Learning and Discovery

Boyan Jovanovic*

September 30, 2009

Abstract

Bayesian inference is a process of eliminating parameter values that do not explain the data and shifting posteriors towards values that work. There is no notion in it of 'discovery'. It seems natural to look for a generalization of the Bayesian method that would allow the support of beliefs to increase over time. I propose such a model here.

PRELIMINARY

1 Introduction

Learning and discovery appear to be two qualitatively different ways in which beliefs change. By 'learning' we shall mean the usual form of statistical updating of beliefs in light of evidence. By a 'discovery' we shall mean a change in beliefs such that the posterior puts positive mass on something on which the prior had a zero mass. Bayes learning is learning from experience alone. Such 'experience learning' leaves room for 'insight' which brings about a new hypothesis never before conceived of.

Decision theory distinguishes events of which an agent is unaware from events that he believes to have zero probability. Dekel, Lipman and Rusticchini (1998), Li (2008), Galanis (2008) and Ozbay (2008, Sec. 4) discuss this point and a string of related papers. Given the rules of probability or, rather, given the form of Bayes rule, a growth of awareness cannot fit into Bayesian inference – a parameter value that is not in the support of the prior cannot be in the support of the posterior. Sims (1971) and Radner (2002) advocated certain rules of thumb about model revision which is a concept close to the treatment of discovery here.

*I thank L. Blume, T. Cogley, S. Galanis, A. Kochov, M. Kredler, R. Lucas, Y. Nyarko, E. Ozbay, R. Radner, J. Schlossberger, J. Stoye, H. Tretvoll and A. Tsyvinski for comments and the Kauffman Foundation and the NSF for support.

Applied work also routinely deals with environments in which there discoveries such as inventions; it typically assumes that agents form beliefs over payoffs, costs or productivity. Models of technological improvement arising from research such as Telser (1982), Muth (1986) and Kortum (1997) proceed in this way, as do research-driven-growth models of Romer (1990) and Aghion and Howitt (1992). Empirical work on patents on a firm's stock price value by Pakes (1985) assumed that the firm invents something; the content of that invention does not matter, only the amount by which it lowers production costs and raises the firm's value. In other studies summarized by Griliches (2000), basic research is held as having a higher probability than applied research of producing an outcome in the right tail of the payoff distribution.

All this relates to the issue of the "directedness" of discovery. Sims (1971) and Radner (2002) offer non-Bayesian treatments of model revision, and under full awareness Rothschild (1974), Jovanovic and Rob (1990) and Jovanovic and Nyarko (1996) used information theory to study directed search.

The paper proposes a model of discovery by which is meant simply that the support of beliefs increases over time, contrary to Bayesian inference. Examples show that it sometimes yields results that seem quite standard, and sometimes not. Bayes decisions arise a special case of the discovery model when A is fixed. The (exogenously-timed) introduction of a new parameter value into the prior leads to a re-evaluation of the evidence and a possibly a dramatic shift in beliefs. Young (02, Sec. 8.3) studies how agents would periodically switch models whenever they fail a statistical test. Kocherlakota (2007) has a related discussion. The model relates to other learning algorithms modeled by Venezia (1985) and Auerswald *et al.* (2000) and to a growth literature on paradigm shift in Bramoullé and Saint Paul (2007).

2 Two examples

This section presents two examples that illustrate why discovery has different implications for observables than Bayes learning under full awareness, but that also indicate how the models can possibly be made equivalent, at least quantitatively.

2.1 Investments with Binomial outcomes

A gamble pays a dollar each time a coin comes up heads, and nothing if the coin comes up tails. Let θ be the probability of heads, and let θ be zero, so that the coin always comes up tails. Let $y_t \in \{0, 1\}$ be the gross payoff to the gamble, so that $y_t = 0$ is observed by the agents for all t. We shall ask how this evidence, tails for ever, will affect the evolution of beliefs of two risk-neutral Bayesian agents, Agent 1 and Agent 2.

Agent 1.—starts out believing dogmatically that the coin is fair, i.e., that $\theta = 0.5$. At date T, the agent discovers that θ could be zero. He never conceived of such a possibility before date T but, having discovered it, he considers it as likely *a priori* as the possibility that $\theta = 1/2$. Absent any evidence, that is, he assigns equal probabilities over the two events. Then his prior is 0.5 on $\theta = 1/2$ and 0.5 on $\theta = 0$.

With these new priors and the evidence that the coin came up tails T periods in a row, his expected gross payoff, $E(y_t)$, falls suddenly from 0.5 for $t < T$ to $\frac{1}{1+2^t}$ for $t \geq T$:

$$E(y_t) = \begin{cases} \frac{1}{2} & \text{for } t < T \\ \frac{1}{1+2^t} & \text{for } t \geq T, \end{cases} .$$

Agent 2.—Now consider agent 2 who, from the outset, is aware of the possibility that $\theta = 0$, who has a prior of μ that $\theta = 0$ and $1 - \mu$ that $\theta = 1/2$. His posterior belief over $\theta = 0$ is

$$\pi(\theta = 0 \mid t \text{ tails and no heads}) = \frac{\mu}{\mu + (1-\mu)\left(\frac{1}{2}\right)^t} = \frac{\mu 2^t}{\mu 2^t + 1 - \mu}$$

Agent 2's expected gross payoff is

$$E^*(y_t) = \frac{1}{2}\left[1 - \frac{\mu 2^t}{\mu 2^t + 1 - \mu}\right] = \frac{1}{2}\frac{1-\mu}{\mu 2^t + 1 - \mu}$$

In Figure 1 we plot $E^*(y_t)$ for six values of $\mu = \{10^{-1}, 10^{-2}, ..., 10^{-6}\}$. We find that even as μ becomes quite small and the possibility of the coin being biased becomes quite remote, the model cannot generate the sudden drop in expected values. The drop takes place mostly over about 10 periods

Value of firm.—If the business was public and if the public had the same beliefs, the market value of that business would experience a sudden crash in the first model, and would take at least 10 periods to do so in the second. On the other hand if one is interested in the behavior of the price of a stock of such a company and is willing to tweak the model and assume that some periods produce a larger number of signals, then a faster market crash is possible in the Bayesian model. For instance, if in Figure 1 we take the right-most curve corresponding to $\mu = 10^{-6}$, , we could assume that at $t = 17$, the agent receives not one signal but ten. This would produce a rapid crash of the stock price, but would interfere on observations at the frequency at which trials actually took place. Calculating value, however, depends on expectations about future discoveries and if discoveries depend on the act of gambling, then gambling would have an option value. The firm's value would, nevertheless, drop precipitously.

Actions.—Suppose that taking the gamble costs 25¢. This can be thought of as the cost of investing in the project and the up-front costs of hiring the factors or production. Then until date T, firm 1 would keep losing money and then after date T investment would cease. For firm 2, investment would cease when E^* drops below 0.25. If we choose μ so that we fit actions, output and profits, clearly the Bayesian model with full awareness is unable to match the stock-price drop.

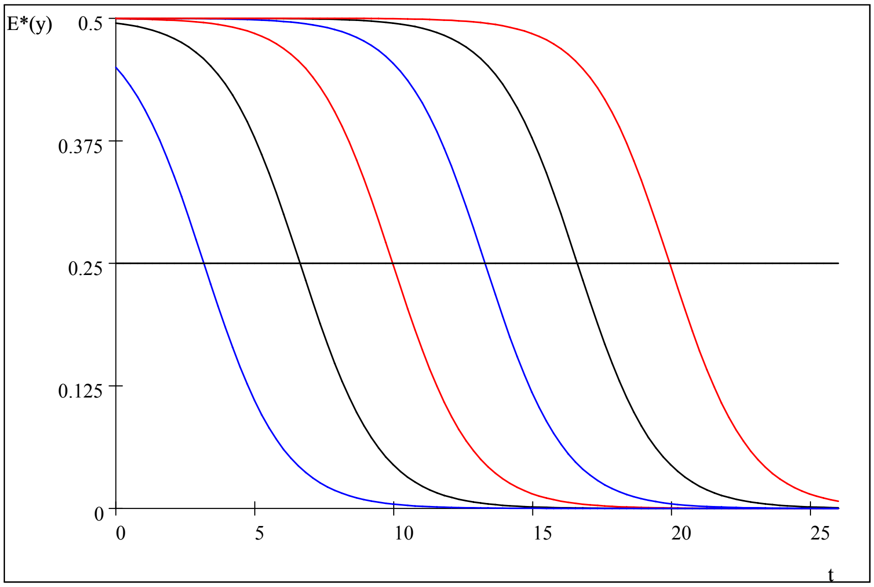

Figure 1: Bayesian learning: Binomial case

2.2 Investments with Normal outcomes

We now present the same two risk-neutral investors with a sequence of normally-distributed gambles which pay

$$y_t = \theta + \varepsilon_t \tag{1}$$

for $t = 0, 1, ...,$ where $\varepsilon \sim N(0, \sigma^2)$. To correspond to the above case, suppose that the true $\theta = 0$. We again ask evidence generated by (1), tails for ever, will affect the evolution of beliefs of two Bayesian agents, Agent 1 and Agent 2.

Agent 1.—The first agent again believes dogmatically that $\theta = 0.5$, and again at date T he discovers that θ could be zero, and assigns equal prior probabilities over $\theta = 0$ and $\theta = 0.5$. Of course his posterior beliefs will suddenly change at this point. Let $\bar{y}_t = \frac{1}{t} \sum_{s=0}^{t-1} y_s$. In contrast to the Binomial case, conditional on the true θ, the evidence, $\bar{y}_t \sim N\left(0, \frac{1}{t}\sigma^2\right)$ is a random variable. There must always be a positive drop at the date T, but the size of the drop of Agent 1's expectations of y_t at date T is then also a random variable, shifting from 0.5 to

$$E(y_T) = \frac{1}{2} \frac{1}{1 + \exp\left\{-\frac{T}{2\sigma^2}\left[\bar{y}_T^2 - \left(\bar{y}_T - \frac{1}{2}\right)^2\right]\right\}}$$

At the median, $\bar{y}_t = 0$, and so the median post-discovery expected payoff is

$$E_{\text{MED}}(y_T) = \frac{1}{2} \frac{1}{1 + \exp\left(\frac{T}{8\sigma^2}\right)}.$$

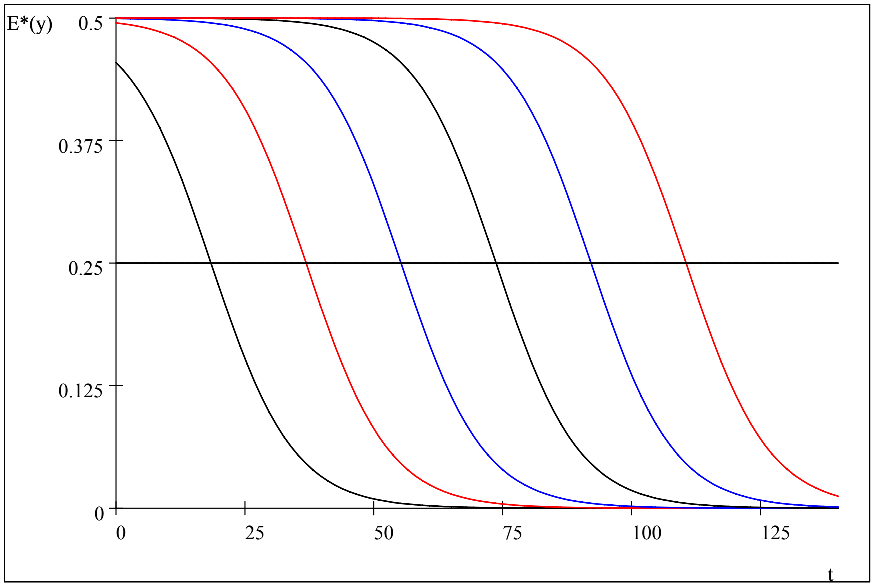

Figure 2: Bayesian learning: Normal case

Agent 2.—From the outset this agent assigns probability μ that $\theta = 0$, and $1 - \mu$ that $\theta = 1/2$. His date-t posterior expectation is

$$E^*(y_t) = \frac{1}{2}\frac{1}{1 + \frac{1-\mu}{\mu}\exp\left\{-\frac{T}{2\sigma^2}\left[\bar{y}_t^2 - \left(\bar{y}_t - \frac{1}{2}\right)^2\right]\right\}} =$$

At the median, $\bar{y}_t = 0$, and so the median post-discovery expected payoff is

$$E^*_{\text{MED}}(y_t) = \frac{1}{2}\frac{1}{1 + \frac{\mu}{1-\mu}\exp\left(\frac{t}{8\sigma^2}\right)}.$$

In Figure 2 we plot $E^*_{\text{MED}}(y_t)$ for $\sigma = 1$, six values of $\frac{\mu}{1-\mu} = \{10^{-1}, 10^{-2}, ..., 10^{-6}\}$.

$$\frac{1}{2}\frac{1}{1 + \frac{1-\mu}{\mu}\exp\left(\frac{t}{8}\right)}.$$

When $\sigma = 1$, learning is very protracted and the full-awareness model cannot replicate a sudden drop in expectations. But when σ is small, information accumulates rapidly, because evidence accumulates at the rate and a combination of a small μ and small σ can indeed generate a precipitous drop in expectations. Of course we are not generally free to choose σ – it is determined by the properties of the y_t sequence. On the other hand, one could add a signal at date t other than y_t, a signal not observed by the analyst, and thereby allow the full-awareness Bayesian model to fit the data.

These two examples focus on a particular sample path that does not reinforce the initial prior belief of Agent 1 so that the shifts in the posteriors can be large, especially if a lot of evidence has accumulated. This probably is the distinguishing feature of any major discovery. In any event, we shall ask whether, more generally, situations in which there is discovery can be modelled in the way in which the support of beliefs does not grow. This boils down to the question of whether sudden shifts in beliefs and expected payoffs are possible under Bayesian learning on sets the support of which is fixed. We shall ask that the Bayesian mode reproduce the time paths of payoffs (e.g., gambling profits), actions (e.g., decisions of whether to put up the ante and gamble again) and value (the present value of profits). The next section formulates and addresses this question more generally.

2.3 Experience learning

Time is discrete. In each period t the decision-maker chooses an action $x_t \in X$. At the end of the period t he observes $y_t \in Y$ and receives a reward $U(x_t, y_t)$. The random variable y_t has density

$$y \sim f(y \mid x, \theta) \tag{2}$$

where $\theta \in \Theta$ is a parameter. To simplify, let $f(y \mid x, \theta) > 0$ for all $(x, y, \theta) \in X \times Y \times \Theta$. Then no observation is logically impossible for any $\theta \in \Theta$, and the support of beliefs over Θ cannot shrink.

In this paper, 'Awareness' mean the support of beliefs. 'Unawareness' of a subset of Θ is equivalent to having beliefs that assign measure zero to that subset.

Awareness-Θ beliefs.—Define the one-step-ahead Bayes map

$$\mu' = B(x, y, \mu) = \frac{f(y \mid x, \theta)\,\mu(\theta)}{\int_\Theta f(y \mid x, \theta)\,d\mu(\theta)}. \tag{3}$$

Thus full-awareness beliefs or just 'full beliefs' are just the usual beliefs with awareness of all of Θ. Date-t full beliefs are then just t-fold iterates of the operator B starting from a full prior $\mu_0(\theta)$ with support Θ.[1]

Awareness-A beliefs.—With awareness A, the agent's actual belief is the restriction of μ to his awareness set, A:

$$m(\theta) = \mu(\theta \mid A) = \begin{cases} \frac{\mu(\theta)}{\mu(A)} & \text{for } \theta \in A, \\ 0 & \text{for } \theta \notin A. \end{cases} \tag{4}$$

It, too, is the t-fold iterate of B from a fictional prior belief $\mu_0(\theta \mid A_t)$. It is fictional because the actual prior awareness is A_0 and not A_t. The real-time sequence of beliefs would depend on the actual sequence (A_t). But since only the agent's current awareness matters, the precise route by which that awareness is reached does not matter for m_A.

[1] All this is precisely as in, say, Easley and Kiefer (1988, p. 1048).

2.3.1 Law of motion for m

We now derive conditions under which m obeys a first-order representation of the form

$$m' = \hat{b}(x, y, A', m), \tag{5}$$

which generalizes (3) to allow for discovery.

If $A' = A$, $\hat{b}$ is the Bayes map, but when $A' \supset A$, determining $m'(A' - A)$ in general requires knowledge of the entire sequence $h^t = (x^t, y^t)$ which would then be used to update m according to (4) with A' in place of A. But (5) will hold if we can invert m_t to obtain all the information about θ that h^t contains. By an application of the inverse function theorem the latter will generically be true if the number of sufficient statistics is less than the number of elements of A.

Let $L_t(h^t, \theta) = \prod_{s=0}^{t-1} f(y_s \mid x_s, \theta)$ denote the likelihood function. Assume that under repeated sampling from (2), the information in (x^t, y^t) has a k-vector of sufficient statistics, $T(h^t) \in \mathbf{T} \subseteq R^k$. Denote the law of motion for T by $T(h^{t+1}) = \tau(x, y, T(h^t))$, or simply by[2]

$$T' = \tau(x, y, T). \tag{6}$$

Since T is sufficient for θ, Fisher-Neyman factorization states that

$$L\left(h^t, \theta\right) = L_1\left(h^t\right) L_2\left(\theta, T\left(h^t\right)\right) \tag{7}$$

for all h^t, and

$$m(\theta) = \mu(\theta \mid T, A) = \frac{L_2(\theta, T)\,\mu_0(\theta)}{\int_A L_2(\theta', T)\, d\mu_0(\theta')} \quad \text{for } \theta \in A \tag{8}$$

Let $\Delta(\Theta)$ be the set of full-awareness beliefs reachable from μ_0 by observing some feasible (h^t).[3].

Let $\Delta_k(\Theta)$ denote the subset of $\Delta(\Theta)$ entailing a support of at least k elements. These are the feasible beliefs of someone who is aware of at least k distinct θ's. Let $\mathcal{M}_k$ be the set of measures over this set.

Now define the map (8) as

$$m = \phi_A(T),$$

so that $\phi_A : T \to \Delta(\Theta)$. The next result concerns the invertibility of this map for A fixed. When can we recover T from m that has a support A?

[2]Examples are provided below.

[3]Thus $\Delta(\Theta)$ is the union of the posteriors obtainable at any date from the supports of h^t as h^∞ ranges over its feasible set.

Proposition 1 *For $m \in \hat{M} \subset \mathcal{M}_k$ as given in (8) assume that the Jacobian of the #A-equation system on the RHS of (8) satisfies*

$$rank \underset{\#A \times k}{\left[\frac{\partial \mu (\theta \mid T, A)}{\partial T} \right]} = k. \tag{9}$$

Then $\phi_A^{-1}(m)$ exists on $\hat{M}$ and the representation (5) exists on $\cdot_k(\Theta)$.

Proof. Then the inverse function theorem implies that for each $m \in \hat{M}$ there is a function $\phi^{-1} : \hat{M} \to R^k$ giving $T = \phi_A^{-1}(m)$. Now let

$$\begin{aligned} m'(\theta) &= \mu\left(\theta \mid \tau\left[x, y, \phi_{A'}^{-1}(m)\right], A'\right) \quad \text{for } \theta \in A' \\ &\equiv \hat{b}(x, y, A', m). \end{aligned}$$

which is of the form (5) as desired. ■

Examples are provided in the Appendix.

2.4 Discovery and the Generalized Bayes map

It remains for us to specify a process for A which we shall combine with that defined by (5).

The law of motion for A.—Discovery thus should depend on A and on m which reflect the agent's awareness and his beliefs over the models he is aware of, and on the possible conflict between these models and the evidence, y^t– such a conflict presumably stimulates new discovery. And A is the support of m. Therefore, the evolution of A depends on m alone, in addition to possibly depending on x_t. Then, for any $A' \in \mathcal{B}(\Theta)$, let

$$\Pr(A_{t+1} = A' \mid x_t = x,\ m_t = m) = \alpha(A', x, m). \tag{10}$$

Combining this with (5), we obtain the generalized Bayes operator

$$m' = b(x, y, m),$$

where

$$b(x, y, m) = \int_{\mathcal{B}(A')} \hat{b}(x, y, A', m)\, \alpha(A', x, m)\, dA'. \tag{11}$$

2.5 Preferences

The agent's utility is $\sum_{t=0}^{\infty} \delta^t U(x_t, y_t)$ but he does not maximize $E\left\{\sum_{t=0}^{\infty} \delta^t U(x_t, y_t)\right\}$ where E is the standard expectations operator. But this operator does not allow extension of mass to sets of zero measure. For reasons we discussed at the outset, the standard updating of probabilities would preclude the agent from realizing that discoveries are possible and, hence, would exclude any motive for the agent to take actions that would raise the chances of making such discoveries. It also would lead to a different (and probably lower) lifetime utility. The agent will evaluate the current reward using awareness A, but he also will recognize the evolution of A in the sense to be made precise below.

The Bellman equation

We write the Bellman equation corresponding to these preferences. It differs from the standard treatment in just one respect: Instead of using the Bayes map (3), we use the generalized Bayes map (11). Preferences are defined recursively, somewhat related to Epstein and Zin (1989):

$$V(m) = \max_{x \in X} \int_Y U(x, y) + \delta V(b(x, y, m)) f(y \mid x, \theta) \, dm(\theta) . \tag{12}$$

This approach has two arguably desirable features

(i) the agent's action be independent of states whose possible existence is not yet discovered.

(ii) he correctly judges the consequences for tomorrow's continuation value of the discoveries that may occur as a result of the actions taken today.

General comparison

When is it equivalent to Bayes?

Parameters vs. Shocks.—It would not do for Bayes counterpart to have only y. Discovery has (y, A) as shocks, and a spanning argument would say you need a second shock for Bayes to have a chance to replicate things.

	Discovery	Bayes
utility	$U(x,y)$	
likelihood	$f(y \mid x, \theta)$; $\theta \in \Theta$	$g(y, z \mid x, \gamma)$; $\gamma \in \Gamma$
prior	$\mu_0(\theta \mid A)$	$\lambda_0(\gamma)$
updating	$\mu'_{\mid A'} = B(x, y, \mu \mid A')$	
awareness growth	$A' = \alpha\left(x, y, \mu_{\mid A}\right)$	$\lambda' = B(x, y, z, \lambda)$
	$m \equiv \mu_{\mid A}$	
	$b(x, y, m) = \int_{\mathcal{B}(\Theta)} \hat{b} \alpha dA'$	
policy function	$x = h(m)$	$x = H(\lambda)$
value	$V(m)$	$W(\lambda)$

where $x = H(\lambda)$ and W solve,

$$W(\lambda) = \max_{x \in X} \left\{ \int_{\Gamma \times Y \times Z} \left[U(x, y) + \delta W(B(x, y, z, \lambda)) \, g(y, z \mid x, \gamma) \right] dy dz d\lambda(\gamma) \right\} \tag{13}$$

Nothing is said about the dimension of Γ relative to that of Θ.

Two more examples

The observables are (x, y). We seek to show that the discovery model can generate the same outcomes for the observables as two standard examples do. The example is a search model, the second an exogenous growth model. We'll show that both can be generated by a process of 'discovery.' With data on outcomes alone, the two models are observationally equivalent for (x, y)

- construct a sample path (x_t, y_t)
- But the agent's V and expected payoffs do not converge

3 Example 1: Search

First we present the search model and then show that the outcomes of search can be generated by a model in which the parameters of a production function are 'discovered' gradually over time.

3.1 Search

Suppose a risk-neutral agent samples y from a distribution $G(y)$, for $y \in [0, 1]$. Suppose that y is the output of a technology if it is used for production. New technologies are sampled (with no resource cost) at the rate of one per period. The technology in use is then the highest sampled to date so that one i.i.d. sample is taken per period. The distribution of the maximum of t draws

$$Y_t = \max_{0 \le i \le t-1} (y_i)$$

is

$$\Pr(Y_t \le Y) = G^t(Y)$$

If one can also produce while sampling, Y_t, then $G^t(Y)$ is also the distribution of output at date t. When there is no sampling cost, this is the model of Burdett (1978) and Muth (1985), and when there is a sampling cost it is the model of McCall (1965).

3.2 Discovery

Now we show that the same decision problem can be framed in the discovery model. Assume that y_t is output produced using the production function

$$y_t = 1 - \left(x_t^p - \theta\right)^2 + \varepsilon_t, \tag{14}$$

where $x^p \in R$ is a 'production' decision (later there will also be a 'discovery' decision, $x^d \in \{0, 1\}$), $\theta \in \Theta \subseteq R$ is a parameter, and $\varepsilon_t \in R$ is an i.i.d. disturbance. The agent does not know θ, however. Instead, he is aware of a finite subset A of Θ, and his prior, denoted by $\mu_A(\theta)$, is uniform over A so that

$$\mu_A(\theta) = \begin{cases} \frac{1}{\#A} & \text{for } \theta \in A \\ 0 & \text{for } \theta \notin A \end{cases} \tag{15}$$

If the agent's awareness were to grow from A to A', say, where A' is a superset of A, then prior beliefs would again be given by (15) but with A' in place of A.

The discovery process.—Let discovery occur via i.i.d. sampling of θ from a distribution with C.D.F. $H(\theta)$ that has support Θ. This means that discovery is 'undirected' in the sense that the true θ is not involved. Then instead of writing (10) and (11) in distribution form, it is easier to write the difference equation for A'

$$A' = A \cup \theta \quad \text{with prob. } dH(\theta) \tag{16}$$

so that awareness remains unchanged with prob. $\int_A dH(\theta)$.

The role of ε.—If $\varepsilon \equiv 0$, the agent would quickly be able to learn the true θ; (14) has two solutions for θ:

$$\theta = x^p \pm \sqrt{1 - y}.$$

These two solutions are illustrated as θ_1 and θ_2 in Figure 3. Therefore one needs just two independent observations of (x, y) to solve for θ, and this would contradict the slower discovery process that we shall posit below.

We now proceed in two steps. First, we show that conditional on deciding to discover a new θ, the distribution sampling, the distribution of y_t as given by (14) will be $G(y)$ different way of deriving $G^t(Y)$ as the distribution of the largest y sampled.

3.2.1 Deriving $G(y)$ via discovery

Additional signals on θ.—Before starting out, the agent has seem T signals s_i with density

$$\xi(s \mid \theta) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{1}{2\sigma^2}(s-\theta)^2} \tag{17}$$

for[4] $i = -T, -T+1, ..., -1$.

[4] The idea is that the agent already has T observations of y at date $t = 0$

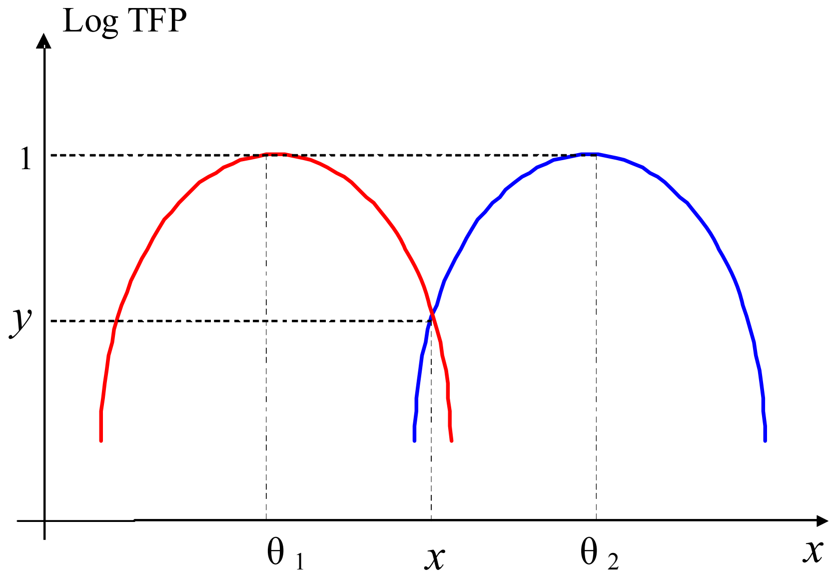

Figure 3: ANY (x, y) IMPLIES THAT $\theta \in \{\theta_1, \theta_2\}$

Optimal decisions.—Since (16) does not involve x^p, the production decision is static. Being risk neutral, the agent solves at each date

$$\min_{x \in R} E\left((x^p - \theta)^2 \mid h^{t,T}, A\right)$$

where $h^{t,T} = \left(x^t, y^t, s^T\right)$ and where

$$x^t = (x_0, ..., x_{t-1}), \; y^t = (y_0, ..., y_{t-1}), \quad \text{and} \quad s^T = (s_{-T}, s_{-T+1}, ..., s_{-1}).$$

The solution is to set x^p to equal the date-t expectation of θ:

$$x_t = E\left(\theta \mid h^{t,T}, A\right). \tag{18}$$

Now let T get large. Then the information in the history of s_t dominates any information that will be received via (y_t). Ignoring the information in y would mean that the expectation in (18) would be taken with respect to the posterior distribution which, in light of the uniform prior reads

$$\mu\left(\theta \mid h^{t,T}, A\right) = \frac{\prod_{i=-T}^{-1} \xi\left(s_i \mid \theta\right)}{\sum_{\theta \in A} \prod_{i=-T}^{-1} \xi\left(s_i \mid \theta\right)}, \tag{19}$$

for $\theta \in A$, where ξ is defined in (17).

Let θ^* denote the true value of θ. Then as $T \to \infty$ and the information becomes perfect, the posterior mean converges to

$$E\left(\theta \mid y^t, s^T\right) \overset{\text{a.s.}}{\to} \arg\max_{\theta \in A} \int \ln \xi\left(s \mid \theta\right) \xi\left(s \mid \theta^*\right) ds, \tag{20}$$

which also is the limit of the maximum-likelihood estimate of θ restricted to A. Although he has unlimited information that should clearly indicate θ^* when $\theta^* \in A$, when $\theta^* \notin A$, the agent obstinately treats that information as having arisen from an unlikely sequence (s_t, ε_t) governed by that value of $\theta \in A$ that best explains s^T.

Then

$$x_t = E\left(\theta \mid y^t, s^T\right) \overset{\text{a.s.}}{\to} \in \arg\min_{\theta \in A_t} |\theta - \theta^*| \tag{21}$$

We include the additional '$\in$' in (21) because there may be more than one argmin. Now let $\sigma^2 \to 0$, in which case, since $\varepsilon \overset{\text{i.p.}}{\to} 0$.[5] In that case we can easily show equivalence of the discovery model and the search model.

Equivalence.—Therefore for each $G(y)$ there exists a discovery process generating the same observations, in probability. If the agent starts with $A_0 = \varnothing$ and if he were to make one discovery per period, call it θ_t such that

$$y_t = 1 - (\theta_t - \theta^*)^2,$$

then the sequences of outputs will be the same at each date.

The technology correspondence.—For each output level $y < 1$, there are two technologies that give rise to it. Therefore we have the technology correspondence

$$\theta = \theta^* \pm \sqrt{1 - y}, \tag{22}$$

which is depicted in Figure 4. If any measurable selection $\psi(y)$ from this correspondence is consistent with $G(y)$, in the sense the the distribution over Θ that it gives rise to, call it $H^\psi(\theta)$, is consistent with G, then then all selections are consistent with G. That is, for any subset $B \subset \Theta$,

$$H^\psi(B) = G\left(\psi^{-1}(B)\right). \tag{23}$$

[5] This simplifies the relation between G and H. The probabilistic removal of ε will mean that we can associate output deterministically with the pair (θ^*, A) as we shall see below. In the theory of search by which new methods of production, as summarized by the cost of production, are drawn from a distribution (Telser (1982), Muth (1986), Kortum (1997)), the technology in use is the most efficient hitherto sampled, and it is used until a still better technology is discovered. In this approach, a sampled technology is an 'inspection good' – having discovered a new technology, the potential user does not have to try it in order to be able to evaluate it and compare it to the other technologies that he knows. Since most technologies must be tried before their full potential is known and is realized, this approach may work if time periods are long enough so that any experimentation can be treated as occupying a negligible fraction of any period, and that the bulk of the time in each period, the technology used is the best hitherto tried.

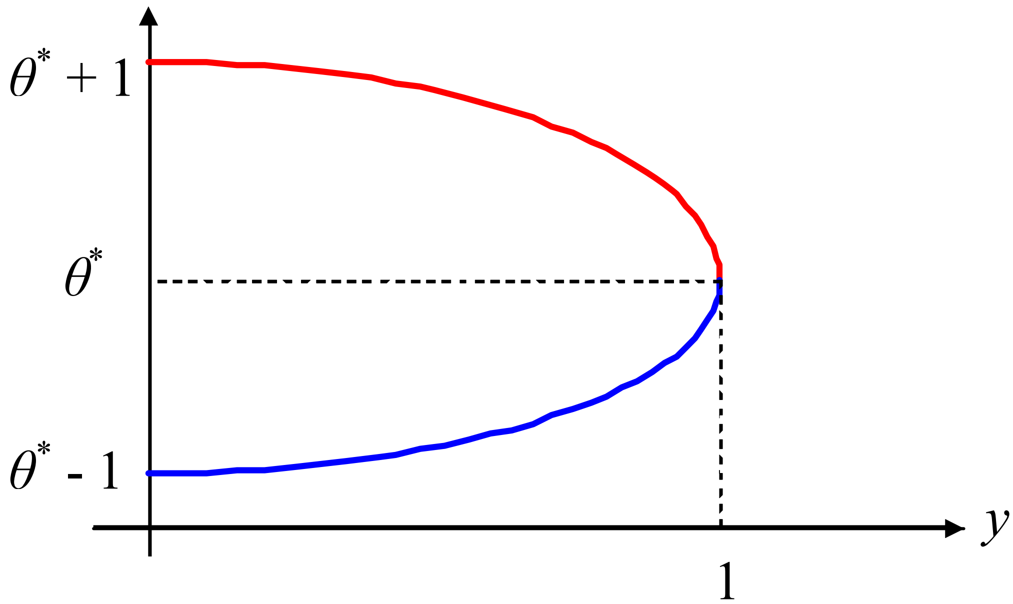

Figure 4: THE TECHNOLOGY CORRESPONDENCE (22)

Let us focus on just two selections from the technology correspondence: The smallest and the largest.

The discovery process that generates $G(y)$.—It is enough that we find just one selection from (22). The selection that maximizes θ of course obtains when, for each y, we take the larger solution in (22) for θ, i.e., $\theta(y) = \theta^* + \sqrt{1-y}$. Then $\theta_t \leq \theta \iff y_t \geq 1 - (\theta - \theta^*)^2$, and the CDF for θ needed to generate G is

$$H^+(\theta) = 1 - G\left(1 - (\theta - \theta^*)^2\right) \qquad \text{for } \theta \in [\theta^*, \theta^* + 1].$$

Alternatively, consider the smaller solution in (22) for θ_t, i.e., $\theta_t = \theta^* - \sqrt{1-y_t}$. Then $\theta_t \leq \theta \iff y_t \leq 1 - (\theta^* - \theta)^2$, and the CDF for θ needed to generate G is

$$H^-(\theta) = G\left(1 - (\theta^* - \theta)^2\right) \qquad \text{for } \theta \in [\theta^* - 1, \theta^*].$$

It does not matter which is chosen or, indeed, if a linear combination of the two is chosen as follows

$$H(\theta;\alpha) = \alpha I_{\{[\theta^*-1,\theta^*]\}} H^-(\theta) + (1-\alpha) I_{\{[\theta^*,\theta^*+1]\}} H^+(\theta) \tag{24}$$

for $\alpha \in [0,1]$. If all we can observe are the y's, the parameter α is not identified, and neither, of course is the selection ψ.

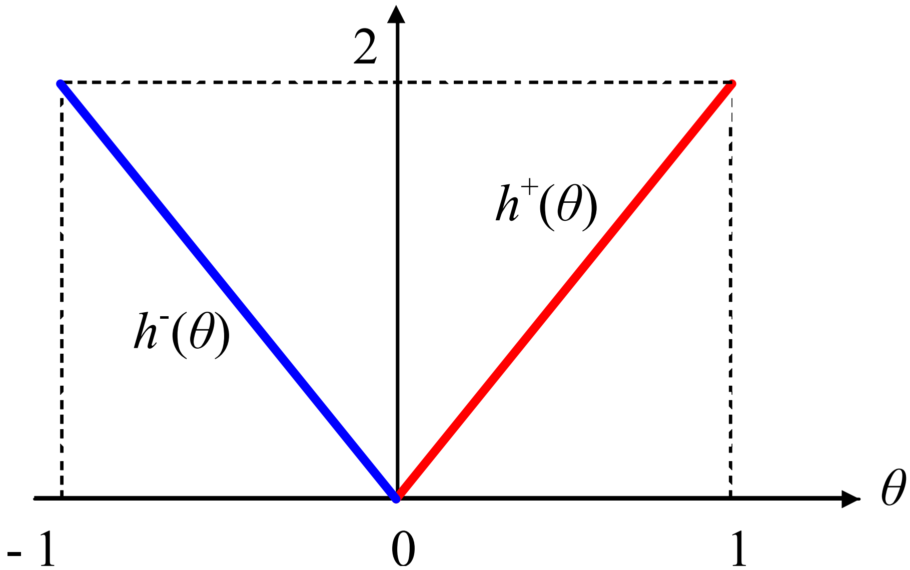

Figure 5: Two discovery densities consistent with a uniform $G(y)$

Example of a uniform G.—To illustrate (24), assume $G(y) = y$ for $y \in [0,1]$, and suppose that $\theta^* = 0$. Then

$$\begin{aligned} H^+(\theta) &= \theta^2 \qquad \text{for } \theta \in [0,1] \text{ and} \\ H^-(\theta) &= 1 - \theta^2 \qquad \text{for } \theta \in [-1,0] \end{aligned}$$

and the densities are

$$\begin{array}{ll} h^+(\theta) = 2\theta & \text{for } \theta \in [\theta^*, \theta^* + 1] \text{ and} \\ h^-(\theta) = -2\theta & \text{for } \theta \in [\theta^* - 1, \theta^*] \end{array} \tag{25}$$

and we illustrate them in Figure 5, where we also have extended their definition to the interval $[-1,1]$. Then any density of the form

$$h^\alpha(\theta) = \alpha h^-(\theta) + (1-\alpha) h^+(\theta) \tag{26}$$

$\alpha \in [0,1]$ also generates the C.D.F. $G(y) = y$.

3.2.2 Sampling costs and the search/discovery decision

While the production decision is static, when there is a sampling cost c, the search decision is dynamic. But having derived $G(y)$ via the distribution H in (24) and, for a specific example, in (25), it is straightforward to add a search-investment decision $x^d \in \{0,1\}$. The decision is $x = (x^p, x^d)$. As $T \to \infty$, so that

$$U(x,y) = y - cx^d.$$

As $\sigma \to 0$,

$$y \to 1 - (x^p - \theta^*)^2 = 1 - \min_{x \in A} (x - \theta^*)^2 \tag{27}$$

That is, as $\sigma \to 0$, the following three truths emerge

- $f(y \mid x, \theta)$ becomes degenerate, and we substitute y out in (12) via (27),
- $\mu(\theta \mid h, A)$ has converged with all its mass to $\arg\min_\theta (\theta - \theta^*)^2$
- V becomes stationary as h, being fully informative about θ^*, no longer matters.

Then (12) reads

$$\begin{aligned} V(A) &= \max_{(x^p, x^d) \in A \times \{0,1\}} \left\{ \left[\left[\int_A \left(U(x, y) + \delta \int_{\mathcal{B}(\Theta)} V(A') \, d\phi(A' \mid A, h, x) \right) d\mu(\theta \mid h, A) \right] \right\} \\ &= 1 - \min_{x^p \in A} (x^p - \theta^*)^2 + \max_{x^d \in \{0,1\}} \left\{ -cx^d + \delta \max \int_\Theta V(A \cup \theta) \, dH(\theta) \right\} \\ &= 1 - [\rho(A, \theta^*)]^2 + \max \left\{ \delta V(A), -c + \delta \int_\Theta V(A \cup \theta) \, dH(\theta) \right\} \end{aligned} \tag{28}$$

where

$$\rho(A, \theta^*) = \min_{\theta \in A} |\theta - \theta^*| \tag{29}$$

is the Hausdorff distance between θ and A in the Euclidean norm. Then since the operator is a contraction, we find that V must be of the form $V(A) = \hat{V}(\rho[A, \theta^*])$. Then (28) reads

$$v(\rho) = 1 - \rho^2 + \max \left\{ \delta v(\rho), -c + \delta \int_0^\infty v(\min(\rho, \rho')) \, dP(\rho') \right\}. \tag{30}$$

where $P(\rho')$ is the distribution of ρ' implied by $H(\theta)$ Change variables from ρ to $\hat{y}(\rho) = 1 - \rho^2$ and note that any $H(\theta)$ satisfying (23) then gives rise to the transform $G(y) = H(\theta)$. Then define w by

$$w\left(1 - \rho^2\right) = v(\rho)$$

Then (30) reads

$$v(y) = y + \max \left\{ \delta v(y), -c + \delta \int_0^\infty v(\max(y, y')) \, dG(y') \right\}$$

and the discovery decision then is

$$x^d = \begin{cases} 1 & \text{for } y < y^* \\ 0 & \text{for } y \geq y^* \end{cases} \tag{31}$$

where y^* solves the equation

$$c = \delta \int_0^\infty \left[v\left(y'\right) - v\left(y\right) \right] dG\left(y'\right). \tag{32}$$

The decision rule characterized by (31) and (32) is the optimal stopping rule for an infinitely-lived agent who samples at a cost c from a distribution G that he knows. then his search decisions and reservation y would be the same as if discovery was i.i.d., with distribution H given in (24).

So far we assumed that the agent's guess about ϕ is correct, which then allows the model to generate search from a distribution G that the agent knows. If the agent does not know ϕ and learns about it along the way, then this would have a counterpart in search theory of the agent learning about G. Rothschild (1974) showed there would be a sequence of stopping set $\left(Y_t^S\right)$ such that the agent stops sampling when $y^t \in Y_t^S$ for the first time. In the discovery model the agent would have stopping sets $\left(\Theta_t^S\right)$ such that he would stop sampling when $\theta^t \in \Theta_t^S$ for the first time.

4 Example 2: Growth

As with Example 1, we first outline the usual assumptions about the stochastic process, and we then show how it may arise via a process of discovery. Only a sketch will be provided here because the development is quite similar to that in Example 1.

4.0.3 Exogenous growth

Let y be output and consumption (there is no saving). Let

$$z_t = \frac{y_{t+1}}{y_t}$$

be the growth factor of output. Suppose, as Mehra and Prescott (1985) do, that z is first-order Markov:

$$\Pr\left(z_{t+1} \leq z' \mid z_t = z\right) = G\left(z', z\right). \tag{33}$$

Assume, additionally, that $z_t \geq 1$.

4.0.4 Discovery

We now derive (33) via discovery. The following procedure works if z_t is non-decreasing. Let utility be iso-elastic

$$U(y) = \frac{y^{1-\gamma}}{1-\gamma}, \text{ and } y_t = \frac{1}{\left|x_t - \theta\right|^k} + \varepsilon_t. \tag{34}$$

where, again, $x_t \in R$ is a decision, $\theta \in \Theta \subseteq R$ an unknown parameter, and ε_t an i.i.d. disturbance. Then

$$U(y) = -\frac{|x_t - \theta|^{k(\gamma-1)}}{\gamma - 1} + \hat{\varepsilon}_t \tag{35}$$

so that

$$k = \frac{2}{\gamma - 1} \Longrightarrow U(y) = -\frac{(x-\theta)^2}{\gamma - 1} + \hat{\varepsilon}_t. \tag{36}$$

Optimal decisions.—Discovery occurs exogenously, independently of the action taken. When (36) holds, it is again optimal to set the decision at the posterior mean of θ. Discovery again occurs exogenously, and we shall assume that a very precise signal of θ is available, as in Example 1, in the sense that T gets large, leading to (18), and (20). Assuming (17) and (36) the optimal policy again given by (21). We shall henceforth send the ε_t to zero in probability and ignore them. As in Example 1, their role is to avoid logical difficulties for the agent. Then output is

$$y_t = \frac{1}{\min_{x \in A_t} |x - \theta^*|^k}$$

Growth.—The growth factor between t and $t+1$ is

$$z_t = \left(\frac{|x_t - \theta^*|}{|x_{t+1} - \theta^*|} \right)^k \tag{37}$$

Since $x_t - \theta^*$ is non-increasing, this construction can work only if $z_t \geq 1$ for all t. Assuming that this is true, we can devise a process for x_t such (37) holds. As in example 1, define ρ again as in (29) to obtain

$$\rho_t = |x_t - \theta^*|$$

as the fraction of the gap between ideal practice, θ^*, and best practice x_t, i.e., the knowledge left undiscovered. Then the stochastic process ρ_t must be non-increasing and satisfies

$$\Delta \ln \rho_t = -\frac{1}{k} \ln z_t \leq 0.$$

If, for example, $z_t \in (z_1, ..., z_N)$ can take on N states and that it is a first-order Markov chain, as in Mehra and Prescott (1985) who assumed that $N = 2$, then $\Delta \ln \rho_t$ is itself a first-order Markov chain taking on values

$$\Delta \ln \rho_t \in \left\{ -\frac{\ln z_1}{k}, ..., -\frac{\ln z_N}{k} \right\}.$$

The process θ_t is once again not uniquely determined, for the same reason that an entire family of densities for θ in (26) is a solution to the search problem.

5 Discussion

Discovery is a growth in A. How might discovery be modeled? The best that we can do here is to try to capture the wisdom contained in the several bodies of theory that deals directly or indirectly with this subject. Let us briefly review the implications of some models of search theory, hypothesis testing theory, and game theory.

relate to market incompleteness models

5.0.5 Search theory

Search theory suggests that we model knowledge growth as sampling randomly from Θ, with the number of samples responding to effort. The search for new θ's would be directed if the agent can choose the distribution of θ from which to sample. Otherwise search is undirected.

Undirected search.—In this approach the action x_t would be the sampling rate. Conditional on a draw, a value θ' would be drawn from Θ or a subset of Θ. Perhaps new ideas $\theta' \in \Theta$ are generated via a stochastic process that is perhaps influenced by beliefs m (which have support A) and by decisions x^t such as R&D effort. Let $\psi(\theta', \theta)$ be a distribution over θ' conditional on θ which states that if the agent thinks of θ today, then θ' will occur to him tomorrow. Let x_t be the number of ideas θ_i with $i = \{1, 2, ..., x_t\}$ sampled and suppose that each idea is drawn from

$$\text{Prob}\{\theta_{i,t} \leq \theta' \in \Theta \mid m_t\} = \int_A \psi(\theta', \theta)\, dm. \tag{38}$$

Then any $\theta_{i,t} \notin A$ is a discovery. If the process in ψ is highly autocorrelated and if the dispersion of θ' is small, the agent can search only a small neighborhood of A, with most of the ideas sampled being duplicates of old ones. Examples of this approach are Telser (1982), Muth (1986) and Kortum (1997).

Directed search.—Bayesian treatments of the bandit problem and its various elaborations are in Rothschild (1974), Jovanovic and Rob (1990) and Jovanovic and Nyarko (1996). There is no discovery in these models in the sense that the parameter space is known from the outset.

5.0.6 Hypothesis testing

Discovery could, however, be more directed, such as the algorithms for model revision proposed by Sims (1971) and more formally by Radner (2002). The agent from time to time enlarges the set of model parameters that he wishes to entertain. One may limit discoveries to those that improve the agent's understanding of the world according to some statistical criterion. He may assign zero probability to any new idea that seems improbable enough in light of the evidence. For example, the discovery it should pass a likelihood-ratio test. Thus one could constrain

$$\text{Admissible discoveries} = \left\{ \theta' \in \Theta \mid \frac{L(h^t, \theta')}{\max_{\theta \in A} L(h^t, \theta)} \geq \lambda \right\}. \tag{39}$$

The parameter $\lambda \geq 0$ indicates the height of the hurdle that a discovery must clear. Indeed, if $\lambda > 1$, with probability 1 discovery must stop short of the true θ.[6]

(Radner 2002): *Deterministic model revision.*—Let N and k be positive integers and let $M(k)$ be the set of N-state Markov chains of order k. Let $\Theta = \cup_{k=0}^{\infty} M(k)$ and $A_k = \cup_{j=0}^{k} M(j)$. Radner assumes that A_k is augmented to A_{k+1} at a pre-specified date t_k where (t_k) is a sequence increasing in k. Since the dimension of the parameter space increases geometrically with k, $t_{k+1} - t_k$ must also increase with k if the agent is to get to learn much of anything. Radner's term 'model revision' is equivalent to a 'discovery' in our sense.[7]

5.0.7 Decision theory

We mentioned a collection of papers in the introduction. As a basis for the intuition regarding to payoff to discovery, Case-based decision theory of Gilboa and Schmeidler (1995) suggests to compare to similar cases. But the difficulty is finding what the similar cases are.

5.0.8 Game theory

In games, a player may take an action to which the other players assign probability zero. That agent then solves an inference problem related to the literature on what

[6] For any λ, the process is in line with the view of Galanis (2008 Section 6, page 17) argues that awareness grows when observing something which you thought was impossible. We can think of "impossibility" as a zero denominator

[7] Let us relate Radner's model to how beliefs are formed in ours and to eq. (4) in particular. Now $M(k)$ can be considered as a zero-measure subset of $M(k+1)$. and $\Delta(M(k))$ the set of measures over $M(k)$. Then letting $\mu_k \in \Delta(M(k))$, let the full-awareness prior be

$$\mu = \sum_{k=0}^{\infty} \alpha_k \mu_k.$$

Substituting this expression for μ on the RHS of (4) leads to

$$\mu_{A_k}(\theta) = \frac{\sum_{j=0}^{k} \alpha_j \mu_j(\theta)}{\int_{A_k} \sum_{j=0}^{k} \alpha_j d\mu_j(\theta)} \quad \text{for } \theta \in A_k, \quad \text{and} \quad \mu_{A_k}(\theta) = 0 \qquad \text{for } \theta \notin A_k.$$

to do when an opponent takes an off-equilibium action. When revelations of elements of $\Theta - A$ are made strategically, it is possible to specify the subjective probability $\hat{\phi}$ using information about the opponent's incentives to reveal various states, as Ozbay (2008) argues. Galanis (2008)....

Preferences will be defined recursively resembling Koopmans (1960).

Now we shall propose a model of how the agent would make decisions to maximize his objectives. Since standard decision theory cannot deal with changes in awareness we need an alternative theory. Here we have two options. The first, taken by Kochov (2009), is to provide axioms that one would find reasonable in this context, and formulate optimal decisions that way. The second, and the route that we shall take here, is to (A) set up a model of decisions in environments where discoveries happen or are likely to happen, such as decisions concerning the search for new technologies and scientific advances (B) compare the decisions that come out of the model to decisions that emerge in existing models of search and growth, and (C) argue that since the discovery model captures elements thought desirable in existing treatments of such problems, the model applies more broadly.

There are two features that a theory of decisions and discovery should capture.

1. *Current decisions must reflect only current awareness.*—The distribution of y_t is known to obey (2), but beliefs over θ are governed by (4) which is conditioned on current awareness, A, and on h^t. Output at date t cannot be produced using technologies yet to be discovered. Thus (x_t, y_t) can depend on (h^t, A^t), but not on future discoveries and events. The agent can make no bets or form any beliefs over states in $\Theta - A$ that he is unaware of. Thus the agent is in a situation that Li (2008) calls 'pure unawareness.' He can, however

2. *Decisions involving intertemporal trade-offs should reflect expected gains from discovery.*—The agent expect to make discoveries in $\Theta - A$ that are likely to improve his decisions and raise his utility. While he cannot imagine any elements of $\Theta - A$, he on average correctly judges the payoffs from investing in discoveries. That is, he can correctly judge the gains from such an investment even if he cannot at present say how and why he thinks these gains will come about. In other words, in his investment decisions he shall act *as if* he knew the law of motion of discoveries. Suppose that this was not so. Suppose that the agent underestimated the returns to discovery. Then discoveries would be scarce and that scarcity would raise their equilibrium value; the rate of return to discovery would then be much higher than rates of return on other investments. In fact, the rate of return to basic research has been only slightly higher than but comparable to the rate of return on physical investment, and why the market's valuation of cumulative research (as measured by Tobin's Q, say) is generally only slightly higher than its valuation of tangible assets.[8] To sum up, there is a tension:

[8] See Ch. 4 of Griliches (2000)

1. If the agent cannot imagine an event, it seems like he should underestimate the probability that it will occur.
2. The agent should be able to imagine that his knowledge will grow even if he cannot name any of its possible destinations. If people consistently underestimated the fruits of discovery, the rate of return to making efforts to discover things would presumably be very high.

6 Conclusion

The paper has defined discovery as an extension of the support of beliefs. It is not the first to take this approach, but perhaps the first with a quantitative orientation.

Several extensions are hoped for. One is an axiomatic justification for the recursive preferences over problems. Another is some general results concerning the correspondence between models of discovery and differently parametrized models with no discovery. Yet another is the issue of bias in the perceived returns to discovery.

References

[1] Aghion Philippe and Peter Howitt. A Model of Growth Through Creative Destruction. *Econometrica* 60, No. 2. (Mar., 1992), pp. 323-351.

[2] Auerswald, Philip, Stuart Kauffman, José Lobo and Karl Shell. "The production recipes approach to modeling technological innovation: An application to learning by doing." *Journal of Economic Dynamics and Control* 24, no. 3 (March 2000): 389-450.

[3] Bramoullé, Yann and Gilles Saint Paul. "Research Cycles." 2007.

[4] Burdett, Kenneth. "A Theory of Employee Job Search and Quit Rates." *American Economic Review* 68, no. 1 (March 1978): 212-20.

[5] Cho, In-Koo & Kenneth Kasa; Simon Fraser University "Learning About Identification." University of Illinois 2009.

[6] Easley, David and Nicholas Kiefer. "Controlling a Stochastic Process with Unknown Parameters." *Econometrica* 56 (1988): 1045-64.

[7] Epstein, Larry and Stanley Zin. "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework." *Econometrica* 57, No. 4 (Jul., 1989): 937-969.

[8] Fagin, Ronald and Joseph Halpern. "Belief, Awareness, and Limited Reasoning." *Artificial Intelligence* 34 (1988) 39-76.

[9] Galanis, Spyros. "Trade and the Value of Information under Unawareness."

[10] Gilboa Itzhak and David Schmeidler. "Case-Based Decision Theory." *Quarterly Journal of Economics* 110, No. 3 (Aug., 1995): 605-639.

[11] Griliches, Zvi. *R&D, Education, and Productivity: A Retrospective.* Harvard University Press, 2000.

[12] Jovanovic, Boyan and Yaw Nyarko. "Learning by Doing and the Choice of Technology," *Econometrica* 64, no. 6 (December 1996): 1299-1310.

[13] Jovanovic, Boyan and Rafael Rob. "Long Waves and Short Waves: Growth through Intensive and Extensive Search." *Econometrica* 58, no. 6 (December 1990): 1391-1409.

[14] Kocherlakota, Narayana. "Model Fit And Model. Selection." *Federal Reserve Bank of St. Louis. Review* July/August 2007, 89(4), pp. 349-60.

[15] Kochov, Asen. "A Model of Limited Foresight." U. Rochester, June 2009

[16] Kortum, Samuel. "Research, Patenting, and Technological Change." *Econometrica* 65, no. 6 (1997): 1389-1419.

[17] Li, Jing. "A Note on Unawareness and Zero Probability." PIER WP 08-022. January 2008.

[18] Mehra, Rajnish and Edward C. Prescott. "The Equity Premium: A Puzzle." *Journal of Monetary Economics* 15, no 2 (March 1985): Pages 145-161

[19] Muth, John F., "Search Theory and the Manufacturing Progress Function." *Management Science* (August 1986), pp. 948-962

[20] Nyarko, Yaw. "Learning in Mis-specified Models and the Possibility of Cycles." *Journal of Economic Theory* 55, no. 2 (December 1991): 416-427.

[21] Ozbay, Erkut. "Unawareness and Strategic Announcements in Games with Uncertainty." UMD February 2008.

[22] Pakes, Ariel. "On Patents, R&D, and the Stock Market Rate of Return." *Journal of Political Economy* 93, no.2 (April 1985): 390-409.

[23] Radner, Roy. "Bayesian Analysis and Model Revision for k'th Order Markov Chains with Unknown k." Stern School, NYU, October 2002.

[24] Rothschild, Michael, 1974. "Searching for the Lowest Price When the Distribution of Prices is Unknown." *Journal of Political Economy* 82, no. 4 (July 1974): 689-711.

[25] Sims, Christopher. "Distributed Lag Estimation When the Parameter Space is Explicitly Infinite- Dimensional." *Annals of Mathematical Statistics* 42, No. 5 (October 1971): 1622-1636.

[26] Telser, Lester. "A Theory of Innovation and Its Effects." *Bell Journal of Economics* 13, no. 1 (1982): 69-92.

[27] Venezia, Itzhak. "On the Statistical Origins of the Learning Curve." *European Journal of Operations Research*, 19 (1985): 191-200.

[28] Young, Peyton. *Strategic Learning and its Limits.* Oxford University Press 2002.

6.0.9 Appendix: Examples of (6)-(9)

TO BE CORRECTED....

Example 1: *Normal distribution.*—Let $y = x + \theta + \varepsilon$ and let $\varepsilon \sim N(0, \sigma_\varepsilon^2)$ so that

$$p(y \mid x, \theta) = \frac{1}{\sqrt{2\pi\sigma_\varepsilon^2}} \exp\left(-\frac{1}{2\sigma_\varepsilon^2}(y - x - \theta)^2\right).$$

and suppose that σ_ε is known. Let $\bar{y} - \bar{x} \equiv z$ be the sample mean of t x-adjusted signals. Then let $T \equiv (z, t) \in R^2$ be the sufficient statistic for θ. Thus $k = 2$. The transition function (6) for T is

$$T' \equiv \tau(x, y, T) = \begin{pmatrix} \frac{1}{t+1}(tz + y - x) \\ t + 1 \end{pmatrix}.$$

Then after canceling $\left(\sqrt{2\pi\sigma_\varepsilon^2}\right)^t$ from both sides, (7) reads

$$\exp\left\{-\frac{1}{2\sigma_\varepsilon^2}\sum_{s=0}^{t-1}(y_s - x_s - \theta)^2\right\} = \exp\left\{-\frac{1}{2\sigma_\varepsilon^2}\sum_{s=0}^{t-1}(y_s - x_s)^2\right\} \exp\left\{\frac{\theta}{\sigma_\varepsilon^2}tz - \frac{\theta^2 t}{2\sigma_\varepsilon^2}\right\}.$$

Let Θ be finite and let the full-awareness prior $\mu_0(\theta) = 1/\#\Theta$ be uniform. Let $A = (\theta_1, ..., \theta_N)$. Now μ_0 cancels in (8) so that, upon taking logs, it reads

$$\ln m(\theta) = \frac{\theta t}{\sigma_\varepsilon^2}\left(z - \frac{\theta^2}{2}\right) - \sum_{i=1}^{N}\frac{\theta_i t}{\sigma_\varepsilon^2}\left(z - \frac{\theta_i^2}{2}\right)$$

The Jacobian in (9) is

$$\begin{bmatrix} \frac{\theta_1 t}{\sigma_\varepsilon^2} - \sum_{i=1}^{N}\frac{\theta_i t}{\sigma_\varepsilon^2} & \frac{\theta_1}{\sigma_\varepsilon^2}\left(z - \frac{\theta_1^2}{2}\right) - \sum_{i=1}^{N}\frac{\theta_i t}{\sigma_\varepsilon^2}\left(z - \frac{\theta_i^2}{2}\right) \\ .. & .. \\ \frac{\theta_N t}{\sigma_\varepsilon^2} - \sum_{i=1}^{N}\frac{\theta_i t}{\sigma_\varepsilon^2} & \frac{\theta_N}{\sigma_\varepsilon^2}\left(z - \frac{\theta_1^2}{2}\right) - \sum_{i=1}^{N}\frac{\theta_i t}{\sigma_\varepsilon^2}\left(z - \frac{\theta_i^2}{2}\right) \end{bmatrix}_{N \times 2}$$

which is non-singular (?) even for $N = 2$, but certainly) for $N \geq 3$ as long as the θ's are distinct. Then any two rows of the matrix can be used and inverted to obtain T.

Example 2: *Binomial distribution.*—Let x be constant (i.e., no decision) and let (2) be given by

$$y = \begin{cases} 1 & \text{with Prob } \theta \\ 0 & \text{with Prob. } 1 - \theta \end{cases}$$

with $\theta \in [0, 1]$. Then in t trials and k successes the likelihood is $\theta^k (1-\theta)^{t-k} \equiv L_2(\theta, T)$, where $T = (k, t)$. The transition (6) reads

$$T' \equiv \tau(y, T) = \begin{pmatrix} k + I_{\{y=1\}} \\ t + 1 \end{pmatrix},$$

where I is the indicator function. Again let Θ be finite and again let $\mu_0(\theta) = 1/\#\Theta$ be uniform. Again let $A = (\theta_1, ..., \theta_N)$. Then after taking logs of both sides, (8) reads

$$\ln m(\theta) = k \ln \theta + (t-k) \ln(1-\theta) + C(A)$$

wherre $C(A, T) = -\ln \sum_{i=1}^{N} \theta_i^k (1-\theta_i)^{t-k}$. Then (9) reads

$$\underset{\text{N} \times 2}{\begin{bmatrix} \ln \theta_1 - \ln(1-\theta_1) + \frac{\partial C}{\partial k} & \ln(1-\theta_1) + \frac{\partial C}{\partial t} \\ .. & .. \\ \ln \theta_N - \ln(1-\theta_N) + \frac{\partial C}{\partial k} & \ln(1-\theta_N) + \frac{\partial C}{\partial t} \end{bmatrix}}$$