

Forward Induction and Dynamic Optimization under Uncertainty*

Katherine Hauck and Tiemen Woutersen

September 10, 2023

Abstract

Backward induction is a “workhorse” to solve dynamic optimization problems. However, this technique assumes full information and works less well when an individual learns through updating prior beliefs or test results. Information sets and beliefs tend to change over time, and it is helpful when dynamic optimization techniques allow for these changes. Our motivating example is a doctor who wants optimal treatment for their patients. Over time, the doctor learns more about the patients, i.e. the doctor updates their prior beliefs, and, the doctor lets treatment depend on tests. This information uncertainty prevents a standard application of backward induction. To solve this problem, we integrate the function of interest with respect to the distribution conditional on the parameters and then integrate again with respect to the prior beliefs. We allow for both updating prior beliefs and future treatments that depend on future tests, and we show that this technique is feasible in many contexts.

Keywords: Dynamic optimization, Bayesian updating, Forward induction

1 Introduction

Information sets and beliefs tend to change over time, and this may complicate dynamic optimization problems. In particular, backward induction, the “workhorse” of dynamic optimization, assumes full information and works less well when an individual learns through updating prior beliefs or through test results. Our motivating

*We thank Jim Powell for very helpful discussions. This work was funded by the Eller Center for Management Innovations in Healthcare at the University of Arizona.

example is a doctor who wants optimal treatment for their patients. Over time, the doctor learns more about the patients, i.e. the doctor updates their prior beliefs, and, furthermore, the doctor lets treatment depend on test results. The information uncertainty from prior beliefs and future test results prevents a standard application of backward induction. To allow for future treatments that will depend on future tests, we integrate the function of interest with respect to the distribution (conditional on the parameters) and then integrate again with respect to the prior beliefs. In particular, we specify the distribution of the outcomes as a vector function of the parameters. We show that the technique we propose is feasible in many contexts beyond optimal treatment plans by doctors.

Suppose that the results of a medical test determine the subsequent treatment of a patient. This situation causes what is called an “endogeneity problem” in economics: the subsequent treatment depends on the test results, so the subsequent treatment cannot be taken as given. This type of endogeneity problem makes a standard application of backward induction difficult because the outcome of the test is only known in the future. A central feature of backward induction is that one calculates the subsequent costs from certain states. However, in Bayesian learning, the subsequent costs may be uncertain and the knowledge about these costs depends on the prior decisions. We illustrate this using an example.

One way to deal with the uncertainty of future events is to have a prior belief about the probability of the future event and to employ Bayesian analysis. However, this method creates a tension if one also uses backward induction. In particular, while backward induction requires later period information to calculate prior period values, Bayesian updating uses prior period information to make later period decisions. We show that in cases of dynamic optimization with information uncertainty, our methodology, which combines prior beliefs and forward induction, may make these optimization problems solvable. Further, we show that our method is an effective way to deal with endogenous paths, such as treatments that depend on test results, because forward induction can incorporate learning about tests and, more generally, parameter uncertainty. The method we propose also solves the endogeneity problem described above in a larger class of maximization of expected utility in sequential decision problems.

1.1 Related Literature

This paper relates to several strands of literature. First, this paper builds upon the dynamic discrete choice literature which incorporates beliefs and learning into the model. Akerberg (2003) creates a related dynamic discrete choice model with beliefs and uses this model to generate predictions in an advertising and consumption setting. Crawford and Shum (2005), Darden (2020), and Mira (2007) all create models in which the priors are known to the researcher. Crawford and Shum (2005) apply their model to prescriptions in the pharmaceutical industry. Darden (2020) applies his stochastic dynamic model to smoking behavior. Mira (2007) implements his model with respect to fertility decisions that may impact infant mortality. Covert (2015) applies a model which includes learning under uncertainty to oil extraction. In a game theory context, Kamenica and Gentzkow (2011) introduce Bayesian persuasion and show that a posterior is Bayes plausible if the expected posterior probability equals the prior predictive distribution.

This paper also relates to Bayesian decision making in medicine. Freedman and Spiegelhalter (1992) apply Bayesian analysis, with particular attention to the choice of priors, to a clinical trial for six different treatments of colorectal carcinoma. Insua et. al (2020) analyze Bayesian design, including both parametric and nonparametric designs, for life tests and maintenance problems like knee replacements. Ashby and Smith (2000) consider the decision-making process of pregnant women deciding to take folic acid supplements. The choice to take folic acid supplements depends on the uncertain likelihood of a neural tube defect in the fetus, and Ashby and Smith model this uncertainty in a Bayesian framework. Müller et al. (2007) propose an algorithm that approximates an infinite dimensional optimal decision rule by a finite dimensional decision rule.

2 Proposition and Examples

The uncertainty facing a doctor or another decision maker includes parameter uncertainty, which can be formalized using Bayesian analysis. Further, the outcome of interest may depend on future tests or other intermediate outcomes. We call the cases of parameter uncertainty or intermediate outcome uncertainty “information un-

certainty.” Such information uncertainty may prevent the application of backward induction as a tool for solving dynamic optimization problems. We find the following procedure to be feasible and efficient in the sense that it uses all information.

We write the function of interest without conditioning on intermediate outcomes. Next, we integrate the function of interest with respect to the distribution conditional on the parameters and then integrate it again with respect to the prior belief. In simple cases (see the survival function in example 4), this integration is equivalent to integrating with respect to the prior predictive distribution. However, in example 1, our procedure is not equivalent to integrating with respect to the prior distribution. Further, our calculations do not require the priors to be updated for every outcome, and this feature will be helpful in the examples below.

Proposition: Backward induction may not be feasible in dynamic optimization problems with information uncertainty. However, in such cases, forward induction and Bayesian learning may still be feasible.

We prove the proposition by stating examples of the impracticality of backward induction (example 3) and of the feasibility of forward induction in all examples.

Example 1: Medical Treatment

Let $Y \in \{0,1\}$ denote whether a medical treatment is a success. Further, the probability of a success is p and the prior on p is uniform on the interval $[0,1]$.

At $t = 0$, the doctor wants to know the probability of a successful treatment. Suppose the test results, which can inform treatment decisions, are only available at $t = 1$. Let W denote the test result, where $W \in \{0,1\}$. Let $P(W = 1) = \frac{p}{2}$, and $P(W = 0) = 1 - \frac{p}{2}$. When the doctor observes the test result W , they can update their prior. Suppose that the treatment is done when $W = 1$ and the treatment is not done otherwise. At $t = 0$, the probability of a successful treatment is

$$\int_0^1 \frac{p}{2} \cdot p dp = \frac{1}{6}.$$

Note that the prior predictive distribution of Y (using just the prior on p) is not helpful to solve the optimization problem here because the test and the outcome of the operation both depend on the same unobserved p .




Figure 1: City Distances Example

Example 2: Traveling between Cities (or Health Outcomes)

Figure 1 shows a simple example which can be understood as depicting the cost of traveling between different cities. At $t = 0$, the agent has a prior belief about the travel costs A and B , and they choose between traveling to city Y_1 , which incurs cost A , and traveling to city Z_1 , which incurs cost B . Let $A|\alpha \sim N(\alpha, 1)$ and let $B|\beta \sim N(\beta, 1)$, where $\pi(\alpha) \sim N(0, 1)$ and $\pi(\beta) \sim N(0, 1)$. The prior beliefs are independent. Further, $A|\alpha$ and $B|\beta$ are independent. Going to a Y_i city always incurs cost A , and going to be Z_i city always incurs cost B , where $i \in \{1, 2\}$. The priors do not imply whether α or β is lower, so without loss of generality, assume that the agent first chooses to travel to Y_1 . When A is large, in particular larger than the threshold τ , then the agent will opt to incur the cost B in the second period. Therefore, the total cost for traveling from “Begin” to “End” in Figure 1 is

$$Cost = A + 1(A < \tau) \cdot A + 1(A > \tau) \cdot B.$$

In expectation, this cost is

$$E(Cost) = \int_{-\infty}^{\tau} \int_{-\infty}^{\tau} \{a + 1(a < \tau) \cdot a + 1(a > \tau) \cdot b\} f(a) f(b) da db.$$

The prior predictive distribution of A is $f(a) = \int_{-\infty}^{\infty} \phi(a - \alpha) \phi(\alpha) d\alpha$, where $\phi(\cdot)$ is the density of the standard normal. This result gives that the prior predictive




Figure 2: City Distances Example with Learning

distribution is normally distributed with mean zero and variance two, which yields

$$\begin{aligned}
 E(\text{Cost}) &= \int_{-\infty}^{\tau} a \cdot f(a) da = \int_{-\infty}^{\tau} \frac{a \cdot \exp\{-\frac{a^2}{4}\}}{\sqrt{2} \cdot \sqrt{2\pi}} da \\
 &= \frac{-\exp\{-\frac{a^2}{4}\}}{\sqrt{\pi}} \Big|_{-\infty}^{\tau} = \frac{-\exp\{-\frac{\tau^2}{4}\}}{\sqrt{\pi}}.
 \end{aligned}$$

Choosing $\tau = 0$ minimizes the expected cost.

We can extend the above example by adding cities Y_3 , Y_4 , Z_3 , and Z_4 . Let A^* be the cost of traveling to Y_3 (Y_4) from Y_2 or Z_2 (Y_3 or Z_3), and let B^* be the cost of traveling to Z_3 (Z_4) from Y_2 or Z_2 (Y_3 or Z_3). Let A^* and B^* be known. In this case, forward induction requires fewer calculations than backward induction.

Example 3: Exploring while Traveling

The value of learning can also be put into this framework, which is illustrated in Figure 2. Let the priors on α and β and the distribution of $A|\alpha$ and $B|\beta$ be the same as in example 2. In a learning context, when an agent has an opportunity to learn more about the costs A and B , they may do so in order to minimize the expected total cost, even if this learning means they may incur a higher cost in the immediate period.

In the example illustrated by Figure 2, the choice of path depends on information learned along the way, which complicates backward induction. This complication arises because backward induction necessarily uses later period information to calculate prior period values, while Bayesian updating uses prior period information to make later period decisions.

In this example, the priors on α and β do not imply a choice between Y_1 or Z_1 . Suppose the individual first travels to Y_1 at cost A . If the cost A is below a threshold τ_L , then this individual will travel to Y_2 at cost A , and finally to Y_3 at cost A . Note that learning the value A is informative beyond updating a prior on α . However, if the cost A is above the threshold τ_L , then the individual will travel from Y_1 to Z_2 and will learn B . The last decision of the individual depends on whether A or B is smaller. The following function gives the cost of traveling,

$$H(A, B, \tau_L) = A + 1(A \leq \tau_L) \cdot 2 \cdot A + 1(A > \tau_L) \cdot \{B + \min(A, B)\}.$$

In this example, we can use the prior predictive distribution because the priors on α and β are independent. In the appendix, we show that the expectation of $H(A, B, \tau_L)$ as a function of τ_L can be written as

$$H(\tau_L) = -\{1 + \Phi(\tau_L/\sqrt{2})\} \frac{\exp\{-\frac{\tau_L^2}{4}\}}{\sqrt{\pi}} - 2 \frac{1 - \Phi(\tau_L)}{\sqrt{2\pi}}.$$

The choice of the threshold τ_L that minimizes the expected cost is $\tau_L = -0.39$. This example shows that learning A and B is informative beyond updating the priors on α and β , as in the previous example. Further, learning B is useful in this case because A and B appear in the final decision. The value derived from this learning causes the threshold τ_L to be lower than the threshold $\tau = 0$ in the last example, illustrating the value of increasing the information set.

This example also illustrates the challenge for backward induction. For backward induction, one needs to calculate the travel cost from Y_2 onward. This travel cost depends on B . However, one does not know if B is known or whether to rely on the prior on B . The same holds for the travel cost from Z_2 onward. Forward induction does not have this issue, and that is why we recommend it in dynamic models with learning.

Example 4: Survival Function

The following example illustrates forward induction in continuous time. Let T denote the survival time of a patient, and let $S_T(t|\alpha)$ denote the survival function

conditional on the parameter α , where

$$S_T(t|\alpha) = \Pr(T \geq t|\alpha) = \exp(-\alpha t).$$

Let $\pi(\alpha)$ denote the prior on α at $t = 0$, where $\pi(\alpha) = \exp(-\alpha)$. Then the prior predictive survival function is

$$\begin{aligned} S_T(t) &= \int_0^\infty S_T(t|\alpha)\pi(\alpha) d\alpha \\ &= \int_0^\infty \exp(-\alpha t) \cdot \exp(-\alpha) d\alpha \\ &= \frac{1}{t+1}. \end{aligned}$$

Suppose that the patient is still alive at time τ . Then the patient and their doctor may also be interested in the survival function that conditions on survival up to τ . Survival up to τ is an endogenous event because this survival is correlated with T . Fortunately, Bayesian analysis that uses the survival functions and the initial prior can easily handle this endogeneity problem. Let Y denote the survival time after τ . Then this survival function is

$$\begin{aligned} S_Y(y) &= \Pr(T \geq \tau + y | T \geq \tau) \\ &= \frac{\Pr(T \geq \tau + y)}{\Pr(T \geq \tau)}. \end{aligned}$$

Using the survival function above yields

$$\begin{aligned} S_Y(y) &= \frac{S_T(\tau + y)}{S_T(\tau)} \\ &= \frac{\tau + 1}{\tau + 1 + y}. \end{aligned}$$

In this example, it is feasible to update the prior for every τ . Note, however, that the calculations above do not require this.

3 Conclusion

In many decision-making scenarios, agents experience information uncertainty and learn new information over time. An example of such a scenario is when doctors are uncertain about a patient’s diagnosis and learn more information via test results. In this case, doctors have a prior belief about a patient’s diagnosis and update this belief based on the results of the test. Such information uncertainty presents a problem in backward induction.

We propose a method to address this issue. Instead of using backward induction to solve dynamic optimization problems, which include updating prior beliefs, we propose an algorithm which uses forward induction. We specify the distribution of the outcomes as a vector function of the parameters, and the function of interest is integrated with respect to the distribution conditional on the parameters and then integrated again with respect to the prior. In simple cases, this integration is equivalent to integrating with respect to the prior predictive distribution. This approach solves the endogeneity problem of later period decisions depending on earlier period information because forward induction can incorporate learning about parameter values. Intuitively, this algorithm (using forward induction) performs better than backward induction in some cases of dynamic optimization with learning because both forward induction and Bayesian updating use prior period information to make later period decisions. In particular, the method we propose can be applied to solve dynamic optimization problems in which agents learn through test results, such as doctors updating their beliefs based on medical tests and using those updated beliefs to make treatment decisions. Another example is a manager making investment decisions and learning more about revenue potential and costs.

References

- [1] Akerberg, D. (2003): “Advertising, Learning, and Consumer Choice in Experience Good Markets: A Structural Examination.” *International Economic Review*, 44, 1007–1040.

- [2] Ashby, Deborah, and Adrian F. M. Smith (2000): “Evidence-based medicine as Bayesian decision-making.” *Statistics in Medicine*, 19, 3291-3305.
- [3] Covert, Thomas (2015): “Experiential and Social Learning in Firms: The Case of Hydraulic Fracturing in the Bakken Shale.”
- [4] Crawford, Gregory S. and Matthew Shum (2005): “Uncertainty and Learning in Pharmaceutical Demand.” *Econometrica*, 73(4), 1137-1173.
- [5] Darden, M. (2017): “Smoking, Expectations, and Health: A Dynamic Stochastic Model of Lifetime Smoking Behavior,” *Journal of Political Economy*, 125: 1245-1721.
- [6] Freedman, Laurence S. and David J. Spiegelhalter (1992): “Application of Bayesian Statistics to Decision Making During a Clinical Trial.” *Statistics in Medicine*, 11, 23-35.
- [7] Hauck, Katherine, and Tiemen Woutersen (2023): “The Homesteading Act and the Process of Learning to Farm in the Late 19th Century.”
- [8] Insua, David Rios, Fabrizio Ruggeri, Refik Soyer, and Simon Wilson (2020): “Advances in Bayesian decision making in reliability.” *European Journal of Operational Research*, 282(1); 1-18.
- [9] Kamenica, Emir and Matthew Gentzkow (2011): “Bayesian persuasion.” *American Economic Review*, 101(6), 2590-2615.
- [10] Müller, Peter, Don A. Berry, Andy P. Grieve, Michael Smith, and Michael Krams (2007): “Simulation-based sequential Bayesian design.” *Journal of Statistical Planning and Inference*, 137(2007), 3140–3150.
- [11] Zhongxiang Dai, Haibin Yu, Bryan Kian Hsiang Low, Patrick Jaillet (2019): “Bayesian Optimization Meets Bayesian Optimal Stopping,” Proceedings of the 36th International Conference on Machine Learning, PMLR 97:1496-1506, 2019.

A Appendix

As in the main text, consider the function

$$\begin{aligned}
 H(A, B, \tau_L) &= A + 1(A \leq \tau_L) \cdot 2 \cdot A + 1(A > \tau_L) \cdot \{B + \min(A, B)\} \\
 &= A + 1(A \leq \tau_L) \cdot 2 \cdot A + 1(A > \tau_L) \cdot B + 1(A > \tau_L) \cdot 1(B < A) \cdot B \\
 &\quad + 1(\tau_L < A < B) \cdot A.
 \end{aligned}$$

In this case, we can use the prior predictive distribution because the priors on α and β are independent. These prior predictive distributions are normal distributions with mean zero and variance two; let $f(a)$ and $f(b)$ denote the densities of these distributions. When evaluating $E\{H(A, B, \tau_L)\}$, note that $E(A) = E(B) = 0$ and $E(B|A) = 0$ for all A so $E\{1(A > \tau_L) \cdot B\} = 0$. Thus, two terms in the last equation have expectation zero. Further, note that

$$\begin{aligned}
 E\{1(A \leq \tau) \cdot A\} &= \int_{-\infty}^{\tau_L} a \cdot f(a) da = \int_{-\infty}^{\tau_L} a \frac{\exp\{-\frac{a^2}{4}\}}{\sqrt{2}\sqrt{2\pi}} d\alpha \\
 &= -\frac{\exp\{-\frac{a^2}{4}\}}{\sqrt{\pi}} \Big|_{-\infty}^{\tau_L} = -\frac{\exp\{-\frac{\tau_L^2}{4}\}}{\sqrt{\pi}}.
 \end{aligned}$$

Similarly, we can calculate the expectation of $1(A > \tau_L) \cdot 1(B < A) \cdot B$ by first calculating that an expectation conditional on A. In particular,

$$E\{1(B < A) \cdot B|A\} = -\frac{\exp\{-\frac{A^2}{4}\}}{\sqrt{\pi}}$$

$$E\{1(A > \tau_L) \cdot 1(B < A) \cdot B\} = E[1(A > \tau_L) \cdot E\{1(B < A) \cdot B|A\}]$$

$$\begin{aligned}
&= E\left[1(A > \tau_L) \cdot \left\{-\frac{\exp\left\{-\frac{A^2}{4}\right\}}{\sqrt{\pi}}\right\}\right] \\
&= -\int_{\tau_L}^{\infty} \frac{\exp\left\{-\frac{a^2}{4}\right\}}{\sqrt{\pi}} \frac{\exp\left\{-\frac{a^2}{4}\right\}}{\sqrt{2}\sqrt{2\pi}} da \\
&= -\frac{1}{\sqrt{2\pi}} \int_{\tau_L}^{\infty} \frac{\exp\left\{-\frac{a^2}{2}\right\}}{\sqrt{2\pi}} da \\
&= -\frac{1 - \Phi(\tau_L)}{\sqrt{2\pi}}.
\end{aligned}$$

Further, we calculate the expectation of $1(\tau_L < A < B) \cdot A$ by first conditioning on B and $B \geq \tau_L$. This conditional expectation is

$$\begin{aligned}
E\{1(\tau_L < A < B) \cdot A|B, B \geq \tau_L\} &= \int_{\tau_L}^B a \cdot f(a) da = \int_{\tau_L}^B a \frac{\exp\left\{-\frac{a^2}{4}\right\}}{\sqrt{2}\sqrt{2\pi}} d\alpha \\
&= -\frac{\exp\left\{-\frac{a^2}{4}\right\}}{\sqrt{\pi}} \Big|_{\tau_L}^B \\
&= -\frac{\exp\left\{-\frac{B^2}{4}\right\}}{\sqrt{\pi}} + \frac{\exp\left\{-\frac{\tau_L^2}{4}\right\}}{\sqrt{\pi}}.
\end{aligned}$$

This gives

$$E\{1(\tau_L < A < B) \cdot A|B\} = 1(B \geq \tau_L) \cdot \left[-\frac{\exp\left\{-\frac{B^2}{4}\right\}}{\sqrt{\pi}} + \frac{\exp\left\{-\frac{\tau_L^2}{4}\right\}}{\sqrt{\pi}}\right].$$

Next, consider the unconditional expectation,

$$\begin{aligned}
E\{1(\tau_L < A < B) \cdot A\} &= E[E\{1(\tau_L < A < B) \cdot A|B\}] \\
&= E\left[1(B \geq \tau_L) \cdot \left\{-\frac{\exp\left\{-\frac{B^2}{4}\right\}}{\sqrt{\pi}} + \frac{\exp\left\{-\frac{\tau_L^2}{4}\right\}}{\sqrt{\pi}}\right\}\right]
\end{aligned}$$

Note that we showed above that $E\left[1(A > \tau_L) \cdot \left\{-\frac{\exp\left\{-\frac{A^2}{4}\right\}}{\sqrt{\pi}}\right\}\right] = -\frac{1 - \Phi(\tau)}{\sqrt{2\pi}}$ so

$$E\left[1(B \geq \tau_L) \cdot \left\{-\frac{\exp\left\{-\frac{B^2}{4}\right\}}{\sqrt{\pi}}\right\}\right] = -\frac{1 - \Phi(\tau_L)}{\sqrt{2\pi}}.$$

Further,

$$E[1(B \geq \tau_L) \cdot \left\{ \frac{\exp\{-\frac{\tau_L^2}{4}\}}{\sqrt{\pi}} \right\}] = \{1 - \Phi(\tau_L/\sqrt{2})\} \left\{ \frac{\exp\{-\frac{\tau_L^2}{4}\}}{\sqrt{\pi}} \right\}.$$

Combining the last two equations yields

$$E\{1(\tau_L < A < B) \cdot A\} = -\frac{1 - \Phi(\tau_L)}{\sqrt{2\pi}} + \{1 - \Phi(\tau_L/\sqrt{2})\} \left\{ \frac{\exp\{-\frac{\tau_L^2}{4}\}}{\sqrt{\pi}} \right\}.$$

We can now consider the function $H(\tau_L) = E\{H(A, B, \tau_L)\}$,

$$\begin{aligned} H(\tau_L) &= E[1(A \leq \tau_L) \cdot 2 \cdot A + 1(A > \tau_L) \cdot \{B + \min(A, B)\}] \\ &= -2 \frac{\exp\{-\frac{\tau_L^2}{4}\}}{\sqrt{\pi}} - \frac{1 - \Phi(\tau_L)}{\sqrt{2\pi}} \\ &\quad - \frac{1 - \Phi(\tau_L)}{\sqrt{2\pi}} + \{1 - \Phi(\tau_L/\sqrt{2})\} \left\{ \frac{\exp\{-\frac{\tau_L^2}{4}\}}{\sqrt{\pi}} \right\} \\ &= -\frac{\exp\{-\frac{\tau_L^2}{4}\}}{\sqrt{\pi}} - 2 \frac{1 - \Phi(\tau_L)}{\sqrt{2\pi}} \\ &\quad - \Phi(\tau_L/\sqrt{2}) \left\{ \frac{\exp\{-\frac{\tau_L^2}{4}\}}{\sqrt{\pi}} \right\} \\ &= -\{1 + \Phi(\tau_L/\sqrt{2})\} \frac{\exp\{-\frac{\tau_L^2}{4}\}}{\sqrt{\pi}} - 2 \frac{1 - \Phi(\tau_L)}{\sqrt{2\pi}} \end{aligned}$$

The argument that minimizes $H(\tau_L)$ is approximately $\tau_L = -0.39$.