GENERAL MODEL

A general model of an economic agent should:

- Require only axiomatic assumptions.
- Assume no specific exchange mechanism or property rights.
- Model information explicitly. In particular, information need not be perfect, and costly search should be the default behavior.
- Allow preferences to change.
- Not require “rational” behavior (von Neumann-Morgenstern utility-maximizing).
- Be the same for producers and consumers. The general case must assume that a need can be met by an agent acting as a producer on its own behalf (build versus buy decision).
- Model dynamic behavior. A dynamic model should reveal important phenomena that might be concealed by equilibrium analysis: transient economic rents, returns to innovation (Schumpeterian competition), returns to learning strategies, inherently non-equilibrium phenomena (for example, asset price bubbles, returns to private information, one-off events such as IPOs, “trendy” goods and terrorist attacks).
- Reduce to the familiar special cases under certain explicit assumptions, such as an agent with fixed preferences and perfect information.
- Retain the budget constraint.
- Model “learning” and “knowledge”.

This last point requires some expansion. Economists accept that “knowledge” and “learning” are of fundamental importance to the economy, yet most current models ignore these terms or at best model very specific cases - even the terminology is often imprecise. Paul A. Samuelson [1985, page 662] notes the importance of rents from innovation: “just as one source of innovational profits is disappearing, another is being born. So these innovational profits will continue to exist”. Stiglitz [2000] argues that “the key question is one of dynamics: how the economy adapts to new information, creates new knowledge, and how that knowledge is disseminated, absorbed and used throughout the economy”.

Current models ignore these terms or at best model very specific cases - even the terminology is often imprecise. The neoclassical producer model assumes constant technology and thus ignores innovation, arguably the single most important aspect of the economy. The neoclassical consumer model assumes fixed preferences. Michael Spence [1970, pages 358 and

359] identifies the hiring process as a learning one, but limits his analysis to the equilibrium case "to avoid studying a system in a continual state of flux". The game theory literature has gone the furthest in addressing dynamic learning behavior (for example, see Fudenberg and Levine [1993]), but focuses on specific cases rather than general results. The terms "information" and "knowledge" are often used interchangeably and without definition. For example, George J. Stigler [1961, page 213] says "One should hardly have to tell academicians that information is a valuable resource: knowledge *is* power".

II. OUTLINE OF THE STATE SPACE MODEL

A. Terminology and Notation

State space models are used to model physical systems whose behavior changes with time. The following terminology is taken from Donald M. Wiberg [1971]:

> The state of a physical object is any property of the object which relates input to output such that knowledge of the input time function for t m t_0 and state at time t= t_0 completely determines a unique output for t m t_0 " [Wiberg, 1971, page 1]

I will represent state space models graphically as shown in Figure 1:

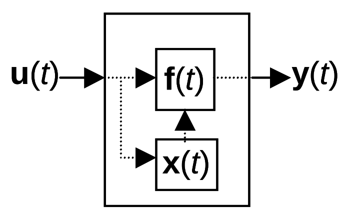

Figure 1. Representation of a State Space Model.

The state variable **x**(t) is an explicit function of time, but also depends implicitly on the starting time t_0, the initial state $\mathbf{x}(t_0)=\mathbf{x}_0$, and the input **u**(t)[1]:

(1) $$\mathbf{x}(t)=\mathbf{g}(t;t_0,\mathbf{x}_0,\mathbf{u}(t))$$

x(t) follows a *trajectory* over time. The output **y**(t) of the model is a function of the input **u**(t) and the current state:

(2) $$\mathbf{y}(t) = \mathbf{f}(t,\mathbf{x}(t),\mathbf{u}(t))$$

x(t) parameterizes the behavior of the agent[2]. **x**(t) depends on the initial state $\mathbf{x}_0$. The *state space* is the set of all possible values of **x**(t).

In this paper, a "signal" is any input or output of a system, in this case **u**(t) or **y**(t). A signal carries "information" if the model cannot predict the signal for all values of *t*. The information contained in the signal may or may not be correlated with expected utility.[3] I define "information" to mean the information contained in the input or output signal of a state space model, although I will often use "information" and "input" interchangeably[4]. The input signal will typically consist of price, quality, contract terms and other information about the environment.

B. State Space Modeling of the Interaction of Two Economic Agents

This section outlines the state space model of an economic agent. A more rigorous treatment is given in Section IV. Figure 2 illustrates the interaction of two economic agents, using state space models:

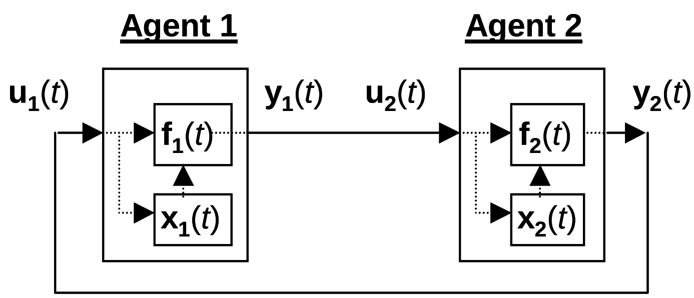

Figure 2. Interaction of Two Agents.

In this model, an agent interacts with the environment only through its input **u**(t) or output **y**(t). The agent's behavior is a function of time, the input and the current state. The state vector **x**(t) represents the current state of the economic agent. It will always contain a term w(t), representing the agent's *energy* or *ability to do work*. For a consumer, this term will include the budget, or ability to pay. By definition, any agent with *zero energy* or w(t)=0 cannot interact with its environment.

To illustrate, let the input to an economic agent be **u**(t)={for sale, brand Z cola, price $1}, and the function **f**(t)={buy anything where price [expected utility}. Then the output **y**(t)={buy, brand Z cola,$1} can only occur if one element of the state vector **x**(t) is $x_i \in \mathbf{x}$ = {brand Z cola, expected utilitym$1} and w(t) m $1.

In principle, the other way this transaction could take place is if the input signal contained the utility information: **u**(t)={for sale, brand Z cola, utility $1}. As I will shown, in our model this cannot happen, because utility is inferred depending on the current state. If our agent is not thirsty, the expected utility may be much less than $1.

III. APPLICATION OF THE STATE SPACE MODEL - AUTOMOBILE PURCHASE EXAMPLE

The application of a state space model is best illustrated with an example. I will first examine George A. Akerlof's famous lemons case [Akerlof, 1970], and then augment the example to explore other phenomena.

Represent each group of traders by a state space model, as shown in Figure 3:

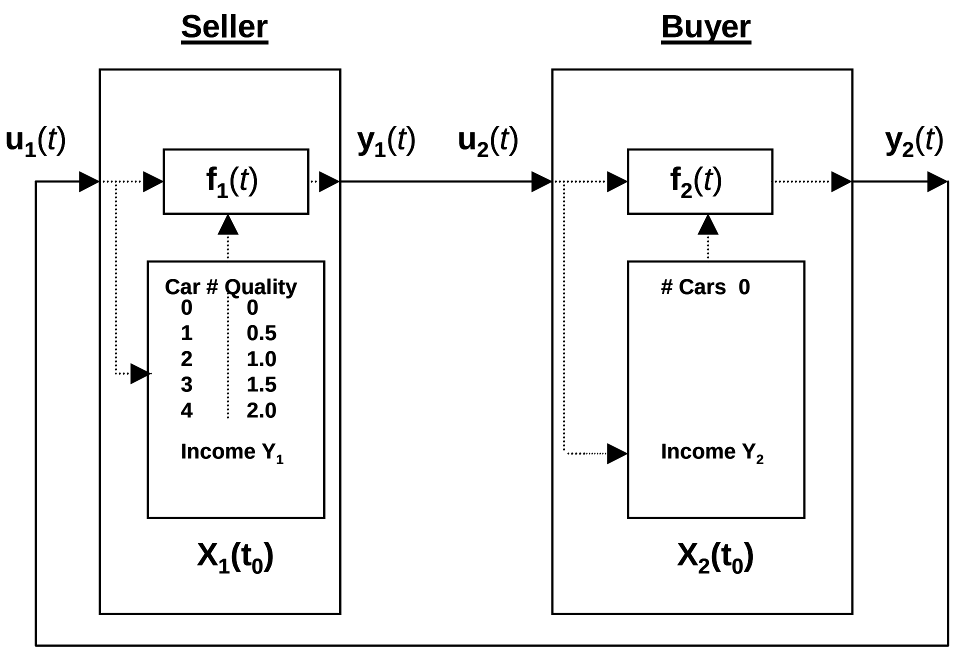

Figure 3. Initial State for Automobile Example

Assume for the purposes of this example that an exchange mechanism exists and the operations "bid", "offer", "buy" and "sell" are defined with the usual meaning (I will generalize later).

A. Information Filtering and the Initial State Vector

Define Group 1 and Group 2 traders as utility maximizing. The state of Group 1 traders $\mathbf{x}_1(t)$ is the list of cars owned by them, the quality of the cars and the income $M_1(t)$. The initial state of Group 1 traders $\mathbf{x}_1(t=t_0)$ = {(car 1, quality 0),(car 2, quality 0.5),(car 3, quality 1),(car 4, quality 1.5),(car 5, quality 2),(income=M_1)}. Let the output of Group 1 traders $\mathbf{y}_1(t)=\mathbf{f}_1(t, \mathbf{x}_1(t),\mathbf{u}_1(t))$ be $\{s_1(t)$,car, $p_1(t)\}$, where $s_1(t) \in$ (for sale, sold), car $\in$ (1..5) represents the car for sale and $p_1(t)$ represents the offer price of the Group 1 traders at time t.

Define the state of Group 2 traders $\mathbf{x}_2(t)$ as the number of cars owned $n_2(t)$, income $M_2(t)$ and quality of the offer at time t, $V_2(t)$. Set the initial state of Group 2 traders:

(3) $\mathbf{x}_2(t=t_0)$={n_2=0,income=M_2,V_2(car)=unknown for all car}.

Let the output of Group 2 traders $\mathbf{y}_2(t)$ be $\{s_2(t),p_2(t)\}$, where $s_2(t) \in$ (no bid, buy) and $p_2(t)$ represents the bid price of the Group 2 traders at time t. Let the utility of a car for Group 2 be 3/2*quality of the car.

Then the output of the Group 2 traders model is $\mathbf{y}_2(t)=\mathbf{f}_2(t, \mathbf{x}_2(t),\mathbf{u}_2(t))$:

(4) $\mathbf{f}_2(t)$= while income $M_2(t)$ m0, if $s_1(t)$="for sale" and $p_1(t)$ [$3/2*V_2(t)$, then $\mathbf{y}_2(t)$=(buy,p)
if $p_1(t) > 3/2*V_2(t)$, then $\mathbf{y}_2(t)$=(no bid)
if $V_2(t)$ =unknown, then $\mathbf{y}_2(t)$=(no bid)

For notational simplicity, I will use discrete time intervals. The state vector of the Group 2 traders is:

(5) $\mathbf{x}_2(t)=\mathbf{g}_2(t;t_0,\mathbf{x}_0,\mathbf{u}(t))=\{n_2(t), M_2(t), V_2(t)\}$

(6) $n_2(t+1)$ $= n_2(t)+1$ if $s_1(t)=$"sold"
$= n_2(t)$ otherwise

(7) $M_2(t+1)$ $= M_2(t)-p_1(t)$ if $s_1(t)=$"sold"
$= M_2(t)$ otherwise

(8) $V_2(t+1)$ $= V_2(t)=\text{quality(car)}$

Set the initial output signal from the Group 1 traders to:

(9) $\mathbf{y}_1(t_0)= \mathbf{u}_2(t_0)=\{\text{for sale,car},p_1=2\}$

In this form the problem is trivial: because quality is unknown it follows from (3), (8) and (4) that there is no state change. That is,

$$\mathbf{x}_2(t)=\mathbf{x}_0 \text{ and } \mathbf{y}_2(t)=\mathbf{y}_2(t=t_0) \text{ for all } t > t_0$$

Definition 1: Define *learning* as a change in state of those elements of the state vector that allow the agent to infer expected utility from the information input.

Proposition 1: For a transaction to take place without learning, there must be some elements of the initial state vector that allow the agent to infer expected utility from the information in the input signal.

Proof of Proposition 1: For a transaction to occur we must have $\mathbf{y}_2(t)$g$\mathbf{y}_2(t_0)$ for some $t>t_0$. But $\mathbf{y}_2(t)=\mathbf{y}_2(t_0)$=(no bid) while $V_2(t)$ is unknown. If in the initial state $V_2(t_0)$ is unknown, then without learning the is no change in state and $V_2(t)= V_2(t_0)$ for all t. Then $\mathbf{y}_2(t)=\mathbf{y}_2(t_0)$ for all t and no transaction will take place. Hence for a transaction to take place $V_2(t_0)$ cannot be unknown.

Definition 2: Define *knowledge* as those elements of the state vector that allow the agent to infer expected utility from the information input.

The functions $\mathbf{f}(t)$ and $\mathbf{g}(t)$ and the initial state $\mathbf{x}_0$ specify a *filter* of the input information. Because this model cannot learn and will only respond when quality is known, the information content of the input signal is *filtered out.* The information content of the input signal is given, but the signal reception (information extraction) depends on the filtering process of the model.

Definition 3: Define *filtering* as the process of extracting information from the input signal.

A general agent model must be able to model learning, because an agent will frequently encounter offers of products unfamiliar to the agent - that is, where no knowledge of expected utility exists in the state variable, or offers unfamiliar to all agents, that is, innovations. Without learning, no consumer agent will ever buy something it has not bought before. Learning requires a filtering process that does not discard the input information when no knowledge of expected utility exists in the state variable.

B. Experience, Learning and Time

Akerlof's model stipulated that Group 2 traders would infer quality from a market price. As this state space model is dynamic, stipulate that the model "infers quality" from the most recent bid price. Then:

(10) $V_2(t+1)$ = p(t) if $\mathbf{y}_2(t)=\{buy,p(t)\}$
= $V_2(t)$ otherwise

Assume the Group 2 traders have some "experience" and believe that average automobile has a quality of one. Then $V_2(t_0)=1$. The model has new "knowledge" about expected utility.

Definition 4: Define *experience* as the value of those elements of the state vector that allow the agent to infer expected utility from the information in the input signal.

The output of the Group 2 traders model is $\mathbf{y}_2(t)=\mathbf{f}_2(t, \mathbf{x}_2(t),\mathbf{x}_2(t=t_0))= \mathbf{f}$(offer price,most recent bid price,income). However, because the offer price $p_1=2$ is greater than 1.5*expected utility then from (4) there is still no state change.

Akerlof's classical analysis technique specified the conditions for an equilibrium market price and showed that there was no solution. This state space analysis analyzes the dynamic behavior of the system, and shows that there is no change from initial conditions. In equilibrium (or *steady state*, defined as the state as $t \to \infty$) the system is in its initial state.

To model learning, assume Group 1 chooses to learn about Group 2 by reducing the price in decrements of 0.1 and observing the response. Then one element of Group 1 state vector $\mathbf{x}_1(t)$ will be (last offer price, response), where response $\in$ (no bid, bid). After five decrements the Group 2 output will be $\mathbf{y}_2(t_0+5)$=(bid,1.5), and Group 1 will sell the four lowest-quality cars. Group 2's utility will change by −3.

Proposition 2: There is a time value of knowledge. Alternately, there is value from faster learning.

Proof of Proposition 2: Let Group 3 be another group of traders that has cars of the same quality as Group 1. Group 3 offers cars to Group 2 independently of Group 1. Assume Group 3 is more aggressive and decrements price by 1/6 each time there is no response by Group 2. Then group 3 will sell the three lowest-quality cars at t_0+3, two periods earlier than Group 1. The value from faster learning will be the two-period time value of 3 if Group 2 continues to buy cars after buying the first three from Group 3. If Group 2 represents the entire market for cars and can only afford four cars then the value of faster learning will be 3.

Learning is central to economic activity in our dynamic model. In the state space model, the result of a "decision" to "learn" must be a change in the output $\mathbf{y}(t)$. But $\mathbf{y}(t)= \mathbf{f}(t,\mathbf{u}(t),\mathbf{x}(t))$. That is, for learning to occur there must be some elements of the state vector $\mathbf{x}(t)$ such that at some time $t=t_1$ and for some input $\mathbf{u}'(t)$, $\mathbf{y}(t_1)=\mathbf{f}(t_1,\mathbf{u}'(t_1),\mathbf{x}(t_1))$={learn}. We can think of the elements of the state vector that trigger learning as the "learning routine" or "learning algorithm" of the agent. This learning routine may be present at $t=t_0$ as part of the initial state $\mathbf{x}_0=\mathbf{x}(t_0)$.

C. Expertise and Rents from Knowledge

We can model expertise by introducing knowledge such as "if engine rattles, expected quality=0.5" into the state vector $\mathbf{x}_2$. Say Group 2 traders "query rattle of each car" with expected quality 1 with no rattle and 0.5 with rattle. Then if the Group 1 traders provide the rattle information (say the three lowest quality cars rattle) and decrement price as before, an "offer 0.5 with rattle" will result in two transactions – the cars with quality 0 and 0.5. The Group 2 traders net change in utility is –0.5, where in the "no expertise case" above the change was –3. The net gain of 2.5 represents a rent from knowledge.

Proposition 3: Knowledge is a potential source of economic rents.

Proof of Proposition 3: Let the cost of acquiring the knowledge "if engine rattles, expected quality =0.5" be c_K. For any one agent the supply of this knowledge is fixed – once the agent possesses this piece of knowledge acquiring "more of it" has no meaning at any cost. The cost c_K is only incurred once and must be finite, because no rational agent would pay an infinite amount. Let the value of the knowledge for one potential transaction be w_K. The total value to the agent can be arbitrarily large depending on the number of cars with rattles that the agent considers buying over its lifetime. There must be some number of potential transactions n for which $c_K < n * w_K$, so knowledge represents a source of economic rents.

D. Information, Knowledge and Physical Work

In section II.B, I introduced a term w(t), representing the agent's *energy* or *ability to do work*. By definition, an agent with w(t)=0 cannot interact with its environment. It will be useful to distinguish between different types of work.

Definition 3: An agent does *physical work* w_p when it changes the physical state of its environment.

For example, the agent can use money to do physical work by purchasing a car (where an exchange mechanism exists), or an agent can do physical work by stealing a car. A change in ownership represents physical work because "ownership" only has meaning if physical work must be done to change it. In economies with property rights, "ownership" means that the state will spend the physical resources to track down a car thief and punish him.

The concept of physical work is useful whether the agent is a "producer" or a "consumer". In either case, the agent is doing physical work to effect a change of state in its environment.

So far we have assumed that learning is costless for the buyer. Group 1 bore the cost of learning about Group 2 by decrementing price, and the "query rattle of each car" was assumed costless. In a general model, we must assume that everything is costly. Then the state vector of a utility-maximizing agent must also contain the expected utility of a query. Let the state vector $\mathbf{x}_2(t)$ be of the form {number of cars, income=M_2, expected value of car, expected value of query}. Then $\mathbf{f}_2(t)$ will have some term such that:

(11) $\mathbf{f}_2(t)$= while $M_2(t)$m0, if $p_1(t)$ [(3/2*V_2(t)-exp. util of query) then $\mathbf{y}_2(t)$=(query rattle)

In general, the agent will have to expend energy (do work) to acquire and process input information, to draw inferences (apply knowledge) and to cause a change in the physical state of the environment.

Definition 4: Define *knowledge work* w_k as work done by the agent that results in a change in those elements of the state vector that allow the agent to infer expected utility from the information in the input signal.

Definition 5: Define *information work* w_i as work done by the agent in processing input information and producing output information that does not result in a change of state in those elements of the state vector that allow the agent to infer expected utility from the information in the input signal.

Proposition 4: To maximize the expected utility of an action or transaction, an agent must optimize the allocation of work over information, knowledge and physical work activities.

Proof of Proposition 4: By definition, the total work w done in any agent action or transaction is

(12) $w=w_i+w_k+w_p$.

w represents the total cost of the action or transaction. Let the benefit of the action or transaction be *b*. Then the change in utility $\Delta U= b-w$. So to maximize utility the agent must minimize *w* – that is, optimize the allocation of work over information, knowledge and physical work activities.

E. Controllability and Observability

The information that the buyer can elicit from the seller depends on the state space analysis concepts of observability and controllability [Wiberg, 1971, page 128]:

> The state of a system is observable if knowledge of the input $\mathbf{u}(t)$ and output $\mathbf{y}(t)$ over a finite time segment $t_0 < \tau$ [t completely determines $\mathbf{x}(t)$.
>
> A state $\mathbf{x}(t)$ of a system is controllable if all initial conditions $\mathbf{x}_0$ at any previous time $\mathbf{t}_0$ can be transferred to $\mathbf{x}_1$ in a finite time by some control function $\mathbf{u}(t,\mathbf{x}_0)$.

A detailed examination of observability and controllability is unnecessary for the purposes of this essay. Observability and controllability of a system depend on the nature of the system. For the purposes of this example, it is sufficient to note that the result of the seller's learning will depend on the characteristics of the buyer.

IV. STATE SPACE DESCRIPTION OF GENERAL CASE

Consider a closed system and denote the agent as A and everything else as the environment B, where A and B are dynamical systems. In the general case no rules or exchange mechanisms are assumed.

The state of the environment $\mathbf{x_B}(t)$ is given by:

$$\mathbf{x_B}(t) = \mathbf{f_B}(t;t_0,\mathbf{x}_B(t_0),\mathbf{u_B}(t))$$

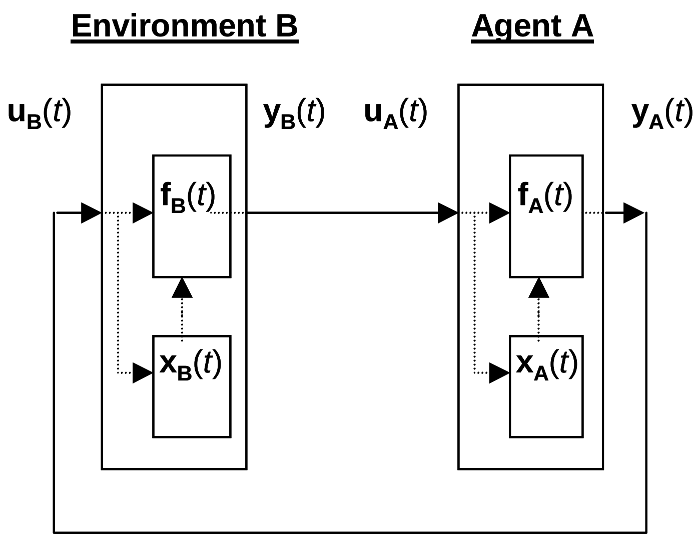

Figure 3. Interaction of an Agent and its Environment.

The agent A is of interest only if it can cause a change in state in the environment B. That is, the input of B must be some function of the output of A. Since we have a closed system, there are no other potential inputs to B and we can set $\mathbf{u_B}(t)= \mathbf{y_A}(t)$ without loss of generality. We can say that A "has the ability to act on" B. In general, the state of the environment $\mathbf{x_B}(t)$ is not directly observable. The output signal $\mathbf{y_B}(t)$ contains all the directly observable information about the environment B. If the agent A has no information about the environment, the result of A's acting on B will in general be unpredictable.

We are most interested in an agent A that has the potential to achieve some desired future state of B as a result of acting on B. Assume that A "can perceive information about" B. We can let $\mathbf{u_A}(t)=\mathbf{y_B}(t)$ without loss of generality. Let $\mathbf{x_B}(t_0)$ be the initial state of the environment B, $\mathbf{x_B}(t_1)$ be the current state of the environment and $\mathbf{x_B}(t_2)$ the agent's desired future state of the environment, $t_2 > t_1 > t_0$. Then the output of the agent $\mathbf{y_A}(t_1)$ is some function $\mathbf{f}_A(t)$ of the observable information about the environment $\mathbf{y_B}(t_1)$ and the state of the agent $\mathbf{x_A}(t_1)$, such that:

$$\mathbf{x_A}(t_2) = \mathbf{g}_A(t;t_1,\mathbf{x_A}(t_1),\mathbf{y_B}(t_1)), \text{ where}$$

$$\mathbf{y_B}(t_1) = \mathbf{f}_B(t, \mathbf{x_B}(t_0),\mathbf{y_A}(t_0))$$

Represent the energy of a state $\mathbf{x_B}$ as $E(\mathbf{x_B})$. By definition, in the physical world a change in state requires work, or energy. So $E(\mathbf{x_B}(t_2))$ g $E(\mathbf{x_B}(t_1))$. The physical work done by the agent is $w_p= E(\mathbf{x_B}(t_2))-E(\mathbf{x_B}(t_1))$. We should also assume that "receiving" and "transmitting" information is costly – denote this cost w_i. In the physical world, energy must be one element in the current state vectors $\mathbf{x_A}$ and $\mathbf{x_B}$. Let $\mathbf{x_A}(t)=\{\mathbf{k_A}(t),w_A(t)\}$. $w_A(t)$ is the budget constraint, expressed as the ability of the agent to do work (we express the budget constraint as "work" w(t) rather than income m(t), but the terms are interchangeable). Then $\mathbf{x_A}(t_2)=\{\mathbf{k_A}(t_1),w_A(t_1)-w_p-w_i\}$. Define $\mathbf{k_A}(t)$ as the "knowledge of the agent A at time t". Then knowledge $\mathbf{k_A}(t)$ and energy $w_A(t)$ parameterize the behavior of agent A. That is, the agent's output is determined not only by the input but by its current state of knowledge and ability to do work (loosely, preferences and budget constraint). The state space of A is the set of all possible $\{\mathbf{k_A}(t), w_A(t)\}$.

So far, we have assumed only that A can act on B and that A can perceive information about B. For generality, we should also assume that A can "learn", at a cost of w_k. That is, $\mathbf{k_A}(t_2)$ g $\mathbf{k_A}(t_1)$. Then $\mathbf{x_A}(t_2)=\{\mathbf{k_A}(t_2),w_A(t_1)-w_p-w_i-w_k\}$. The output term $\mathbf{y}(t)$ is comprised of information and work. $\mathbf{k_A}(t)$ represents "knowledge", in the sense of "conditional probabilistic beliefs" about the relationship of the input signal to marginal utility (see Spence [1973]). Note, however, that in this general case there is no assumption of equilibrium and these beliefs need not be "correct".

The agent's optimization problem is to achieve the desired change in state at minimum cost. This requires minimizing $w_p+w_i+w_k$.

V. Assumptions in the Neoclassical Consumer Model

The neoclassical model is really a special case that contains many hidden assumptions about information and knowledge that can be revealed using a state space model. I will illustrate this with the case of an street vendor selling apples. I will use "information" to mean inputs to the state space model, "knowledge" to refer to elements of the state vector and "work" to refer to physical work (for example, expenditure). In these terms, a model of the street vendor is rich with information and knowledge:

Information: shiny round red objects are sitting on a bench in front of me, and a placard has Apples: $1 written on it.

Knowledge: those shiny round red things are apples; I can eat apples; eating an apple will provide the taste sensations of sweetness and juiciness and some reduction in hunger; I would feel better after these sensations; the vendor is willing to transfer ownership of the apple to me for $1; transfer of ownership means I can do with the apple whatever I want; these sensations are worth more to me than $1; simply taking the apple could result in my arrest, a fine and criminal conviction.

Work: the work required to obtain the apple is $1 plus the effort of the transaction.

There are many assumptions that may not in general be valid:

Information: this model assumes the agent can acquire adequate information. However, the agent may be blind, or illiterate, or unable to read English, or simply looking the wrong way. *Filtering* affects the outcome of the transaction.

Knowledge: this scenario assumes the agent can recognize an apple, knows the utility of apples in general, knows the utility of these apples in particular (quality), and knows that purchasing an apple has higher utility than simply taking one. Possible exceptions might occur if the agent has never seen an apple before, or knows that apples are out of season and therefore that the utility of these apples is likely to be lower than average, or knows that he will never be in town again, and can run faster than the vendor, and so decides to steal the apple. The agent's *knowledge* affects the outcome of the transaction.

Physical work: the intrinsic assumption that the agent has the required budget is included in the neoclassical model, because it takes the budget constraint into account.

Seller agent actions could significantly alter the transaction. Information work might include shouting "Apples: $1", perhaps in multiple languages, thus increasing the number of people who have information about the offer. Knowledge work might include educating potential buyers by providing a taste sample.

The neoclassical assumptions - that utility is perfectly and costlessly observable and that a perfect market exchange mechanism exists - obscure how information filtering and knowledge can affect outcomes.

A. Analytical Relationship of State Space Model to Neoclassical model

In the neoclassical model, preferences are fixed (for producers, technology is fixed). Marginal utility is not a function of the current state but of quantity offered. To model this with a state space model every element of the state variable except for the budget must remain constant.

In the language of the state space model of an economic agent, preferences relate input information to expected utility. Because there is no learning, to replicate the neoclassical model all preferences must exist in the initial state vector. In the state space model, learning means that the relationship between input information and expected utility can change. That is, "preferences" can change.

	Neoclassical Model	State Space Model
Information	Information is a public good	Perceived information depends on the "filtering" of the agent
Preferences	Preferences are fixed	"Revealed preferences" depend on knowledge, which changes due to learning, expertise etc.
Marginal Utility	Decreasing marginal utility as a function of quantity offered	Changing marginal utility as a function of current state (e.g. quantity consumed over time)

Table 1. Comparison of Neoclassical and State Space Models.

VI. SALES AND MARKETING AS LEARNING

One of the major benefits of a state space approach is the ability to model sales and marketing. In the neoclassical model, all necessary information and knowledge is present, and the agent need only choose. Sales and marketing activities are unnecessary.

Using the state space approach, sales and marketing can be modeled as information and knowledge work designed to achieve a desired change in a buyer's state. There are two actors in our automobile model (I will generalize later)– the seller and the buyer, so there are four possible ways the buyer can learn:

- "Buyer information search": The buyer can cause its own input $\mathbf{u_2}$(t) to change; that is, search for more information.
- "Buyer knowledge acquisition": The buyer can cause the state $\mathbf{x_2}$(t) to change.
- "Seller information transfer": The seller can cause the buyer's input $\mathbf{u_2}$(t) to change by changing $\mathbf{y_1}$(t).
- "Seller knowledge transfer": The seller can cause the buyer's state $\mathbf{x_2}$(t) to change by changing $\mathbf{y_1}$(t).

A. "Buyer information search"

In this case, the buyer applies a stimulus to the seller state space model, observes the response and makes inferences. For example, the stimulus can include questions, negotiation, offers intended to "separate" seller types and so on. There is knowledge involved in searching efficiently (for example, hypothesis testing). This important issue of *learning efficiency* will be explored later in this essay. The inferences drawn depend on **x**(t), or "current state knowledge": different inferences may be drawn in different states. For example, "trust"– if x_i ={Bob has never been wrong before}, then the agent's response to input from Bob may differ from the response to a stranger. In this model, buyer information search conveys information to the seller and may elicit a response (a change the terms of the offer, for example). *Passive search* conveys no information to the seller. *Active search* may be faster but conveys information to the seller (for example, queries about a target company's valuation may drive up the stock in anticipation of a takeover bid).

More generally, the buyer applies a stimulus to the environment. This can entail simply going to the supermarket, where the buyer searches for an item, or the seller signals, hoping to catch the buyer's eye. Sales "channels" such as distributors and retailers can reduce the buyer's search cost. The issue of who has the power to control what the customer sees is an important one in retailing.

B. "Buyer knowledge acquisition"

If the buyer input $\mathbf{u_2}$(t) is constant, the buyer can change state (that is, change beliefs or get more knowledge) by buying the good ("learning by experience") or from research. As with all learning strategies, there are techniques to minimize the cost to the buyer, such as asking for a trial period, or learning from other customers' experience.

C. "Seller information transfer"

In most cases, sellers will "advertise", for example $\mathbf{y_1}$(t)={for sale, good X,$p}. Typical strategies include mass media advertising, direct mail, in store promotions and product demonstrations. "Seller information transfer" also includes searching on behalf of the buyer – that is, helping a buyer find an item, providing feedback/reinforcement to the buyer ("that looks great on you"), describing the product in flattering terms and signaling – for example, by hiring attractive people to wear the product to influence the buyer's perception of input data.

Sellers can add information to their output signal to help buyers estimate expected utility. Spence [1973] examined the equilibrium case, and demonstrated that to convey information about utility in equilibrium, the cost of the signal should be negatively correlated with expected utility. In our more general dynamic model, the response of the buyer to seller information transfer will depend on the buyer's state $\mathbf{x}(t)$. Put another way, the efficacy of signaling depends on the buyer's knowledge.

Techniques to minimize the cost of seller information transfer may include branding.

D. "Seller knowledge transfer"

The seller can attempt to change the buyer's state by discounting, free samples, free trial period, or a "returns" policy. They can also change beliefs through persuasion, such as product endorsements.

E. Sales Techniques

Salespeople also use a combination of these techniques. For example, they provide stimuli to elicit responses, hence learning about the buyer's current state and changing offers accordingly.

F. Marketing

So far, I have examined buyer learning, but all of these techniques apply to the seller. The business term "marketing" refers to the set of seller learning strategies:

- "Seller information transfer": The seller can cause the buyer's input $\mathbf{u}_2(t)$ to change by changing $\mathbf{y}_1(t)$
- "Seller knowledge transfer": The seller can cause the buyer's state $\mathbf{x}_2(t)$ to change by changing $\mathbf{y}_1(t)$
- "Seller information search": The seller can cause its own input $\mathbf{u}_1(t)$ to change; that is, search for more information about the buyer
- "Seller knowledge acquisition": The seller can cause its own state $\mathbf{x}_1(t)$ to change; that is, learn about the buyer

G. General Case

So far I have assumed that the seller and buyer are "directly connected" – that is, participating in a market that transmits information efficiently and where no information is lost. In general, imperfect connections, filtering and knowledge gaps will affect the outcome of buyer and seller actions. The range of possible outcomes in the general case are shown in Tables 2 and 3:

	Change in..			
Scenario	U2	X2	Y2	EXPLANATION
1	N	N	N	Buyer not connected, does not “see” anything
2	Y	N	N	Buyer filters out the input information
3	Y	Y	N	Buyer learns something but does not act on it
4	Y	N	Y	Buyer acts on new information
5	Y	Y	Y	Buyer learns and acts

Table 2. Possible Buyer Responses to a Change in Seller Output.

Similarly, the Seller state space model can react in various ways to the Buyer’s response:

	Change in..			
Scenario	U1	X1	Y1	EXPLANATION
1	N	N	N	Seller not connected, does not “see” anything
2	N	Y	N	Seller learns from lack of perceived response but does not act on it
3	N	Y	Y	Seller learns from lack of perceived response and acts on it
4	Y	N	N	Seller filters out the response information
5	Y	Y	N	Seller learns from the response but does not act on it
6	Y	N	Y	Seller responds to new information with existing knowledge
7	Y	Y	Y	Seller learns and acts

Table 3. Possible Seller Responses to Buyer’s Response.

H. Sales and Marketing Costs

Sales and marketing costs can be large, particularly when neither party knows expected utility. When the good is an input (say, a tool to be used in a home improvement project) and the customer has imperfect knowledge, neither the seller nor the buyer knows the utility of the good. The seller does not know what ultimate utility the buyer desires. This is a classic problem in business – “if you can tell me exactly what you’re trying to achieve, I can tell you what you need”. In many cases, there is no way of knowing the buyers utility, hence the need for “experience”. The cost of experience is reduced by returns policies of sellers.

VII. CONTRACTS

Using state space analysis, a contract is simply another way to cause a desired output from a seller’s state space model. Instead of the buyer agent applying input information to a seller and hoping for the desired output, the buyer specifies the desired output. Where a product does not exist to meet a buyer’s need, a contract between buyer and seller to specify a desired output may be costly for both but increase utility of both. The alternative is for producers/sellers to bear the cost of learning the buyer’s needs. The cost of the contracting process must be less than the cost of other learning strategies. In the case of unique products such as office buildings, “trial and error” learning is prohibitively expensive. Use of contracts reduces learning costs enough that products are produced that would not otherwise be cost-effective.

VIII. IMPACT OF INFORMATION AND KNOWLEDGE ON PRODUCTION

The state space model applies to producers as well as consumers. Although the definition of "input as information" seems more restrictive that that of the usual production function – where inputs are physical items such as steel and labor - there is no loss of generality. Instead of physical inputs, the input is of the form of information about the current state of the physical inputs. For example, instead of a roll of sheet steel being considered an "input" of production for some producer B, it is necessary only that one element of the state vector describing the steel be x_i=(owned by B).

In general, an agent might choose to produce rather than purchase. This is unlikely for cars but certainly feasible for many services, such as home repairs. The "build or buy" decision will be determined by information and knowledge costs.

Consider an example of installing track lighting. The buyer has information in his current state vector, such as the location of the current power supply and the desired orientation of lighting. This information might be costly for an electrician to acquire – he might have to visit the client's house. On the other hand, the electrician has knowledge (expertise) that it may be costly for the client to acquire. The buyer has to do less physical work in terms of travel time, but the electrician may be more efficient due to his expertise. The opportunity cost of the client's labor may be higher than the cost of the electrician during the week, but less on the weekend. The build/buy decision will be based on an assessment of the relative information, knowledge and physical work costs of the build and buy options. Thus, the state space model updates a concept of "production" designed around manufacturing rather than the service economy.

IX. LEARNING – GENERAL CASE

If our agent's objective is to maximize expected utility, it must be able to estimate the expected utility of an investment in learning. This is true both for consumers and for producers – that is, innovators. Using the example from Spence [1973] to illustrate, the utility of a decision to hire a graduate will be the expected marginal product of the graduate *plus the expected utility of learning about the "graduate" signal*. In Spence's equilibrium analysis there is no learning and hence no expected utility from learning. This is *not* true in general – it is a special case due to the equilibrium assumption. That is, *every expected utility calculation should take into account the expected returns from learning*. Alternatively, the opportunity cost to not learning should be considered. Learning involving state change will provide a stream of returns over time. In general, there will be rents attributable to learning. These rents will vanish only in the equilibrium case.

A. Learning to Learn

Optimizing the expected utility of learning is itself a learning process. That is, the agent must "learn how to learn". For example, an optimal agent will respond to stimuli, but the size and speed of the response will vary based on how the agent learns. Fast response may lead to instability (over-reaction or "overshoot"). Slow response may reduce returns from learning.

It is possible that agents under- or over-invest in learning. Because learning involves an investment of time as well as money, there is an added complexity to estimating the utility of learning. The opportunity cost of learning resource may be difficult to establish. It may also be difficult to evaluate the cost of delay while learning takes place. Further, the utility of learning depends on many factors, including the appropriability of rents. Because of the complexity of "investment in learning" decision it is possible that many managers "go with the gut".

X. INTERACTION OF INFORMATION, KNOWLEDGE, WORK AND TIME

I have been careful throughout this paper to speak in terms of "work" instead of "change of ownership". I use "work" to mean "ability to act" - not just income but the ability to change physical state, including the ability to secure access to scare resources.

Information, knowledge and work interact in subtle ways. For example, rents from ownership of scarce resources depend on knowledge. In an efficient market for scarce resources, the price should reflect the discounted present value of all future income. However, ownership of a physical asset entitles the owner to the stream of future rents that might accrue due to new knowledge. These rents are named for Edith Penrose [1995]. The power of a dictator can secure rents even from new oil discoveries made with new technology. Every purchase of a physical asset should contain an option value for utility not yet known, yet this value cannot be determined in an "efficient" market as information and knowledge are imperfect, by definition. Conversely, new knowledge can render a resource worthless.

Another perspective on the fungibility of information, knowledge and work is that in principle, every agent faces a "buy or build" decision. An economic agent can secure any resource given information on its location, knowledge of how to secure it and the ability to do the necessary physical work. This is particularly true of services, where there are trade-offs in terms of gathering information and physically transporting either the customer or the service provider.

Because in principle any resource can be secured with information, knowledge and physical work, it is only when we introduce the concept of ownership that a market mechanism is needed. Ownership implies that it is prohibitively costly to acquire resources without the owners consent. Property rights must exist in part because it is more efficient for the state to do the physical work to protect assets than for the individual agents to defend their assets.

Inventories are a good example of the interaction of information, knowledge and physical work. Inventories are the physical manifestation of imperfect information. Better information on demand can reduce inventory requirements, a fact often used to justify investment in Internet-based supply chain technology.

Information, knowledge and physical work costs are not independent. Information acquisition costs are a function of knowledge – for example, "knowing what to look for". Superior knowledge can reduce the physical work required to achieve a desired state change.

Finally, because information and knowledge can be sources of economic rents, their value depends critically on time.

XI. IMPLICATIONS AND APPLICATIONS

A. Filtering

Information is filtered based on current knowledge, so information that can challenge existing knowledge may be treated as "noise" and filtered out. There is much anecdotal evidence that people tend to "filter out" things they don't know or don't want to hear – "there's none so blind as those that won't see". The search process is often thought of as active, but filtering is a passive search process. The Internet has reduced the cost of connecting with potential customers, and hence shifted the cost of filtering to the individual. Filtering may be one reason why new product introductions are so expensive.

B. Endowments

The dependence of output on initial state may explain the value of endowments and the persistence of class advantages. For example, the success of stockbrokers depends on access to affluent clients. Affluent parents are more likely to send their children to prep schools and private universities, where they make contacts with the children of other affluent individuals. If information acquisition cost is high (that is, it is costly to search for people willing to give you their money to invest), then the initial endowment of contacts is valuable and represents a source of rents.

C. Competing through Learning

There is a time value of knowledge. If we allow the agent to have "memory" – that is, retain previous states – an agent's ability to learn may depend on how much it has learned in the past. That is, the returns to learning may be cumulative and represent a potential source of rents.

D. Business Strategy

The objective of business strategy should be to maximize rents. There are three factors of production – information, knowledge and physical work – and two sides to business transactions – buyer and seller. Then there are six generic strategies:

Seller information advantage

Seller earns rents from information not available to competitors - for example, credit card companies' information on customers' purchasing patterns, stockbrokers' pre-existing relationships with a potential investors.

Seller knowledge advantage

Seller earns rents from proprietary knowledge, such as patents, proprietary technology, and unique knowledge of customer needs due to preferred access.

Seller work advantage

A seller earns rents from ownership of a scarce resource such as real estate or an oil deposit.

Buyer information advantage

Seller earns rents because a buyer has better access to information from that seller than from others. For example, fund management companies compete fiercely for workplace access to investors.

Buyer knowledge advantage

Seller earns rents because a buyer has better knowledge about the utility of his product than competing products. Software is one example - buyers tend to buy new versions of products with which they are familiar rather than entirely new products.

Buyer work advantage

Seller earns rents because the buyer has to do less physical work with that seller than with others – for example, due to proximity.

Sources of advantage

All of these advantages must be derived either from a seller's initial state or through better learning than competitors'. For example, in an efficient market the lease of an advantaged store location would be priced to reflect fully all potential future revenues, and no economic rents could exist. Rents could be earned only if the buyer's initial state or learning had produced a "knowledge asymmetry".

The economics of learning can be an important part of competitive strategy. Economies of scale and network effects in information acquisition can provide rents. Because monopolies on information and knowledge are usually temporary, lock-in is key to appropriating rents.

E. Sales

In Section VI, I modeled sales as learning. Two examples of this are technical sales and solution selling. In technical sales, salespeople and "Application Engineers" are a major channel for dissemination of ideas. Vendors typically bear the cost of educating the customers – through presentations, training and free trials.

Solution selling is an important feature of modern business. In the past, companies used to sell components. For example, information technology companies sold mainframes, storage, monitors, keyboards, operating system software and application software. However, buyers want utility – for example, an effective human resources information system - not a bundle of components. In IT, there is a lot of work involved in getting from input to utility. Often the relationship between the input information (features of individual products) and utility is not clear. The cost of the learning process in making a purchasing decision increases with the number of suppliers considered. Knowledge is costly to acquire, and in IT people are typically trained on specific software and hardware. The incremental cost of acquiring knowledge for new products can be much greater than learning about upgrades to familiar products.

Computer companies have now invested heavily in providing "solutions". Solution selling economizes on the customer's learning cost. Because learning has a large fixed cost component, a supplier's marginal cost in learning the customer's needs may be much less than the customer's cost to learn about the supplier's technology

Learning is an integral part of the transaction for the many cases where the utility of a good is the result of some production process. In some cases, there is no way of knowing the utility of a good in advance because the buyer's utility is subjective. A person shopping for clothing may not want a "red dress", but something that "looks great". The seller often bears this cost, in the form of a returns policy. Interior decorators are familiar with the need for many iterations to satisfy their clients. There is a real cost to the seller when the buyer's knowledge is not perfect. The buyer will pay less for the good due to the uncertainty associated with the ultimate utility.

F. Marketing

The neoclassical model makes no allowance for "persuasion". Marketing has been discussed as a way of reducing search costs and as a "signal" of quality. In our model, marketing can also change knowledge, or preferences. Consumer marketing can persuade people that an item is "cool", or conveys sophistication or provides some other emotional benefit. In the industrial or business market, marketers can educate the customer to the utility of the product. The state space model of economic agents seems to provide a much more complete explanation for the importance of marketing in the economy.

G. Organization Design

A firm can be represented as a state space model, like any other economic agent. (A more realistic representation would be that of the firm as a collection of agents, with the state variable of the firm $\mathbf{X}(t)$ being a function of the input and the current states of all of the agents employed by the firm. This raises the interesting question of the circumstances in which the firm's agents might be "individually smart" but the resulting firm's behavior "collectively dumb". However, for my current purposes a single model of the firm will suffice). Then some of the primary activities of the firm are to obtain information about the environment, apply knowledge to filter the information, apply knowledge to decide what to do with the filtered information, and do work to implement the decision. The question of organizational design is how to allocate resources to information, knowledge and physical work activities to maximize profits. It turns out that the organizational structure must depend on the environment.

Consider an example from the newspaper business. In the news business the Associated Press and Reuters exist because of economies of scope in capturing information. An Australian newspaper may not be able to justify a full-time correspondent in Vienna because of the low frequency of news of interest to the Australian audience. The editor of the Australian newspaper uses her knowledge of her market to filter the input information from the news service and to infer the utility of newswire stories to her audience. However, in so doing the newspaper editor foregoes the rents that might be earned from an "exclusive" on some story in Vienna that happens to be of great interest to Australians. She has no ability to obtain new information to learn about some developing story in Vienna – she has delegated "search" to the newswire correspondent. That organizational decision is based on judgment (knowledge) about the likelihood and size of temporary rents originating from Viennese news. That knowledge is based on her model of the behavior of the Viennese news environment.

In a firm, the rents from an "exclusive" may be able to be locked in – for example, the licensing of some innovative product. There is a continuing conflict between minimizing the costs of information and maximizing rents. Unless the state vector of the environment is observable, there will be no perfect solution because, by definition, the information output of the environment will be unpredictable.

H. The Internet

The Internet was expected to destroy intermediaries. Filtering and learning represent two possible reasons this didn't happen.

Internet technology reduces the cost of information transfer. Because the cost is reduced, volume increases and so do the resources needed to filter the information. The natural filtering process of high cost of information transmission was lost.

Information is not utility. There is a learning process required to acquire the knowledge necessary to infer utility from information. Put another way, easy access to information is only valuable if consumers "know what they need to know". Retail environments and salespeople provide valuable learning. Feedback from salespeople, seeing what other people are buying, seeing what is displayed prominently, browsing to "see what appeals" are all examples of customers learning about their own utility. Because of the importance of learning, the Internet is most valuable for lowering the transaction costs of repeat purchases, or where the buyer has enough knowledge to infer utility from information with a high degree of confidence.

XIII. CONCLUSIONS

The state space representation seems a natural way to model economic behavior. A person can gather information with her eyes, ears and nose, create and store knowledge in her brain, and do physical work directed by the information and her knowledge to achieve a desired result. For a firm, a state space representation reflects the "business process" model: Input, Processing, Output. Further, a state space model seems useful in explaining the prevalence of sales and marketing in a modern economy.

A state space model of an economic agent has many advantages. No *a priori* assumptions are necessary. The agent need not be rational or utility maximizing. Information need not be perfect or costless to acquire or to transmit. Preferences can change. Learning is possible. Further, a state space model allows the tools of state space analysis to be used to analyze important characteristics of the system, such as controllability, observability and stability.

The state space approach may have important business applications. Competitive advantage must derive from superior information, superior knowledge and/or superior ability to do work. Learning may be a significant competitive strategy, and tactics to acquire and protect temporary information and knowledge monopolies may be important. Organizational design may be strongly influenced by the need process information, learn, apply knowledge or do physical work.

There are many avenues for further work: optimal learning models, the economics of learning (for example, scale and network effects in search), stability, controllability and observability of economic systems, computer simulation of agent interactions and transient phenomena.

The most important benefit of the state space approach may be in placing information, knowledge, innovation and learning at the center of economic analysis.

REFERENCES

Akerlof, George A., "The Market for "Lemons": Quality, Uncertainty and the Market Mechanism", *The Quarterly Journal of Economics*, Volume 84, Issue 3 (August, 1970), 488-500

Fudenberg, Drew and Levine, David K., “Steady State Learning and Nash Equilibrium”, *Econometrica*, Volume 61, Issue 3 (May 1993), 547-573

Kahneman, Daniel and Tversky, Amos, “Prospect Theory: An analysis of decision under risk”, *Econometrica*, Volume 47 (1979), 263-291

Penrose, Edith, *The Theory of the Growth of the Firm*, Third Edition (New York, Oxford University Press, 1995)

Samuelson, Paul a. and Nordhaus, William D., *Economics*, Twelfth Edition (McGraw Hill, 1985)

Spence, Michael. “Job Market Signaling”, *The Quarterly Journal of Economics*, Volume 87, Issue 3 (August 1973), 355-374

Stigler, George J., “The Economics of Information”, *The Journal of Political Economy*, Volume 69, Issue 3 (June 1961), 213-225

Stiglitz, Joseph E. “The Contributions of the Economics of Information to Twentieth Century Economics”, *The Quarterly Journal of Economics*, November 2000, 1441-1478

Wiberg, Donald M., *State Space and Linear Systems* (McGraw Hill, 1971)

Williamson, Oliver E. “The Lens of Contract: Private Ordering”, *The American Economic Review*, May 2002, 438-443

*** Percipient Consulting**

[1] The purpose of this paper is to illustrate the utility of state space analysis in modeling the behavior of economic agents. For clarity and brevity, I will use an informal approach to notation and omit formal definitions. A more rigorous treatment is available in mathematics and signal processing texts, for example Wiberg (1971)

[2] It can be helpful to think of a computer, where the output is determined not just by the input and some algorithm but also its current state that parameterizes the algorithm

[3] In this paper I will use the terms "information" and "signal" differently – and more precisely – than the information economics literature. In general, the usage is drawn from the signal processing literature.

[4] The information economics literature does not explicitly define "information". It is defined implicitly as "information about utility", in the context of equilibrium analysis where *information about expected utility is asymmetric.* In this paper, agents with different information simply have different inputs.