ELSEVIER

European Journal of Operational Research 105 (1998) 118–129

EUROPEAN JOURNAL OF OPERATIONAL RESEARCH

Theory and Methodology

Action-timing problem with sequential Bayesian belief revision process [1]

Jae-Hyeon Ahn [a,*], John J. Kim [b]

[a] AT&T Laboratories, Room 7E-530, 600 Mountain Avenue, Murray Hill, NJ 07974, USA

[b] Hunter College of the City University of New York, New York, NY 10021, USA

Received 20 April 1995; accepted 20 November 1996

Abstract

We consider the problem of deciding the best action time when observations are made sequentially. Specifically we address a special type of optimal stopping problem where observations are made from state-contingent distributions and there exists uncertainty on the state. In this paper, the decision-maker's belief on state is revised sequentially based on the previous observations. By using the independence property of the observations from a given distribution, the sequential Bayesian belief revision process is represented as a simple recursive form. The methodology developed in this paper provides a new theoretical framework for addressing the uncertainty on state in the action-timing problem context. By conducting a simulation analysis, we demonstrate the value of applying Bayesian strategy which uses sequential belief revision process. In addition, we evaluate the value of perfect information to gain more insight on the effects of using Bayesian strategy in the problem. © 1998 Elsevier Science B.V.

Keywords: Decision theory; Stochastic dynamic programming; Bayesian analysis; Action-timing problem; Simulation

1. Introduction

We consider the problem of deciding the best time to act when observations are made sequentially. More precisely, we analyze a special type of optimal stopping problem, called the *action-timing problem.* In the action-timing problem, observations are made sequentially from a certain distribution. Following the observation, the decision maker makes an irrevocable decision whether to take the observation under consideration or reject it for another observation. If the decision maker takes the observation, the decision-making process ends with the observation. If the decision maker rejects it, the same recursive decision-making process will continue with one less decision opportunity. For a decision made too early, a chance of better observation can be lost in the later time. For a decision made too late, a better observation could have been lost already. Therefore, decision maker has to choose the best action time while balancing the cost and benefit of acting now versus later.

* Corresponding author. Fax: + 1-908-582-2456; e-mail: ahn@ulysses.att.com.

[1] This research was supported by the PSC-CUNY Research Foundation.

PII S0377-2217(97)00036-2

Howard formulated and solved the action-timing problem using a dynamic programming framework (Howard, 1966). Navarro analyzed it with multiple stopping opportunities using a decision-analytic framework, and applied it to evaluate R & D investment opportunities (Navarro, 1987). Ahn extended the decision-analytic framework using a Markov process to analyze a situation where decision prospects change over time (Ahn, 1993) and applied it to the medical decision problem (Ahn and Hornberger, 1996). Unlike the *secretary problem* (Freeman, 1983; Samuels, 1991) which seeks to maximize the probability of selecting the best candidate by using ordinal measure, the action-timing problem seeks to maximize the expected utility by using cardinal measure.

In the action-timing problem, it is normally assumed that the distribution where observations are made is known with certainty. However, we often find that in reality we don't have the exact information on the distribution where observations are made. To illustrate the uncertainty involved in the decision environment, let's consider an entrepreneur who has a technology to sell.

An entrepreneur believes that he has a technology which can be applied usefully in a specific industry. Believing in the possibility of applying the technology usefully, the entrepreneur plans to sell the technology. Because the technology has certain proprietary characteristics, the entrepreneur approaches potential buyers sequentially and gives them the exclusive privilege for a specific time to evaluate the technology. After the evaluation time expires, a potential buyer offers a certain price for the technology and the entrepreneur decides whether to accept or reject the offer. The technology could either revolutionize the way products are produced or only have a marginal benefit. But, the entrepreneur faces information asymmetry: he does not know exactly how much cost saving can be achieved by applying the technology. But, actual manufacturers know their cost structures and potential savings.

In the above case, the offers made by potential buyers to the entrepreneur reveals some information about how well the technology is perceived by buyers. If an entrepreneur observes high offers, then the entrepreneur would believe that the technology may indeed be revolutionary. Conversely, if an entrepreneur observes low offers, then the entrepreneur would believe that the technology would have only a marginal benefit. But the entrepreneur doesn't know for certain how well the technology is perceived: he can only conjecture how well the technology is perceived. That is one source of uncertainty rooted in the information asymmetry between entrepreneur and potential buyers. Another source of uncertainty is the random observations from a specific distribution. Even though potential buyers have the same perception to the quality of the technology, their offers may be significantly different from each other because people value the same thing in different ways depending on their preferences.

Albright addressed a model updating parameters through Bayesian updating scheme for a single distribution function (Albright, 1977), but information asymmetry is not addressed in earlier models (Howard, 1966; Navarro, 1987; Ahn, 1993). In this paper, we develop a methodology which can addresses information asymmetry as well random observations. To address the information asymmetry issue, we introduce *states* and use a sequential Bayesian belief revision process to update our belief on state based on the previous observations. The uncertainty related to the random observations is represented for each state as a distribution where observations are made.

The rest of the paper is organized as follows. In Section 2, we develop an action-timing model with sequential Bayesian belief revision process and derive the explicit optimal solution. In Section 3, we walk through the decision-making process with the actual Bayesian belief revision process. In Section 4, we demonstrate the value of the sequential Bayesian belief revision process and value of perfect information. Finally, Section 5 provides concluding remarks.

2. Model with sequential Bayesian belief revision process

Let X_1, X_2, ..., be sequentially observed random variables from a probability density function $f(x)$. Suppose that x_k be the kth observation already made sequentially from $f(x)$, where $k \geq 1$. So x_1 would be the first observation, x_2 would be the second observation and so on. Decision will be made whether the decision maker takes the kth ($k \geq 1$) observed value x_k or rejects it in a sequential manner. Suppose that n number of decision stages are given ($n \geq 1$). In each stage, the decision maker makes one observation. Therefore, at most n observations can be made. If the decision maker decides to take the observed value, then the observation is the final value which the decision maker will take. If the decision maker rejects the value, then it is discarded and another observation is made to continue the process. The whole process ends when the decision maker either takes an observed value or runs out of all the decision stages.

Consider a strategy, with n remaining decision stages, such that the decision maker takes the observed value x_k if the value is grater than or equal to $d(n)$ and rejects it if x_k is less than $d(n)$. That is:

$$\begin{cases} \text{Accept} \quad x_k & \text{if } x_k \geq d(n) \\ \text{Reject} \quad x_k & \text{if } x_k < d(n). \end{cases} \tag{1}$$

If the decision maker rejects the value x_k, then another observation is made with $n-1$ remaining decision stages. Therefore, the problem is to find $d^*(n)$, the optimal strategy with n remaining decision stages, which maximizes the expected utility from this decision process. Because rejected value is discarded and there is usually finite decision stages, decision maker has to balance the risk and benefit of rejecting the observed value.

2.1. Sequential Bayesian belief revision process

Suppose a situation where observations are made from a distribution which is dependent on certain state. That is, depending on state, the decision maker observes values from different probability density functions. For example, offers for the exclusive technology license may be dependent on the degree of the innovation of the technology. Observations on stock prices may be dependent on the state of the specific company or economy. Sequential medical test results may be dependent on the state of health (existence of cancer or not) in medicine. Therefore, we will assume that probability density function $f(x)$, where observations are made, is different state by state.

Let's assume that there are m states. For each state S_i ($i = 1, \ldots, m$), suppose a corresponding probability density function $f_i(x)$, where $f_i(x) > 0$ for $\forall x \in R$. When observations are made sequentially, it is not known to the decision maker from which distribution they are made.

The uncertainty on the probability distribution function itself makes the problem more complex. The decision maker has to consider not only the uncertainty related to the random observations from a distribution, but some inference on the probability density function itself. Even though there is no information on density function itself, the idea is that we can use the prior observations to make inferences on the density function. Because there is one probability density function corresponding to each state, inference on the density function itself would be same as the inference on the state. The process of using prior observations sequentially to revise the belief on state would be called the *sequential Bayesian belief revision process*.

To show how the Bayesian belief revision process works, let's represent the observation vector as $O_k = \{x_1, x_2, \ldots, x_{k-1}, x_k\}$ with prior observations $x_1, x_2, \ldots, x_{k-1}$, and current observation x_k. The interpretation of these observations from non-Bayesian (or conventional statistical) approach is significantly different from Bayesian approach. Non-Bayesian approach would consider the observation vector O_{k-1} ($k \geq 1$) mere statistical realizations of values from a distribution which is unknown. On the other hand, Bayesian approach would use O_{k-1} to infer the distribution where observation vetor are made. Following the Bayesian approach, we will demonstrate how we can use the prior observations O_{k-1} and current observation x_k to update our belief on each state and use the information to derive the optimal strategy.

Suppose that we made a new observation x_k. Based on the observation vector O_{k-1} and x_k, we want to update our belief on state i. We can obtain the posterior belief on state i by performing a massive calculation of the impact of all previous observations on the posterior belief. But the calculation will be very complex if there we combine Eq. (4) and Eq. (5) to get

$$\frac{\pi(S_i|O_k)}{\pi(\neg S_i|O_k)} = \frac{\pi(S_i|O_{k-1})}{\pi(\neg S_i|O_{k-1})} \frac{\Pr(x_k|S_i)}{\Pr(x_k|\neg S_i)}, \quad \text{or} \quad O_d(S_i|O_k) = O_d(S_i|O_{k-1}) L(x_k|S_i). \tag{6}$$

Here, we can derive the condition when we can support the current belief as a new observation x_k becomes available. From the definition of posterior odds, $\pi(S_i|O_k)$ is represented in terms of $O_d(S_i|O_k)$ as

$$\pi(S_i|O_k) = \frac{O_d(S_i|O_k)}{1 + O_d(S_i|O_k)}.$$

Because $\pi(S_i|O_k)$ is the increasing function of $O_d(S_i|O_k)$, $\pi(S_i|O_k) > \pi(S_i|O_{k-1})$ if and only if $O_d(S_i|O_k) > O_d(S_i|O_{k-1})$. From Eq. (6), it means that upon the arrival of the new information x_k, $\pi(S_i|O_k) > \pi(S_i|O_{k-1})$ if and only if $L(x_k|S_i) > 1$. On the other hand, our prior belief on state i decreases upon the arrival of the new observation x_k, if the likelihood ratio $L(x_k|S_i)$ is less than 1. If $L(x_k|S_i)$ is 1, then the new observation x_k doesn't change our prior belief on the state. Thus, Eq. (6) shows in a very simple term how the belief revision process is related to the likelihood ratio.

2.2. Action-timing problem with sequential Bayesian belief revision

With the above framework of incorporating prior observations into the process of updating posterior belief on the state, we develop a model for the action-timing problem. Let's assume that we made an observation x_k ($k \geq 1$) at decision stage n. Index k implies that there exists $k-1$ prior observations. Now, we have to derive the optimal decision criteria whether to accept the observation x_k or not. When the observation is rejected, another observation will be made in the next decision stage with cost C ($C \geq 0$) and the value in the next decision stage will be discounted with unit period time preference α ($0 < \alpha \leq 1$). Note that even though the new observations may become available in the future and our belief on the state i would be updated based on them, $\pi(S_i|O_k)$ is the only and the best information we have. So, we use posterior belief $\pi(S_i|O_k)$ to derive strategies in the later stages.

The expected value at decision stage n is calculated by taking the expectation weighted by the belief on each state. If we apply arbitrary decision criterion $d(n)$ at decision stage n, then the expected value will be represented as

$$\sum_{i=1}^{m} \pi(S_i|O_k)\left[\Pr_i(X_n \geq d(n)) E_i(X_n|X_n \geq d(n)) + \alpha \Pr_i(X_n < d(n))\left(\sum_{j=1}^{m} \pi(S_j|O_k) V_j(n-1) - C\right)\right]. \tag{7}$$

Note that $V_i(n)$ is the expected value if decision maker is in state i and follows the optimal strategies from decision stage $n-1$ and it is represented as

$$V_j(n-1) = \int_{d^*(n-1)}^{\infty} x f_j(x)\,\mathrm{d}x + \alpha \int_{-\infty}^{d^*(n-1)} f_j(x)\,\mathrm{d}x \big(V_j(n-2) - C\big). \tag{8}$$

The first expression in Eq. (7) is the conditional expected value when X_n is greater than or equal to $d(n)$ and the second expression is the conditional expected value when X_n is less than $d(n)$. Eq. (7) can be represented again as

$$\sum_{i=1}^{m} \pi(S_i|O_k)\left[\int_{d(n)}^{\infty} x f_i(x)\,\mathrm{d}x + \alpha \int_{-\infty}^{d(n)} f_i(x)\,\mathrm{d}x \left(\sum_{j=1}^{m} \pi(S_j|O_k) V_j(n-1) - C\right)\right]. \tag{9}$$

The explicit derivation of the solution as shown in Eq. (12) was possible because of the simplicity of the exponential distribution. Derivations for normal and gamma distributions are equally possible. For general distribution functions, we can use numerical integration if it is difficult to derive the explicit solutions (Wolfram, 1994).

3. Decision-making process with sequential Bayesian belief revision

The typical decision-making process with sequential Bayesian belief revision is as following [Fig. 1].

Step 1: Observation: An observation is made from a distribution which the decision maker doesn't know during the whole decision-making process.

Step 2: Bayesian belief revision: Using the observation x_k $(k \geq 1)$, prior belief on state i, $\pi(S_i|O_{k-1})$ is revised to get the posterior belief on state i, $\pi(S_i|O_k)$ as described in Eq. (3).

Step 3: Calculation of $d^*(n)$: We calculate $d^*(n)$ as shown in Eq. (12).

Step 4: Decision: If $x_k \geq d^*(n)$ at decision stage n, x_k is accepted and the whole decision process ends. If not, it is rejected for another observation. Then, k becomes $k+1$, n becomes $n-1$ and the process goes back to step 1.

To demonstrate how the whole decision-making process works with sequential Bayesian belief revision, we simulate the events of making observations from an exponential distribution. For simulations, we use uniform random generator and use inverse transformation technique (Hiller and Lieberman, 1995) to generate random numbers for exponential distribution.

Example. Let's consider an example of two states; state 1 and state 2. Suppose that state 1 and 2 follow exponential distributions with parameters $\lambda_1 = 0.1$, $\lambda_2 = 0.2$, respectively. So, the probability density functions from which we make observations are represented as

$$f_1(x) = \begin{cases} 0.1\mathrm{e}^{-0.1x} & x \geq 0, \\ 0, & \text{otherwise} \end{cases} \qquad f_2(x) = \begin{cases} 0.2\mathrm{e}^{-0.2x}, & x \geq 0, \\ 0, & \text{otherwise} \end{cases}$$

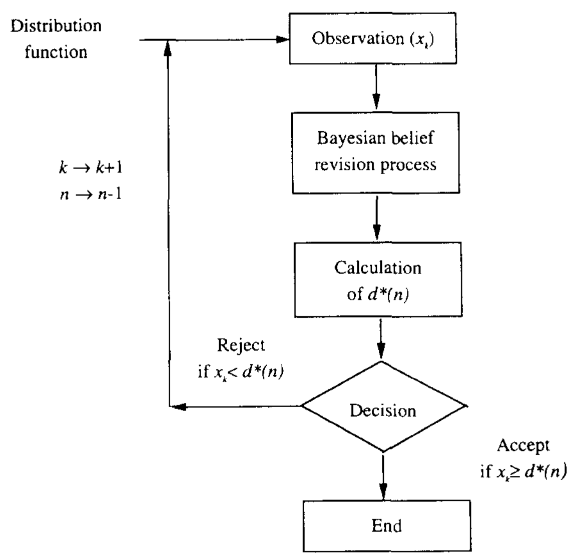

Fig. 1. Decision-making process with sequential Bayesian belief revision.

Let's assume that the unit period time preference $\alpha = 0.99$, the observation cost $C = 0.5$, salvage value $V_1(0) = V_2(0) = 0$ and the number of remaining decision stages $n = 15$. For convenience, we assume that the probability of being in state 1 or state 2 is equally likely. So, $\pi(S_1|O_{k-1}) = \pi(S_2|O_{k-1}) = \frac{1}{2}$, $k = 1$. With those assumptions, decision-making process proceeds as following. In the very beginning of the process, we start with $k = 1$ and $n = 15$.

Step 1: From the random number generation module, make an observation which happens to be 11.5709 (in this case, state 1 was selected and all the rest observations will also be made from state 1).

Step 2: From Eq. (3), posterior belief on state 1 is given as

$$\pi(S_1|O_k) = \Pr(S = S_1|x_k, O_{k-1}) = \frac{\pi(S_1|O_{k-1})0.1\mathrm{e}^{-0.1x_k}}{\pi(S_1|O_{k-1})0.1\mathrm{e}^{-0.1x_k} + \pi(S_2|O_{k-1})0.2\mathrm{e}^{-0.2x_k}}. \tag{15}$$

Because the observed value x_k is 11.5709, it is substituted in Eq. (15) and $\pi(S_1|O_k)$ is updated from 0.5 to 0.6139.

Note that the probability statement in Eq. (3) can be also considered as probability density in the continuous case as shown in Eq. (15). Because, for small Δx, we can approximate $\Pr(x_k \le X_k \le x_k + \Delta x|S_j)$ as $\Delta x\, f(x_k|S_j)$, Eq. (3) becomes

$$\pi(S_i|O_k) = \Pr(S_i|x_k, O_{k-1}) = \frac{\pi(S_i|O_{k-1})\Pr(x_k \le X_k \le x_k + \Delta x|S_i, O_{k-1})}{\sum_{j=1}^{m} \pi(S_j|O_{k-1})\Pr(x_k \le X_k \le x_k + \Delta x|S_j)}$$

$$= \frac{\pi(S_i|O_{k-1})\Delta x f(x_k|S_i)}{\sum_{j=1}^{m} \pi(S_j|O_{k-1})\Delta x f(x_k|S_j)} = \pi(S_i|O_{k-1})\frac{f(x_k|S_i)}{\sum_{j=1}^{m} \pi(S_j|O_{k-1})f(x_k|S_j)}.$$

Step 3: $d^*(n)$ is calculated as 16.1868 from Eq. (12).

Step 4: Because $x_k\,(= 11.5709) < d^*(15)\,(= 16.1868)$, the decision maker rejects x_k and goes back to step 1 with one less decision stage and one more observation. We continue to iterate the process until either an observed value is accepted or exhaust the decision stages.

Table 1 shows the observations, updated probabilities and the optimal decision criterion $d^*(n)$ for each decision stage. The first 3 observations didn't really change the belief on state 1, but the observation at $n = 12$ decreased it. Observation at $n = 11$ supported the belief on state 1. At $n = 10$, we observed 24.003 which also supported the belief on state 1. The value observed at $n = 10$ was sufficiently large enough to accept it. The whole decision process ended with the value 24.003. The net and discounted value back to initial stage is 20.42 which is a little bit less than the actual observation or 24.003. This is because (1) we rejected 5 times each costing 0.5 and (2) time preference. If there is (1) no cost for observation, (2) no time discount and (3) no limited number of decision stages ($n = \infty$), then we can see that decision maker will want to wait indefinitely until the maximum possible value (in this case, ∞) is observed.

Table 1
Decision-making process

	Remaining decision stage (n)					
	15	14	13	12	11	10
Observation (x_k)	11.5709	6.9453	7.0628	2.1033	14.3147	24.003
Updated belief on state 1 (S_1)	0.6139	0.6143	0.6174	0.4989	0.6757	0.9540
$d^*(n)$	16.1868	16.0721	15.9681	14.2716	16.3656	20.1605
Decision	reject	reject	reject	reject	reject	accept

From the perspective of belief updating, it is interesting to know how the new observation changes our belief on state. From Eq. (6), it is necessary and sufficient that we observe observation x_k, satisfying

$$L(x|S_i)=\frac{\Pr(x|S_i)}{\Pr(x|\neg S_i)}\geq 1,\quad \text{that is}\quad \Pr(x|S_i)\geq\frac{\sum_{j\neq i}^{m}\Pr(S_j)\Pr(x|S_j)}{1-\Pr(S_i)}. \tag{16}$$

In the above example of the two state case, we can calculate the intervals that increase or decrease our belief on state 1 as

$$\frac{\Pr(x_k|S_1)}{\Pr(x_k|\neg S_1)}=\frac{\lambda_1 e^{-\lambda_1 x_k}}{\lambda_2 e^{-\lambda_2 x_k}}\geq 1,\quad \lambda_1<\lambda_2. \tag{17}$$

By taking natural logarithm on both sides of Eq. (17), substituting values for $\lambda_1=0.1$, $\lambda_2=0.2$ and rearranging it, we have:

$$\begin{cases}\text{Support the belief on state 1,} & \text{if } x>6.9315,\\ \text{Support the belief on state 2,} & \text{if } 0\leq x<6.9315,\\ \text{Indifferent,} & \text{if } x=6.9315.\end{cases} \tag{18}$$

4. Bayesian and non-Bayesian strategy: A comparative analysis

Bayesian strategy is an adaptive strategy in a sense that optimal strategy and belief on the state are changed as the new observations become available. In Sections 2 and 3, we demonstrated how we can derive the optimal strategy and update belief on the state with the new observations. On the other hand, non-Bayesian strategy is a strategy that the optimal strategy and initial belief on the state are not changed with the arrival of the new observations. In this section, we will demonstrate the superiority of the Bayesian strategy which was developed in the previous sections. In Section 4.1, we compare Bayesian strategy and non-Bayesian strategy by simulation. In Section 4.2, we calculate the value of perfect information and find their implications for Bayesian and non-Bayesian strategy. For comparison purposes, we use the exponential distribution case discussed in Section 3.

4.1. Performance comparison between Bayesian and non-Bayesian strategy

Let's represent the optimal expected value for Bayesian strategy as $V_B(n)$ and non-Bayesian strategy as $V_N(n)$. To calculate $V_B(n)$, we run 10,000 simulations. For each simulation, 15 random numbers are generated from a distribution which was chosen randomly with probability $\Pr(S_i)$. $\Pr(S_i)$ is the known probability that observations will be made from state i. Then, we simulate the decision-making activities following the process described in Section 3. At the end of the process, we record a chosen value as the result of the simulation. The same process is repeated 10,000 times. Then, we calculate $V_B(n)$ by averaging out the 10,000 simulation results.

To calculate the expected value for non-Bayesian strategy, $V_N(n)$, we use Eq. (12) while maintaining the initial belief on state regardless of the observations, that is, $\pi(S_i|O_k)=\Pr(S_i)$. Thus, we can derive $V_N(n)$ without doing simulations.

Fig. 2 shows the difference between $V_B(n)$ and $V_N(n)$, or $V_B(n)-V_N(n)$. The probability p in Figs. 2–4 is the known probability $\Pr(S_i)$ with which observations will be made from state 1. For $n=1$, $V_B(1)$ has no significant meaning. At $n=1$ which is the final decision stage, we take whatever value we observe. Therefore $V_B(1)$ is just the expected value of 10,000 random observations from the two distributions which are also chosen randomly with probability 0.5 for each state. $V_N(1)$ is just the mathematical calculation of $V_B(1)$. For a sufficiently large number of simulations, $V_B(1)$ and $V_N(1)$ will be the same. For $n = 2$, we don't see much difference since there is only one opportunity to materialize the value of Bayesian belief revision process. For $n \geq 3$, we observe that $V_B(n)$ is greater than $V_N(n)$. We also observe that the value difference between $V_B(n)$ and $V_N(n)$ generally increases as n increases. For small n, there is a risk of misleading belief on states because of the small data set. The small data set may not be enough to update our belief sufficiently. For larger n, we can update our belief on states sufficiently, and we experience a greater value difference between Bayesian and non-Bayesian strategy.

At the same time, Fig. 2 shows that the value difference is greater when p value is close to 0.5 and less when p value is close to 0 or 1. For p value close to 0 or 1, using the previous observations doesn't really help in many cases. However, for p value close to 0.5, better inference on the real state using Bayesian strategy really pays off.

Generally, Bayesian approach will work better than non-Bayesian approach when uncertainty exists and the uncertainty is better understood by using previous observations. Bayesian strategy could perform poorly if prior observations mislead the belief on states. However, on the average, the Bayesian belief revision process updates the belief on the true state and helps decision maker to derive the corresponding optimal strategy better than the non-Bayesian approach. Fig. 2 clearly demonstrates the value of Bayesian strategy over non-Bayesian strategy for sufficiently a large number of simulations.

4.2. *Value of perfect information*

Value of *perfect* information on the state is the maximum value that a decision maker is willing to pay if he or she knows the *exact* state where observations are made. If the decision maker is certain of the state before observations are made, contingent strategy can be derived and exercised upon the knowledge of the state.

Value of perfect information for the non-Bayesian strategy, $\mathrm{VPI}_N(n)$, is calculated as the difference between value with perfect information on state for non-Bayesian strategy ($\mathrm{VWPI}_N(n)$) and value without information on state for non-Bayesian strategy ($V_N(n)$) or $\mathrm{VPI}_N(n) = \mathrm{VWPI}_N(n) - V_N(n)$. $\mathrm{VWPI}_N(n)$ can be mathematically calculated with the information on the exact state where observations are made. For each initial belief on state 1 ($= p$), we calculate the value difference between perfect information on state and without information. Fig. 3 shows $\mathrm{VPI}_N(n)$. We can see that $\mathrm{VPI}_N(n)$ increases as n increases. We also observe that the value of information is greatest at $\Pr(S_1) = 0.5$ over the whole range of $\Pr(S_1)$. This is intuitive since the information on state is more valuable when there is greater uncertainty on the state.

From the entropy perspective, $\Pr(S_1) = 0.5$ is the point that maximizes the entropy of the realization of state 1. To make this point more clear, consider the realization of state 1 as random variable X with probability p. Then, the entropy of the random variable X (Cover and Thomas, 1991) is calculated as

$$H(X) = -\sum_{x \in X} p(x)\log_2 p(x) = -p\log_2 p - (1-p)\log_2(1-p). \tag{19}$$

From Eq. (19), we can easily see that entropy of X is maximized at $p = 0.5$ and takes value 0 at $p = 0$ and 1. This analysis supports the observations in Fig. 3 that $\mathrm{VPI}_N(n)$ takes the maximum value at $p = 0.5$ and decreases to zero when there exists no uncertainty. That is, $\mathrm{VPI}_N(n)$ takes the maximum value when there exists the greatest uncertainty on state which is represented by entropy measure.

Fig. 4 shows the difference between value with perfect information on state for non-Bayesian strategy ($\mathrm{VWPI}_N(n)$) and $V_B(n)$. We can see that there is not much difference. That is, the Bayesian strategy works as well as the non-Bayesian strategy with perfect information on the state. Because the Bayesian belief updating process in a sense provided enough information on the states even though it is not the perfect information, Bayesian strategy performs as good as the non-Bayesian strategy with perfect information on the state. This implies that perfect knowledge on states does not provide valuable information if we use Bayesian strategy.

From the entrepreneur's point of view who wants to sell a specific technology, the value of perfect information is the maximum amount of money he wants to invest to know how well the technology is perceived by potential buyers. If the entrepreneur sticks with non-Bayesian strategy, he would like to pay more as he is not sure whether the technology is revolutionary one or has a marginal benefit (as p approaches 0.5). Also, the entrepreneur would pay more if there are more potential buyers (as n increases). On the other hand, if the entrepreneur applies Bayesian strategy, perfect information on state wouldn't help much. So, the entrepreneur would not bother to spend money to know how the technology is perceived by the buyers. Instead, the entrepreneur can get the almost same quality of information on state by analyzing the prior offers.

5. Conclusion

In this paper, we developed an action-timing problem with sequential Bayesian belief revision process. The process allowed us to update belief on state by using the previous observations which become available sequentially. By taking advantage of the conditional independence of the observations for a given distribution, the belief updating process turned out to be a simple recursive form. With the sequential Bayesian belief revision process, the action-timing model was developed to address the information asymmetry issue and the explicit optimal strategy for the model was derived.

In the following simulation analysis, we clearly demonstrated the value of the Bayesian strategy which used sequential belief revision process. Bayesian strategy showed a consistently higher expected value than the conventional non-Bayesian strategy (Fig. 2). The value difference became more substantial as decision stages (n) increased and entropy on state increased (13% increase of expected value when $n = 15$ and $p = 0.5$). Further analysis on the value of perfect information on states for non-Bayesian strategy also showed that it increased as decision stages (n) increased and entropy on state increased (Fig. 3). However, perfect information on states for Bayesian strategy turned out not to be the valuable one. Because the Bayesian belief revision process already used the information from the prior observations in deriving the optimal strategy, perfect information on states was not that useful. Therefore, we didn't need to pay for the perfect information on state if we apply Bayesian strategy with sequential belief revision process (Fig. 4).

The methodology developed in this paper provides a new theoretical framework for addressing the information asymmetry issue in the action-timing problem context. The methodology can be applied usefully in the decision environment, such as in the areas of technology acquisitions, asset transactions and negotiation processes. The decision maker will be able to derive the optimal threshold level in his or her problem and better prepare for the actual transactions and negotiations.

References

Ahn, J.-H. (1993), "Optimal option-exercising strategy with an application for kidney transplant decisions", unpublished Ph.D. dissertation, Department of Engineering-Economic Systems, Stanford university.

Ahn, J.-H. and Hornberger, J.C. (1996), "Involving patients in the cadaveric kidney transplant allocation process: A decision-theoretic perspective", *Management Science* 42/5, 629–641.

Albright, C.S. (1977), "A Bayesian approach to a generalized house selling problem", *Management Science* 24/4, 432–440.

Cover, T.H. and Thomas, J.A. (1991), *Elements of Information Theory*, John Wiley and Sons.

Freeman, P.R. (1983), "The secretary problem and its extensions: A review", *International Statistical Review* 51/2, 189–206.

Hiller, F.S. and Lieberman, G.J. (1995), *Introduction to Operations Research*, 6th Ed., McGraw-Hill Publishing Co.

Howard, R.A. (1966), "Dynamic programming", *Management Science* 12/5, 317–348.

Navarro, E.J. (1987), "Decision analysis with uncertain future opportunities", unpublished Ph.D. dissertation, Department of Engineering-Economic Systems, Stanford university.

Samuels, S.M. (1991), "Secretary problems", in: B.K. Ghosh and P.K. Sen (eds.), *Handbook of Sequential Analysis*, Marcel Dekker, 381–405.

Wolfram, S. (1994), Mathematica™, Addison-Wesley.