



Lecture 8

Introduction to Discrete Time Dynamic Optimization

Randall Romero Aguilar, PhD

This draft: November 7, 2018

Universidad de Costa Rica
SP6534 - Economía Computacional

Table of contents

1. Introduction
2. Markov Decision Models
3. Bellman Equation
4. Collocation Method
5. Discrete Action Examples

Introduction

- We now begin our study of dynamic economic models.
- Dynamic economic models often present complications rarely encountered together in other scientific models:
 1. humans are capable of assessing how their actions today will affect them in the future.
 2. many aspects of human behavior are unpredictable.
 3. the predictable component of human behavior is often complex.
- Thus, most useful dynamic models in economics are future-looking, stochastic, and highly nonlinear.

- The complications inherent in forward-looking, stochastic, nonlinear models make it impossible to obtain closed-form solutions to all but a small number of dynamic economic models.
- However, advances in computers now make it possible for economists to analyze a broad range of analytically insoluble dynamic models using numerical methods.

We examine three classes of discrete time dynamic economic models:

- dynamic decision making by a single economic agent
- dynamic decision making by multiple economic agents
- dynamic equilibrium models of decentralized decision making by many individuals

Examples of dynamic decision making by single economic agents involving **discrete** or **mixed** actions:

- timber stand owner deciding whether to clearcut his stand
- producer deciding whether to replace a physical asset
- capitalist deciding whether to enter or exit an industry
- unemployed worker deciding whether to accept a job offer
- financial investor deciding whether to exercise a put option
- borrower deciding whether to default on a loan
- agent deciding whether to make a bank transaction
- subsistence household deciding how many children to have

Examples of dynamic decision making by single economic agents involving pure **continuous** actions:

- individual deciding how much to consume and save
- planner deciding how much of a renewable resource to harvest
- mine owner deciding how much ore to extract
- reservoir authority deciding how much water to release
- central bank attempting to stabilize the economy
- entrepreneur planning production, investment, and inventory
- agricultural producer managing a livestock enterprise

Examples of dynamic decision making involving multiple economic agents:

- capacity expansion in an oligopolistic market
- smallholder income risk-sharing arrangements
- national grain marketing boards competing on world markets

Examples of dynamic decentralized economic behavior:

- asset returns in a pure exchange economy
- futures markets for primary commodities
- government commodity price stabilization program

Markov Decision Models

The **discrete time Markov decision model** has this structure:

- At the beginning of every period t , an agent observes the pre-determined state of an economic process s_t , takes an action x_t , and earns a **reward** $f_t(s_t, x_t)$ that depends on the state and action.
- The **state** of the economic process in period $t + 1$ will depend on the state and action taken in period t and a purely exogenous random shock ϵ_{t+1} that is realized after action is taken in period t :

$$s_{t+1} = g_t(s_t, x_t, \epsilon_{t+1}).$$

- The agent seeks to maximize the present value of current and expected future rewards over a time horizon T , discounted at a per-period rate ρ .

State Space

- The **state space** S enumerates the states attainable by the economic process.
- A **continuous state model** possesses state variables whose ranges are intervals of the real line.
- A **discrete state model** possesses state variables whose ranges are finite.
- A **mixed state model** possesses both continuous and discrete state variables.

Action Space

- The **action space** $X(s)$ enumerates the actions that may be taken by the agent when the economic process is in state s .
- A **continuous action model** possesses action variables whose ranges are intervals of the real line.
- A **discrete action model** possesses action variables whose ranges are finite.
- A **mixed action model** possesses both continuous and discrete action variables.

Reward Function

- The **reward function** $f : S \times X \mapsto \mathfrak{R}$ gives the reward earned in the current period in terms of the current state and action.

$$f_t(s_t, x_t)$$

- The **objective** is to maximize the present value of current and expected future rewards over a time horizon T , discounted at a per-period factor $\delta = \frac{1}{1+\rho}$.

$$V_t(s_t) = \max_{\{x_\tau\}_{\tau=t}^T} \mathbb{E}_t \left[\sum_{\tau=t}^T \delta^{\tau-t} f_\tau(s_\tau, x_\tau) \right]$$

- Notice that the maximum attainable **value** V_t depends on the state s_t .

Transition Function

- The **transition function** $g : S \times X \times \Omega \mapsto S$ gives next period's state in terms of the current state and action and possibly a random shock realized after the decision is made.

$$s_{t+1} = g_t(s_t, x_t, \epsilon_{t+1}).$$

- A discrete time Markov decision model may have a **finite horizon** ($T < \infty$) or an **infinite horizon** ($T = \infty$).
- If the horizon is finite, one must specify a **terminal value** describing a final reward or “salvage” value earned by the agent in period $T + 1$ as a function of the state in period $T + 1$, after the final decision is taken in period T .

Model Formulation

To fully specify a Markov decision model, one must clearly identify

- The state variables and state space
- The action variables and action space
- The reward function
- The transition function
- The time horizon and, if finite, the terminal value

Bellman Equation

Optimal Policy

Given a discrete time Markov decision model, an **optimal policy** is a sequence of rules $\{x_t^*\}$ that prescribe the action $x = x_t^*(s)$ that the agent must take in period t if the process is in state s in order to maximize the present value of current and expected future rewards.

Principle of Optimality

The discrete time Markov decision problem may be analyzed using **Bellman's Principle of Optimality**, which was articulated by Richard Bellman in 1957 as follows:

“An optimal policy has the property that, whatever the initial state and decision are, the remaining decisions must constitute an optimal policy with regard to the state resulting from the first decision.”

Dynamic Programming

- Bellman's Principle of Optimality motivates a strategy for solving dynamic decision models called **dynamic programming**
- Dynamic programming is superior to alternative approaches to dynamic optimization because, in a unified framework, it can readily handle
 - deterministic and stochastic models
 - discrete and continuous time models
 - discrete and continuous state models
 - discrete and continuous action models
 - models with constraints
- Dynamic programming introduces the concept of the **value function** and characterizes it as the solution to a functional equation known as **The Bellman Equation**.

The Value Function

- Let $V_t(s)$ denote the maximum attainable present value of current and expected future rewards, given the economic process is in state s in period t .
- We call $V_t : S \mapsto \mathfrak{R}$ the **period- t value function**.
- The value functions are unknown a priori and must be derived from the underlying model.

The Bellman Equation

The Principle of Optimality implies that the value functions V_t must satisfy the **Bellman equation**

$$V_t(s) = \max_{x \in X(s)} \{f_t(s, x) + \delta E_t V_{t+1}(g_t(s, x, \epsilon))\}, \quad s \in S.$$

where $\delta = \frac{1}{1+\rho}$ is called the **discount factor**.

The Bellman equation is derived from the objective function:

$$\begin{aligned}
 V_t(s_t) &= \max_{\{x_\tau\}_{\tau=t}^T} \mathbb{E}_t \left[\sum_{\tau=t}^T \delta^{\tau-t} f_\tau(s_\tau, x_\tau) \right] \\
 &= \max_{\{x_\tau\}_{\tau=t}^T} \mathbb{E}_t \left[f_t(s_t, x_t) + \sum_{\tau=t+1}^T \delta^{\tau-t} f_\tau(s_\tau, x_\tau) \right] \\
 &= \max_{\{x_\tau\}_{\tau=t}^T} \left\{ f_t(s_t, x_t) + \mathbb{E}_t \left[\sum_{\tau=t+1}^T \delta^{\tau-t} f_\tau(s_\tau, x_\tau) \right] \right\} \\
 &= \max_{\{x_\tau\}_{\tau=t}^T} \left\{ f_t(s_t, x_t) + \delta \mathbb{E}_t \mathbb{E}_{t+1} \left[\sum_{\tau=t+1}^T \delta^{\tau-(t+1)} f_\tau(s_\tau, x_\tau) \right] \right\} \\
 &= \max_{x_t} \left\{ f_t(s_t, x_t) + \delta \mathbb{E}_t \max_{\{x_\tau\}_{\tau=t+1}^T} \mathbb{E}_{t+1} \left[\sum_{\tau=t+1}^T \delta^{\tau-(t+1)} f_\tau(s_\tau, x_\tau) \right] \right\} \\
 &= \max_{x_t} \{ f_t(s_t, x_t) + \delta \mathbb{E}_t V_{t+1}(s_{t+1}) \} \\
 &= \max_{x_t} \{ f_t(s_t, x_t) + \delta \mathbb{E}_t V_{t+1}(g_t(s_t, x_t, \epsilon_{t+1})) \}
 \end{aligned}$$

- The Bellman equation is a **recursive functional equation**.
- It is a functional equation because the unknowns, the value functions V_t , are functions, not vectors in \mathbb{R}^n .
- It is recursive because V_t is defined in terms of V_{t+1} .
- The Bellman equation compactly captures the fundamental tradeoff that a rational, future-regarding, dynamically optimizing agent must make between immediate rewards $f_t(s_t, x_t)$ and expected future rewards $\delta E_t V_{t+1}(s_{t+1})$.

Solving the Bellman equation, finite-horizon problem

- Because the Bellman equation is **recursive**, we need to know V_{t+1} in order to find V_t .
- But in the **finite-horizon** problem, following the recursion:

$$V_t(s_t) = \max_{x_t \in X(s_t)} \{f_t(s_t, x_t) + \delta E_t V_{t+1}(s_{t+1})\}$$

$$V_{t+1}(s_{t+1}) = \max_{x_{t+1} \in X(s_{t+1})} \{f_{t+1}(s_{t+1}, x_{t+1}) + \delta E_{t+1} V_{t+2}(s_{t+2})\}$$

⋮

$$V_{T-1}(s_{T-1}) = \max_{x_{T-1} \in X(s_{T-1})} \{f_{T-1}(s_{T-1}, x_{T-1}) + \delta E_{T-1} V_T(s_T)\}$$

$$V_T(s_T) = \max_{x_T \in X(s_T)} \{f_T(s_T, x_T) + \delta E_T V_{T+1}(s_{T+1})\}$$

once we have specified a **terminal** value V_{T+1} , we can solve for V_t by **backward induction**.

Stationarity

- An infinite-horizon discrete time Markov decision model is **stationary** if the reward function, transition function, and shock distributions are independent of time t .
- An infinite-horizon discrete time Markov decision model is stationary, neither the value function nor the optimal policy will depend on time t .
- In this case, the Bellman equation takes the form of a **functional fixed-point equation** whose single unknown is the stationary value function V :

$$V(s) = \max_{x \in X(s)} \{f(s, x) + \delta E_{\epsilon} V(g(s, x, \epsilon))\}, \quad s \in S.$$

- It is a functional fixed-point equation because it defines the unknown value function V in terms of itself.

Existence of Solution

By the Contraction Mapping Theorem, the Bellman equation of a stationary infinite-horizon discrete time Markov decision model possesses a unique solution if the discount factor δ is less than one and the reward function f is bounded.

Collocation Method

- We now show how to solve Bellman equations via collocation.
- For simplicity, we limit our introductory discussion to infinite-horizon models with a continuous one-dimensional state space.
- However, the strategy presented generalizes to finite-horizon models and more general state spaces.

- To compute an approximate solution to the Bellman equation via collocation:
- Approximate the value function by a linear combination of n known basis functions $\phi_1, \phi_2, \dots, \phi_n$

$$V(s) \approx \sum_{j=1}^n c_j \phi_j(s)$$

and fix the basis function coefficients c_1, c_2, \dots, c_n by requiring the value function approximant to satisfy the Bellman equation, not at all possible states, but rather at n judiciously chosen collocation nodes s_1, s_2, \dots, s_n in S .

- Collocation replaces the Bellman equation with a system of n nonlinear equations

$$\sum_{j=1}^n c_j \phi_j(s_i) = \max_x \left\{ f(s_i, x) + \delta E_\epsilon \sum_{j=1}^n c_j \phi_j(g(s_i, x, \epsilon)) \right\}$$

$V(s_i)$ $V[g(s_i, x, \epsilon)]$

$i = 1, 2, \dots, n$ in n unknowns, $c_j, j = 1, 2, \dots, n$.

- The nonlinear equation system may be compactly expressed in vector form as the **collocation equation**

$$\Phi c = v(c)$$

whose unknown is the basis coefficient vector c .

- Here, Φ is the $n \times n$ matrix whose typical element

$$\Phi_{ij} = \phi_j(s_i)$$

is the j^{th} basis function evaluated at the i^{th} node.

- And v is the function from \mathfrak{R}^n to \mathfrak{R}^n whose typical element

$$v_i(c) = \max_{x \in X(s_i)} \{f(s_i, x) + \delta E_\epsilon \sum_{j=1}^n c_j \phi_j(g(s_i, x, \epsilon))\},$$

is the optimal value of the maximization problem embedded in the Bellman equation at the i^{th} collocation node, replacing the value function with its approximant.

- The collocation equation may be solved using a variety of nonlinear equation solution methods.
- For example, one may write the collocation equation as a **fixed-point problem**

$$c = \Phi^{-1}v(c)$$

and solve for c using function iteration

$$c \leftarrow \Phi^{-1}v(c).$$

- Alternatively, one may write the collocation equation as a **rootfinding problem**

$$\Phi c - v(c) = 0$$

and solve for c using Newton's method

$$c \leftarrow c - [\Phi - v'(c)]^{-1}[\Phi c - v(c)].$$

- Here, $v'(c)$ is the $n \times n$ Jacobian of v at c , whose typical element may be computed using the Envelope Theorem:

$$v'_{ij}(c) = \frac{\partial v_i}{\partial c_j}(c) = \delta E_\epsilon \phi_j(g(s_i, x_i, \epsilon)).$$

- Can also use Broyden's method instead of Newton's method.

Evaluating the accuracy of the solution

- The accuracy of the computed approximate solution is assessed by inspecting the **residual**.
- For a given basis coefficient vector c , the residual is a function of the state variable s :

$$R_c(s) \equiv \sum_{j=1}^n c_j \phi_j(s) - \max_x \{ f(s, x) + \delta E_\epsilon \sum_{j=1}^n c_j \phi_j(g(s, x, \epsilon)) \}$$

- If the approximant were exact, the residual would be zero everywhere.
- However, in general, the residual will be nonzero, except at the collocation nodes, where it is zero by construction.
- The accuracy of the approximant is measured by how much the residual differs from zero at non nodal states.

Easier said than done...

- Discussion above is general, but obscures some practical issues:
 - First, the state and action spaces could be **multi-dimensional** and involve a mixture of continuous and discrete variables.
 - Second, evaluation of v requires a finite-dimensional **nonlinear maximization** operation, about which little has been said.
 - Third, solving the collocation equation requires an **initial guess** for the unknown coefficients.
- Mixed multi-dimensional state and action spaces and finite-dimensional nonlinear maximization are handled in a straightforward fashion by the CompEcon Toolbox.
- But algebra is cumbersome and requires multi-dimensional indexing, which we have chosen to avoid for clarity.

Practical decisions

- When applying the collocation method, the analyst faces a number of practical decisions:
 - First, choose **basis functions** and **collocation nodes**.
 - Second, choose **algorithm** for solving the collocation equation.
 - Third, choose numerical **quadrature technique**, if needed.
 - Fourth, provide an **initial guess** for the basis coefficients.
- A careful analyst may try different basis-node families, different degrees of discretization, different solution algorithms, and different initial guesses for the coefficients to assure robustness of results.
- A careful analyst always, always, always checks the residual!

The curse of dimensionality in the collocation method

- The collocation method suffers from the **curse of dimensionality**.
- Specifically, the computational effort required grows exponentially with the dimension of the state space.
- For example, if the state space has d continuous dimensions and you choose n basis functions for each state variable dimension, then the collocation equation has n^d single equations and the same number of unknowns.
- For this reason, in practice, the main limitation to solving dynamic optimization problems numerically is keeping the state dimension manageable.

Discrete Action Examples

MODEL 1: Timber Harvesting

- At the beginning of each period t , a timber stand owner observes the biomass of his stand s_t and must decide whether to harvest and sell it, or allow it to grow for another period.
- The revenue received per unit of harvested biomass is a constant p and the cost of clearcutting and replanting is a constant κ .
- A stand with biomass s_t that is not clearcut in period t will have biomass $s_{t+1} = h(s_t)$ the following period; a stand that is clearcut will have biomass $s_{t+1} = h(0) > 0$.
- What clearcutting policy maximizes the value of the stand?




Figure 1: Biomass Growth Function

Formulation

This is an infinite-horizon, deterministic model with the following structural features:

- One continuous state variable, stand biomass

$$s_t \in [0, \bar{s}],$$

where \bar{s} is the maximum attainable biomass.

- One binary action variable, the clearcut decision

$$j_t = \begin{cases} 1, & \text{clearcut stand} \\ 0, & \text{otherwise.} \end{cases}$$

- The reward is current profit

$$\pi_t = \begin{cases} 0, & \text{if } j_t = 0 \\ ps_t - \kappa, & \text{if } j_t = 1 \end{cases}$$

where $p > 0$ and $\kappa > 0$.

- State transitions are governed by

$$s_{t+1} = \begin{cases} h(s_t), & \text{if } j_t = 0 \\ h(0), & \text{if } j_t = 1 \end{cases}$$

where $h(0) > 0$, $h' > 0$ and $h'' < 0$ on $s \in [0, \bar{s}]$, $h(\bar{s}) = \bar{s}$.

Bellman Equation

The value of a stand with biomass s satisfies the Bellman equation

$$V(s) = \max\left\{ \underset{\text{doesn't clearcut}}{\delta V[h(s)]}, \quad \underset{\text{clearcuts}}{ps - \kappa + \delta V[h(0)]} \right\}$$

Here,

- If owner does not clearcut the stand, she earns no profit and begins the following period with a stand of biomass $h(s)$ worth $V[h(s)]$.
- If owner clearcuts the stand, she earns profit $ps - \kappa$ and begins the following period with a stand of biomass $h(0)$ worth $V[h(0)]$.

Alternatively, the Bellman equation can be written

$$V(s) = \max\left\{ \underset{\text{doesn't clearcut}}{V_0(s)}, \underset{\text{clearcuts}}{V_1(s)} \right\}$$

where

$$V_0(s) \equiv \delta V[h(s)]$$

$$V_1(s) \equiv ps - \kappa + \delta V[h(0)]$$

are the **action-contingent value functions**, which give the values contingent on not clearcutting and clearcutting, respectively.

- The stand owner clearcuts the stand if

$$V_1(s) > V_0(s)$$

or, equivalently, if

$$ps - \kappa > \delta \{V[h(s)] - V[h(0)]\}.$$

- That is, the owner clearcuts and replants if the net revenue this period exceeds the capital loss in stand value.
- A myopic owner that clearcuts simply if net revenue is positive will clear cut sooner than he should.
- The critical biomass s^* at which a stand should be clearcut is characterized by $V_0(s^*) = V_1(s^*)$.

- The collocation method calls for the value function to be approximated by a linear combination of n basis functions ϕ_k :

$$V(s) \approx \hat{V}(s) \equiv \sum_{k=1}^n c_k \phi_k(s).$$

- The n coefficients c_k are fixed by requiring the value function approximant to satisfy the Bellman equation at n judiciously chosen nodes s_i .

This requires solving the n collocation equations

$$\sum_{k=1}^n c_k \phi_k(s_i) = \max_{V(s_i)} \left\{ \begin{array}{l} \delta \sum_{k=1}^n c_k \phi_k(h(s_i)), \quad \text{doesn't clearcut} \\ p s_i - \kappa + \delta \sum_{k=1}^n c_k \phi_k(h(0)) \quad \text{clearcuts} \end{array} \right\},$$

$i = 1, 2, \dots, n$, for the n unknown basis function coefficients c_k ,
 $k = 1, 2, \dots, n$.

Example 1:

Timber Harvesting, linear
approximation

- Let us solve the model numerically for the following parameterization:

- $h(s) = s + \gamma(\bar{s} - s), \bar{s} = 0.5, \gamma = 0.1$

- $p = 1$

- $\kappa = 0.2$

- $\delta = 0.9$

Computing a Linear Approximation by “Hand”

- To better understand the collocation method, let us begin by constructing a linear approximation to the value function:

$$V(s) \approx \hat{V}(s) \equiv c_1 + c_2 s.$$

- Given two collocation nodes s_1 and s_2 , the two collocation equations become

$$\begin{aligned} c_1 + c_2 s_1 &= \max\{\delta(c_1 + c_2 h(s_1)), \quad p s_1 - \kappa + \delta(c_1 + c_2 h(0))\} \\ c_1 + c_2 s_2 &= \max\{\delta(c_1 + c_2 h(s_2)), \quad p s_2 - \kappa + \delta(c_1 + c_2 h(0))\}. \end{aligned}$$

- If we choose $s_1 = 0.2$ and $s_2 = 0.4$, insert all parameter values, and simplify, then

$$c_1 + 0.2c_2 = \max\{0.9c_1 + 0.207c_2, \quad 0.0 + 0.9c_1 + 0.225c_2\}$$
$$c_1 + 0.4c_2 = \max\{0.9c_1 + 0.369c_2, \quad 0.3 + 0.9c_1 + 0.405c_2\}.$$

- Although cumbersome, these two nonlinear equations can be solved explicitly for c_1 and c_2 , yielding the value function approximant

$$\hat{V}(s) \equiv 0.0387 + 0.5525s.$$

- One can then derive approximants for the conditional value functions

$$\hat{V}_0(s) = \delta \hat{V}(h(s)) \quad = -0.0597 + 0.4475s$$

$$\hat{V}_1(s) = ps - \kappa + \delta \hat{V}(h(0)) \quad = -0.1403 + s.$$

- Equating the two yields the critical biomass level $s^* = 0.36$ at which a stand should be clearcut.

- Open the Jupyter notebook [dp/01a Timber Harvesting--2 nodes](#), which already contains the model parameters:

```
price = 1.0 # price of biomass
kappa = 0.2 # clearcut-replant cost
smax  = 0.5 # stand carrying capacity
gamma = 0.1 # biomass growth parameter
delta = 0.9 # discount factor
```

- The script also already contains some code that generates graphs and other output.

To solve the collocation equation in Python,

Step 1: code the growth function

```
def h(s): return s + gamma*(smax - s)
```

Step 2: code the value function approximant

```
def vhat(c, s): return c[0] + c[1]*s
```

Step 3: code the conditional value function approximants

```
def vhat1(c,s):  
    return price*s - kappa + delta * vhat(c,h(0))  
def vhat0(c,s):  
    return delta * vhat(c, h(s))
```

Step 4: and code the residual

```
def resid(c,s=snodes):  
    return vhat(c,s) - np.maximum(vhat0(c,s),vhat1(c,s))
```

Step 5: Solve the collocation equation with two collocation nodes

```
snodes = np.array([0.2, 0.4])  
cc = NLP(resid).broyden(np.zeros(2))
```

Step 6: Compute the critical biomass

```
scrit = NLP(lambda s: vhat0(cc,s)-vhat1(cc,s)). \  
    broyden(0.0)[0]
```

Conditional Value Functions

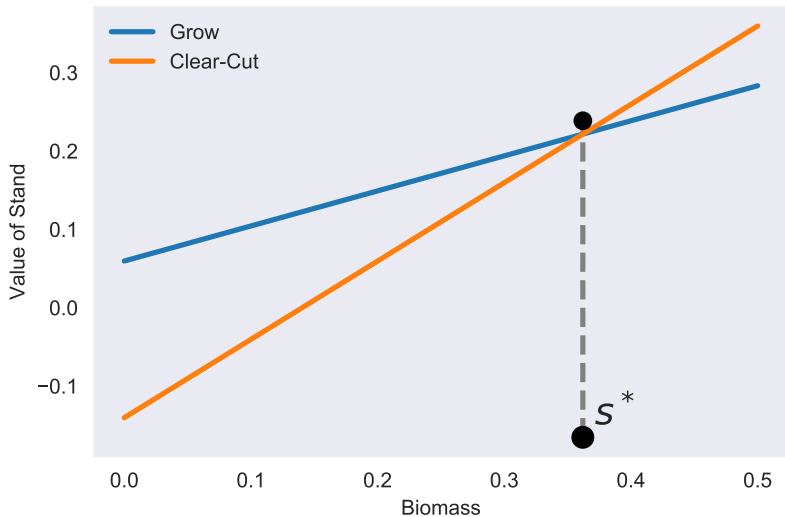


Figure 2: Conditional Value Functions and Critical Biomass, 2 Basis Functions




Figure 3: Bellman Equation Residual, 2 Basis Functions

- The Bellman equation residual, expressed as a percent of the underlying value, reaches magnitudes of around 35 percent, which is unacceptably high.

Example 2:

Timber harvesting, cubic spline
approximation

Let us now open the Jupyter notebook [dp/01b Timber Harvesting--cubic spline](#) so that we solve the model using 200 cubic spline basis functions.

Step 1: code the growth function

```
def h(s): return s + gamma*(smax - s)
```

Step 2: code the value function approximant

```
ns = 200
vhat = BasisSpline(ns,0,smax,k=3)
```

Step 3: code the conditional value function approximants

```
def vhat1(s):
    return price*s - kappa + delta * vhat(h(0))
def vhat0(s):
    return delta * vhat(h(s))
```

Step 4: and code the residual

```
def resid(c,s=vhat.nodes):  
    vhat.c = c  
    return vhat(s) - np.maximum(vhat0(s), vhat1(s))
```

Step 5: Solve the collocation equation

```
cc = NLP(resid).broyden(vhat.c)
```

Step 6: Compute the critical biomass

```
scrit = NLP(lambda s: vhat0(s)-vhat1(s)).broyden(0.0)[0]
```


- The Bellman equation residual using 200 cubic spline basis functions reaches magnitudes of around 0.02 percent, four orders of magnitude better than with two basis functions.
- The computed critical biomass level at which a stand should be clearcut is $s^* = 0.31$.

Conditional Value Functions

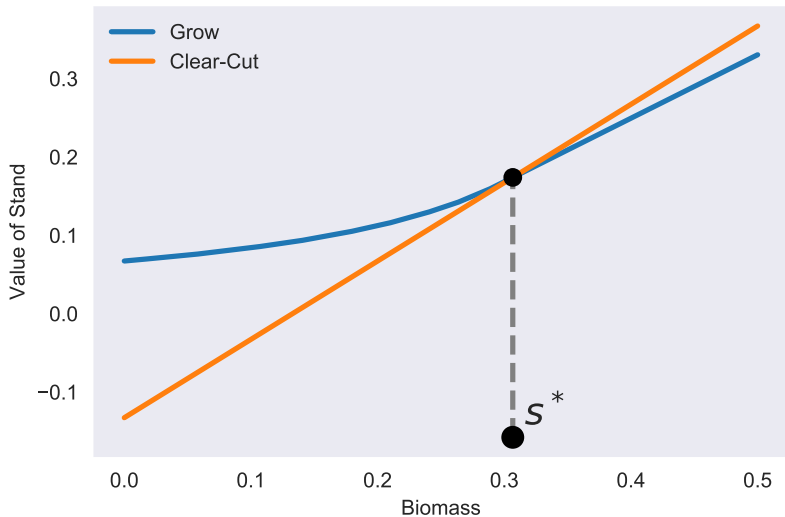


Figure 4: Conditional Value Functions and Critical Biomass, 200 Basis Functions




Figure 5: Bellman Equation Residual, 200 Basis Functions

Parametric Analysis

- The critical biomass level and rotation cycle are 0.31 and 10 periods, respectively.
- How do these values change if..
 - market price falls by half? 0.44, 21.
 - unit tax of 0.2 is imposed on sale of timber? 0.35, 12.
 - cost of clearcutting and replanting falls by half? 0.21, 6.
 - tax of 0.1 is imposed on clearcutting? 0.38, 14.
- How would you adapt the model and code to allow for a probability q that a fire could destroy the timber stand each period?

MODEL 2: Asset Replacement

- At the beginning of each period t , a manufacturer observes the age of an asset a_t and must decide whether to replace it with a new one.
- An asset that is a periods old produces $q(a)$ units of output.
- Safety regulations require that the asset be replaced once it reaches an age of A periods.
- The net cost of liquidating an old asset and purchasing a new one is a constant κ .
- The asset earns a net profit p_t per unit of output in period t , which evolves according to $p_{t+1} = h(p_t, \epsilon_{t+1})$ where the ϵ are serially i.i.d. shocks.
- What replacement policy maximizes firm profits?




Figure 6: Asset Productivity

Formulation

This is an infinite-horizon, stochastic model with the following structural features:

- One continuous state variable, net unit profit

$$p_t \in (0, \infty),$$

and one discrete state variable, asset age

$$a_t \in \{1, 2, 3, \dots, A\}.$$

- One binary action variable, the replacement decision

$$j_t = \begin{cases} 1 & \text{replace asset} \\ 0 & \text{otherwise,} \end{cases}$$

which is subject to the constraint that $j_t = 1$ if $a_t = A$.

- The reward is current profit

$$\pi_t = \begin{cases} p_t q(a_t), & \text{if } j_t = 0 \\ p_t q(0) - \kappa, & \text{if } j_t = 1. \end{cases}$$

- State transitions are governed by

$$p_{t+1} = h(p_t, \epsilon_{t+1})$$
$$a_{t+1} = \begin{cases} a_t + 1, & \text{if } j_t = 0 \\ 1, & \text{if } j_t = 1. \end{cases}$$

Bellman Equation

The value of owning an asset of age a , given net unit profit p , satisfies the Bellman equation

$$V_a(p) = \max \left\{ \underset{\text{keep it}}{V_{a0}(p)}, \quad \underset{\text{replace it}}{V_{a1}(p)} \right\}$$

where

$$V_{a0}(p) \equiv pq(a) + \delta E_\epsilon V_{a+1}(h(p, \epsilon))$$

$$V_{a1}(p) \equiv pq(0) - \kappa + \delta E_\epsilon V_1(h(p, \epsilon))$$

are the values contingent on keeping and replacing an asset of age a , respectively.

- If the manufacturer keeps an asset of age a , he earns $pq(a)$ and begins the following period with an asset that is one period older and worth $V_{a+1}(\tilde{p})$, where \tilde{p} is the following period's net unit profit.
- If he replaces the asset, on the other hand, he earns $pq(0) - \kappa$ and begins the following period with an asset that is one period old and worth $V_1(\tilde{p})$.

- The manufacturer replaces the asset if

$$p(q(0) - q(a)) - \kappa > \delta E_{\epsilon} (V_{a+1}(\tilde{p}) - V_1(\tilde{p})),$$

that is, if the net revenue this period exceeds the expected capital loss.

- A myopic manufacturer that replaces the asset simply if net revenue is positive will replace too soon or too late, depending on whether the expected capital loss is positive or negative.
- The critical unit profit contribution p_a^* at which an asset of age a is replaced is characterized by $V_{a0}(p_a^*) = V_{a1}(p_a^*)$.

- The collocation method calls for the value function to be approximated by a linear combination of n basis functions ϕ_k :

$$V_a(p) \approx \hat{V}_a(p) \equiv \sum_{k=1}^n c_{ak} \phi_k(p).$$

- The $A \times n$ coefficients c_{ak} are fixed by requiring the value function approximants to satisfy the Bellman equation at n judiciously chosen nodes p_i .

This requires solving the $A \times n$ nonlinear collocation equations

$$\hat{V}_a(p_i) \equiv \max\{\underset{\text{keep it}}{\hat{V}_{a0}(p_i)}, \underset{\text{replace it}}{\hat{V}_{a1}(p_i)}\},$$

$a = 1, 2, \dots, A, i = 1, 2, \dots, n$, where

$$\hat{V}_{a0}(p_i) = p_i q(a) + \delta E_\epsilon \hat{V}_{a+1}(h(p_i, \epsilon))$$

$$\hat{V}_{a1}(p_i) = p_i q(0) - \kappa + \delta E_\epsilon \hat{V}_1(h(p_i, \epsilon))$$

for the $A \times n$ basis coefficients c_{ak} , $a = 1, 2, \dots, A$,
 $k = 1, 2, \dots, n$.

Example 3:

Asset replacement

- Open and run CompEcon demo program [dp/02 Asset Replacement Model](#).
- This demo solves the asset replacement model assuming
 - $q(a) = \alpha_0 + \alpha_1 a + \alpha_2 a^2$, $\alpha_0 = 50$, $\alpha_1 = -2.5$, $\alpha_2 = -2.5$;
 - $h(p, \epsilon) = \bar{p} + \gamma(p - \bar{p}) + \epsilon$, $\bar{p} = 1$, $\gamma = 0.5$;
 - ϵ serially i.i.d. normal, mean 0, standard deviation $\sigma = 0.15$;
 - $A = 6$, $\kappa = 40$, $\delta = 0.9$.
- The value function is approximated by a linear combination of $n = 200$ cubic spline basis functions defined on the interval $[0, 2]$.
- The unit profit contribution shock ϵ is discretized using an $m = 5$ node Gaussian quadrature scheme.

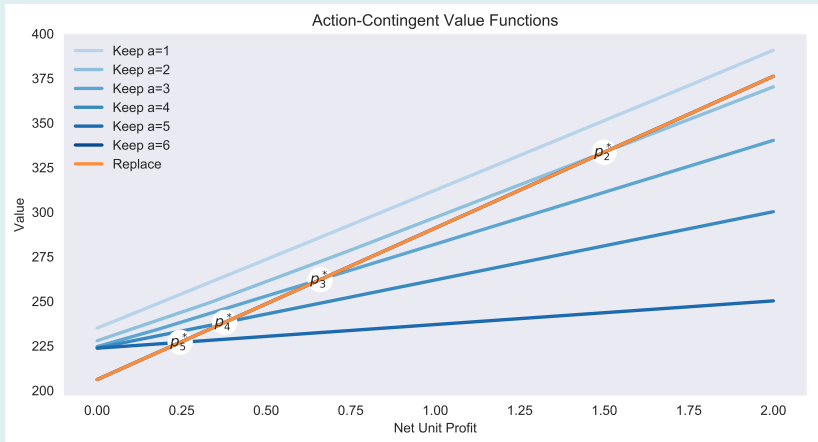


Figure 7: Conditional Value Functions and Critical Unit Profit Contributions, by Asset Age

- The critical unit profit contributions at which the asset should be replaced are 1.50, 0.66, 0.38, and 0.25, for assets of age 2, 3, 4, and 5, respectively (an asset of age 1 is never replaced).
- The critical unit profit contributions p_a^* were computed numerically by solving $\hat{V}_{a0}(p_a^*) = \hat{V}_{a1}(p_a^*)$.

Bellman Equation Residual




Figure 8: Bellman Equation Residuals

Parametric Analysis

- The average asset age over the first fifty periods is 2.01.
- How does this value change if...
 - replacement cost falls by half? 1.61.
 - longrun mean profit contribution rises by 20%? 1.99.
 - manufacturer discounts the future less ($\delta = 0.95$)? 2.00.
- How would you adapt the model and code to allow for investing in maintenance that enhances asset productivity?

MODEL 3: Industry Entry and Exit

- At the beginning of each period t , a firm observes its potential profit over the current period p_t , which may be negative, and must decide whether to operate.
- The firm faces no fixed costs or shutdown costs, but incurs a startup cost κ if it reopens after a period of inactivity.
- Short-run potential profit evolves according to $p_{t+1} = h(p_t, \epsilon_{t+1})$ where the ϵ are serially i.i.d shocks.
- What entry-exit policy maximizes the value of the firm?

Formulation

This is an infinite-horizon, stochastic model with the following structural features:

- One continuous state variable, potential profit

$$p_t \in (-\infty, \infty),$$

and one binary state variable, the operational status of the firm last period

$$i_t = \begin{cases} 1 & \text{active} \\ 0 & \text{idle.} \end{cases}$$

- One binary action variable, the operational status of the firm this period

$$j_t = \begin{cases} 1 & \text{active} \\ 0 & \text{idle.} \end{cases}$$

- The reward is current net profit

$$\pi_t = p_t j_t - \kappa(1 - i_t) j_t.$$

- State transitions are governed by

$$p_{t+1} = h(p_t, \epsilon_{t+1})$$

$$i_{t+1} = j_t.$$

Bellman Equation

The value of the firm, given operational status i and potential profit p , satisfies the Bellman equation

$$V_i(p) = \max\{V_{i0}(p), V_{i1}(p)\}$$

idle active

where

$$V_{i0}(p) \equiv \delta E_\epsilon V_0(h(p, \epsilon))$$

$$V_{i1}(p) \equiv p - \kappa(1 - i) + \delta E_\epsilon V_1(h(p, \epsilon))$$

are the values contingent on being idle and active this period, respectively.

- An idle firm earns no profit and begins the following period with a value $V_0(\tilde{p})$, where \tilde{p} is the following period's potential profit.
- An active firm, on the other hand, earns profit p , incurring a startup cost κ if it was idle the preceding period, and begins the following period with a value $V_1(\tilde{p})$.
- The critical potential profits at which an idle firm reopens, p_0^* , and at which an active firm shuts down, p_1^* , are characterized by $V_{i0}(p_i^*) = V_{i1}(p_i^*)$, for $i = 0, 1$.

- An idle firm re-opens if

$$p + V^*(p) > \kappa,$$

and an active firm remains open if

$$p + V^*(p) > 0,$$

where

$$V^*(p) \equiv \delta E_\epsilon (V_1(h(p, \epsilon)) - V_0(h(p, \epsilon))) > 0,$$

is the **option value** of being active.

- Since $V^*(p) > 0$, an idle firm may reopen and an active firm may remain open, even though it loses money in the shortcut.

- The collocation method calls for the value function to be approximated by a linear combination of n basis functions ϕ_k :

$$V_i(p) \approx \hat{V}_i(p) \equiv \sum_{k=1}^n c_{ik} \phi_k(p).$$

- The $2 \times n$ coefficients c_{ik} are fixed by requiring the value function approximants to satisfy the Bellman equation at n judiciously chosen nodes p_k .

This requires solving the $2 \times n$ nonlinear collocation equations

$$\hat{V}_i(p_k) \equiv \max\{\underbrace{\hat{V}_{i0}(p_k)}_{\text{idle}}, \underbrace{\hat{V}_{i1}(p_k)}_{\text{active}}\},$$

$i = 0, 1, k = 1, 2, \dots, n$, where

$$\hat{V}_{i0}(p_k) = \delta E_\epsilon \hat{V}_0(h(p_k, \epsilon))$$

$$\hat{V}_{i1}(p_k) = p_k - \kappa(1 - i) + \delta E_\epsilon \hat{V}_1(h(p_k, \epsilon)),$$

for the $2 \times n$ basis coefficients $c_{ik}, i = 0, 1, k = 1, 2, \dots, n$.

Example 4:

Industry entry and exit

- Open and run CompEcon demo program [dp/03 Industry Entry-Exit Model](#).
- This demo solves the industry entry-exit model assuming
 - $h(p, \epsilon) = \bar{p} + \gamma(p - \bar{p}) + \epsilon$, $\bar{p} = 1$, $\gamma = 0.7$;
 - ϵ serially i.i.d. normal, zero mean, standard deviation $\sigma = 1$;
 - $\kappa = 10$, $\delta = 0.9$.
- The value function is approximated by a linear combination of $n = 250$ cubic spline basis functions on $[-20, 20]$.
- The profit shock ϵ is discretized using an $m = 5$ node Gaussian quadrature scheme.

An idle firm reopens at $p_0^* = 2.10$ and an active firm shuts down at $p_1^* = -2.30$.

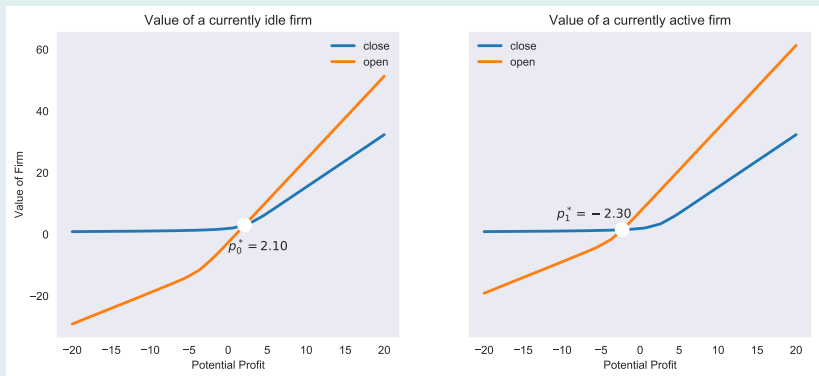


Figure 9: Conditional Value Functions and Critical Short-Run Potential Profits

The ergodic probability of operation, that is, the percentage of periods the firm is active in the longrun, is 92%.




Figure 10: Probability of operation




Figure 11: Bellman Equation Residuals

- How does the ergodic probability of operation change if...
 - longrun mean profit falls in half? 51%.
 - cost of reopening an idle firm doubles? 89%.
 - standard deviation of the profit shock doubles? 65%.
- How would you adapt the model and code to allow for shutdown costs and fixed costs that depend on whether firm is active or inactive?

MODEL 4: Job Search

- At the beginning of each period t , an infinitely-lived worker is either employed or unemployed.
- If employed, he must decide whether to continue to work at the prevailing wage w_t or quit and be idle.
- If unemployed, he must decide whether to search for a job and receive an unemployment benefit u or be idle.
- An idle worker receives neither a wage nor an unemployment benefit, but enjoys a pure leisure benefit v .

(cont)

- An unemployed worker who searches for a job will find one for the following period with probability p_0 .
- An employed worker who works will keep his job for the following period with probability p_1 .
- The wage evolves according to $w_{t+1} = h(w_t, \epsilon_{t+1})$ where the ϵ are serially i.i.d shocks.
- Under what conditions will an employed worker quit and will an unemployed worker search?

Formulation

This is an infinite-horizon, stochastic model with the following structural features:

- One continuous state variable, the prevailing wage

$$w_t \in [0, \infty),$$

and one binary state variable, the worker's employment status

$$i_t = \begin{cases} 1 & \text{employed} \\ 0 & \text{unemployed.} \end{cases}$$

- One binary action variable, the activity of the worker this period

$$j_t = \begin{cases} 1 & \text{active (work or search)} \\ 0 & \text{idle.} \end{cases}$$

- The reward is current benefits

$$\pi_t = w_t i_t j_t + u(1 - i_t)j_t + v(1 - j_t).$$

- State transitions are governed by

$$w_{t+1} = h(w_t, \epsilon_{t+1})$$

$$i_{t+1} = \begin{cases} 0 \text{ with probability } 1, & \text{if } j_t = 0 \\ 1 \text{ with probability } p_0, & \text{if } j_t = 1, i_t = 0 \\ 0 \text{ with probability } 1 - p_0, & \text{if } j_t = 1, i_t = 0 \\ 1 \text{ with probability } p_1, & \text{if } j_t = 1, i_t = 1 \\ 0 \text{ with probability } 1 - p_1, & \text{if } j_t = 1, i_t = 1 \end{cases}$$

Bellman Equation

The worker's expected lifetime benefits, given his employment status i and prevailing wage w , satisfies the Bellman equation

$$V_i(w) = \max\{V_{i0}(w), V_{i1}(w)\}$$

idle active

where

$$V_{i0}(w) \equiv v + \delta \mathbb{E}_\epsilon V_0(h(w, \epsilon))$$

$$V_{i1}(w, s) \equiv wi + u(1 - i) + \delta E_\epsilon \left[p_i V_1(h(w, \epsilon)) + (1 - p_i) V_0(h(w, \epsilon)) \right]$$

are the values contingent on being idle and active, respectively.

- An idle worker earns only leisure benefits and begins the following period with lifetime prospects $V_0(\tilde{w})$, where \tilde{w} is the following period's prevailing wage.
- An unemployed active worker earns an unemployment benefit u and begins the following period with lifetime prospects $V_0(\tilde{w})$ or $V_1(\tilde{w})$ with probabilities $1 - p_0$ and p_0 , respectively.
- An employed active worker earns a wage w and begins the following period with lifetime prospects $V_0(\tilde{w})$ or $V_1(\tilde{w})$ with probabilities $1 - p_1$ and p_1 , respectively.

- The critical wage w_0^* at which an unemployed worker searches is characterized by the condition that $V_{00}(w_0^*) = V_{01}(w_0^*)$.
- The critical wage w_1^* at which an employed worker quits is characterized by the condition that $V_{10}(w_1^*) = V_{11}(w_1^*)$.

- An unemployed worker searches for a job if

$$u + V_0^*(w) > v$$

and an employed worker continues to work if

$$w + V_1^*(w) > v$$

where

$$V_i^*(w) \equiv \delta p_i E_\epsilon (V_1(h(w, \epsilon))) - V_0(h(w, \epsilon)),$$

are the **option values** of being active for an unemployed worker, $i = 0$, or an employed worker, $i = 1$.

- Since $V_i^*(w) > 0$, an unemployed worker may search for a job, even though his leisure benefit exceeds his unemployment benefit, and an employed worker may continue working, even though his leisure benefit exceeds his wage.

- The collocation method calls for the value function to be approximated by a linear combination of n basis functions ϕ_k :

$$V_i(w) \approx \hat{V}_i(w) \equiv \sum_{k=1}^n c_{ik} \phi_k(w).$$

- The $2 \times n$ coefficients c_{ik} are fixed by requiring the value function approximants to satisfy the Bellman equation at n judiciously chosen nodes w_l .

This requires solving the $2 \times n$ nonlinear collocation equations

$$\hat{V}_i(w_l) \equiv \max\{\hat{V}_{i0}(w_l), \hat{V}_{i1}(w_l)\},$$

$i = 0, 1, l = 1, 2, \dots, n$, where

$$\hat{V}_{i0}(w_l) = v + \delta E_\epsilon \hat{V}_0(h(w_l, \epsilon))$$

$$\hat{V}_{i1}(w_l) = w_l i + u(1 - i) + \delta E_\epsilon \left[p_i \hat{V}_1(h(w_l, \epsilon)) + (1 - p_i) \hat{V}_0(h(w_l, \epsilon)) \right],$$

for the $2 \times n$ basis coefficients $c_{ik}, i = 0, 1, k = 1, 2, \dots, n$.

Example 5:

Job search

- Open and run CompEcon demo program [dp/04 Job Search Model](#).
- This demo solves the job search model assuming
 - $h(w, \epsilon) = \bar{w} + \gamma(w - \bar{w}) + \epsilon$, $\bar{w} = 100$, $\gamma = 0.4$;
 - ϵ are serially i.i.d. normal, mean 0, standard deviation $\sigma = 5$;
 - $u = 90$, $v = 95$, $p_0 = 0.2$, $p_1 = 0.9$, $\delta = 0.95$.
- The value function is approximated by a linear combination of $n = 150$ cubic spline basis functions on $[0, 200]$.
- The wage shock ϵ is discretized using an $m = 15$ node Gaussian quadrature scheme.

- The critical wage at which an unemployed worker will search is $w_0^* = 93.8$.
- The critical wage at which an employed worker will quit is $w_1^* = 79.4$,
- The critical wages were computed numerically by solving $\hat{V}_0(w_i^*, i) = \hat{V}_1(w_i^*, i)$, for $i = 0, 1$.



Figure 12: Conditional Value Functions and Critical Wage, Unemployed vs. Worker




Figure 13: Bellman Equation Residuals

Parametric Analysis

- The ergodic employment rate is 37%.
- How does this value change if...
 - unemployment benefit falls to 80? 100%.
 - average wage rate rises 20%? 33%.
 - standard deviation of wage shock rises to 8? 42%.
 - probability of finding a job when searching rises to 30%? 25%.
- How would you adapt the model and code to allow for a maximum time limit on unemployment benefits?

MODEL 5: American Option Pricing

- An American put option gives an investor the right, but not the obligation, to sell a specified asset at a specified strike price K within a specified number of periods T .
- At the beginning of each period $t \leq T$, the investor observes the price of the underlying asset P_t and must decide whether to exercise the option, earning a payoff $K - P_t$, or to postpone exercise for at least one more period.
- The log asset price $p_t = \log P_t$ evolves according to $p_{t+1} = p_t + \epsilon_{t+1}$ where the ϵ are serially i.i.d. normal with mean μ and variance σ^2 .
- If the investor is risk-neutral, and thus maximizes expected payoff discounted at the risk-free rate of return, then how does the critical asset price at which the investor should exercise the option vary over time?

Formulation

This is a finite-horizon, stochastic model with the following structural features:

- One continuous state variable, the log asset price

$$p_t \in (-\infty, \infty).$$

- One binary action variable, the exercise decision

$$j_t = \begin{cases} 1 & \text{exercise} \\ 0 & \text{do not exercise.} \end{cases}$$

- The reward is the option payoff

$$\pi_t = \begin{cases} 0, & \text{if } j_t = 0 \\ K - \exp(p_t), & \text{if } j_t = 1. \end{cases}$$

- State transitions are governed by

$$p_{t+1} = p_t + \epsilon_{t+1}$$

Bellman Equation

The value of an unexercised option in period t , given the log asset price p , satisfies the recursive Bellman equation

$$V_t(p) = \max\{ \underset{\text{exercise}}{K - \exp(p)}, \quad \underset{\text{do not exercise}}{\delta E_\epsilon V_{t+1}(p + \epsilon)} \},$$

$t = 0, 1, 2, \dots, T$, subject to the terminal condition $V_{T+1} \equiv 0$.

- If the put option is exercised, the investor earns a one-time payoff $K - \exp(p)$ equal to the difference between the strike price and the asset price.
- If the put option is not exercised, the investor earns no payoff and begins the following period with an option worth $V_{t+1}(\tilde{p})$, where \tilde{p} is the log asset price the following period.
- The critical log asset price p_t^* at which the option should be exercised in period t is characterized by

$$K - \exp(p_t^*) = \delta E_\epsilon V_{t+1}(p_t^* + \epsilon).$$

Numerical Solution

- The collocation method calls for the time-contingent value functions to be approximated by linear combinations of n basis functions ϕ_k :

$$V_t(p) \approx \hat{V}_t(p) \equiv \sum_{j=1}^n c_{tk} \phi_k(p),$$

$$t = 0, 1, 2, \dots, T.$$

- The $T + 1 \times n$ coefficients c_{tk} are fixed by requiring the value function approximants to satisfy the Bellman equation at n judiciously chosen nodes p_i .

This requires recursively solving the $T + 1 \times n$ nonlinear collocation equations

$$\hat{V}_t(p_i) = \max\left\{ \underset{\text{exercise}}{K - \exp(p_i)}, \quad \underset{\text{do not exercise}}{\delta E_\epsilon \hat{V}_{t+1}(p_i + \epsilon)} \right\},$$

$t = 0, 1, 2, \dots, T, i = 1, 2, \dots, n$, subject to the terminal condition $V_{T+1} \equiv 0$.

The critical asset prices p_t^* at which the put option should be exercised in period t were computed numerically by solving $K - \exp(p_t^*) = \delta E_\epsilon \hat{V}_{t+1}(p_t^* + \epsilon)$.

Example 6:

American option pricing

- Open and run CompEcon demo program [dp/05 American Put Option Pricing Model](#).
- This demo solves the American put option model assuming $T = 300$, $K = 1$, $\mu = 0.0001$, $\sigma = 0.008$, and $\delta = 0.9998$.
- The value functions are approximated by linear combinations of $n = 500$ cubic spline basis functions on $[-1, 1]$.
- The log price shock ϵ is discretized using an $m = 15$ node Gaussian quadrature scheme.

American Put Option Optimal Exercise Boundary

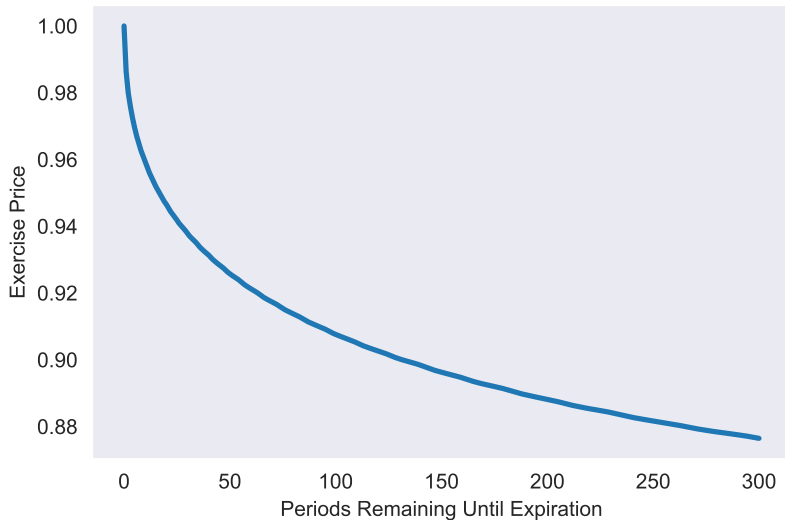


Figure 14: American Put Option Optimal Exercise Boundary

Parametric Analysis

- The critical exercise price 300 periods from expiration is 0.88.
- How does this value change if...
 - price volatility doubles? 0.75.
 - price drift doubles? 0.91.
 - strike price is \$1.10? 0.96.
- How would you adapt the model and code to price a call option, which gives an investor the right, but not the obligation, to buy?

