Lecture 9

Discrete Time Continuous Action Dynamic Economic Models

Randall Romero Aguilar, PhD

This draft: November 28, 2018

Universidad de Costa Rica
SP6534 - Economía Computacional

Table of contents

Euler Conditions

Euler Conditions

- We now turn our attention to discrete time Markov decision models with purely continuous state and action spaces and continuously differentiable reward and transition functions.
- The optimal solutions to such models can be characterized by “first-order” equilibrium conditions called the *Euler conditions*.
- Euler conditions help us understand the essential features of a dynamic decision problem and offer us an alternative way to solve for the optimal policy.
- Let d_s and d_x denote the dimensions of the state and action variables, respectively.

- We first derive the Euler conditions under the simplifying assumptions that the model is infinite-horizon and deterministic and the actions are unconstrained.
- Under these assumptions, the Bellman equation takes the relatively simple form:

$$V(s) = \max_x \{f(s,x) + \delta V(g(s,x))\}, \qquad s \in S.$$

- The Euler conditions involve the derivative of the value function

$$\lambda(s) \equiv V'(s).$$

- We call $\lambda : S \mapsto \Re^{d_s}$ the *shadow price function* because it gives the prices that the dynamically optimizing agent imputes to each of the d_s state variables.

- The Euler conditions are derived by applying the K-K-T and Envelope Theorems to the Bellman equation.
- The K-K-T Theorem implies that the optimal action x, given state s, satisfies the equimarginality condition

$$0 = f_x(s,x) + \delta\lambda(g(s,x))g_x(s,x).$$

- The Envelope Theorem implies

$$\lambda(s) = f_s(s,x) + \delta\lambda(g(s,x))g_s(s,x).$$

- Here, f_s, f_x, g_s, and g_x denote partial derivatives.

- The Euler conditions imply that along an optimal path

$$\begin{aligned} 0 &= f_x(s_t, x_t) + \delta\lambda_{t+1} g_x(s_t, x_t) \\ \lambda_t &= f_s(s_t, x_t) + \delta\lambda_{t+1} g_s(s_t, x_t). \end{aligned}$$

Steady-State

- The model may possess a well-defined *steady-state* to which the optimized economic process converges over time:

$$
\begin{aligned}
s^* &\equiv \lim s_t \\
x^* &\equiv \lim x_t \\
\lambda^* &\equiv \lim \lambda_t.
\end{aligned}
$$

- The steady-state state s^*, action x^*, and shadow price λ^*, if they exist, must satisfy the Euler and state stationarity conditions:

$$
\begin{aligned}
0 &= f_x(s^*, x^*) + \delta\lambda^* g_x(s^*, x^*) \\
\lambda^* &= f_s(s^*, x^*) + \delta\lambda^* g_s(s^*, x^*) \\
s^* &= g(s^*, x^*).
\end{aligned}
$$

- The steady-state conditions pose a finite-dimensional nonlinear equation that can be solved numerically, and often analytically, without having to solve the Bellman Equation.
- Knowledge of the steady-state is useful for understanding the longrun tendencies of the optimized economic process.
- Knowledge of the steady-state is also useful when developing initial guesses for numerical solution algorithms for the model.

Euler Equation

- If the state transition function g does not depend on the state s, the shadow price function may be eliminated as an unknown and the Euler conditions may be reduced to a single functional equation in a single unknown, the optimal policy x:

$$0 = f_x(s, x(s)) + \delta f_s(g(x(s)), x(g(x(s))))g'(x(s)).$$

- This equation, when it exists, is called the *Euler equation*.

- The Euler equation implies that along an optimal path

$$0 = f_x(s_t, x_t) + \delta f_s(s_{t+1}, x_{t+1}) g'(x_t)$$

- and, in steady-state

$$\begin{aligned} 0 &= f_x(s^*, x^*) + \delta f_s(s^*, x^*) g'(x^*) \\ s^* &= g(x^*). \end{aligned}$$

Stochastic Euler Conditions

- If the model is stochastic, the Euler conditions take the form:

$$0 = f_x(s,x) + \delta E_\epsilon \left[\lambda(g(s,x,\epsilon))g_x(s,x,\epsilon)\right]$$

and

$$\lambda(s) = f_s(s,x) + \delta E_\epsilon \left[\lambda(g(s,x,\epsilon))g_s(s,x,\epsilon)\right]$$

for all $s \in S$.

- Stochastic models lack simple steady-states, as the states visited by the process and the actions taken by the agent will vary over time due to random shocks.
- However, under mild conditions, if the model is stationary, the process over time will visit states according to a well-defined *ergodic distribution*.
- Although the erodic distribution can be formally characterized via a functional equation, the easiest way to visualize it is to solve and simulate the model over a large number of periods.

- Although stochastic models lack simple steady-states, it is nonetheless useful to derive the steady-state of the model with its shock fixed at its mean.
- Knowledge of the *deterministic steady-state* of the non-stochastic version of the model is useful when developing initial guesses for numerical solution algorithms.
- Although the deterministic steady-state need not be the mean of the ergodic distribution, the two should be reasonably close if the ergodic distribution is more or less symmetric.

Bound Constraints

- If the actions are subject to simple bounds that are differentiable functions of the state variable

$$a(s) \leq x \leq b(s)$$

  then the Euler conditions take the form of a functional complementarity problem

$$a_i(s) \leq x_i \leq b_i(s)$$

$$x_i > a_i(s) \Longrightarrow \mu_i(s) \geq 0$$

$$x_i < b_i(s) \Longrightarrow \mu_i(s) \leq 0$$

$$\begin{aligned}\lambda(s) = f_s(s,x) &+ \delta E_\epsilon \left[\lambda(g(s,x,\epsilon)) g_s(s,x,\epsilon)\right] \\ &+ \min(\mu(s),0)a'(s) + \max(\mu(s),0)b'(s)\end{aligned}$$

- … where

$$\mu(s) \equiv f_x(s, x) + \delta E_\epsilon \left[\lambda(g(s, x, \epsilon)) g_x(s, x, \epsilon)\right]$$

- Here, μ_i, measures the current and expected future reward from a marginal increase in the i^{th} action variable x_i.
- At the optimum, μ_i must not be positive if x_i is below its upper bound, otherwise additional rewards obtain by raising x_i.
- Similarly, μ_i must not be negative if x_i is above its lower bound, otherwise additional rewards obtain by reducing x_i.
- If x_i is strictly between its bounds, μ_i must be zero.

Linear-Quadratic Control

Linear-Quadratic Control Model

- A linear-quadratic (L-Q) control model is a discrete time continuous state-action Markov decision model with:
    - quadratic reward function
    - linear transition function
    - unconstrained actions
- The L-Q control model is one of few discrete time continuous state continuous action Markov decision models with known closed-form solution.

- The L-Q control model reward and transition functions take the form

$$
\begin{aligned}
f(s,x) &= F_0 + F_s s + F_x x + \tfrac{1}{2}s'F_{ss}s + s'F_{sx}x + \tfrac{1}{2}x'F_{xx}x \\
\\
g(s,x,\epsilon) &= G_0 + G_s s + G_x x + \epsilon.
\end{aligned}
$$

- Variables

$$
\begin{aligned}
s &= d_s \times 1 \text{ state vector} \\
x &= d_x \times 1 \text{ action vector} \\
\epsilon &= d_s \times 1 \text{ exogenous shock vector}
\end{aligned}
$$

- Parameters

F_0	1×1	F_s	$1 \times d_s$	F_x	$1 \times d_x$
F_{ss}	$d_s \times d_s$	F_{sx}	$d_s \times d_x$	F_{xx}	$d_x \times d_x$
G_0	$d_s \times 1$	G_s	$d_s \times d_s$	G_x	$d_s \times d_x$

F_{ss} symmetric, F_{xx} negative definite symmetric, and $E\epsilon = 0$.

- One may show by induction that the optimal policy and shadow price functions of the stationary infinite-horizon L-Q control model are linear in the state variable:

$$
\begin{aligned}
x(s) &= x^* + \Gamma(s - s^*) \\
\lambda(s) &= \lambda^* + \Lambda(s - s^*).
\end{aligned}
$$

- Here, Γ is $d_x \times d_s$, Λ is symmetric $d_s \times d_s$, and s^*, x^*, and λ^* are the deterministic steady-state state, action, and shadow price.
- The optimal policy and shadow price functions of the L-Q model do not depend on higher moments of the shock — this is known as the *certainty-equivalence property*.

- The deterministic steady-state state s^*, action x^*, and shadow price λ^* can be computed by jointly solving the Euler conditions and the state stationarity condition, a linear equation:

$$\begin{bmatrix} F'_{sx} & F_{xx} & \delta G'_x \\ F_{ss} & F_{sx} & \delta G'_s - I_n \\ G_s - I_n & G_x & 0 \end{bmatrix} \begin{bmatrix} s^* \\ x^* \\ \lambda^* \end{bmatrix} = \begin{bmatrix} -F'_x \\ -F'_s \\ -G_0 \end{bmatrix}$$

Riccati Equation

- Λ is characterized by the vector fixed-point *Riccati equation*:

$$\Lambda = -H_{sx} H_{xx}^{-1} H_{sx}' + H_{ss}$$

where

$$H_{ss} = \delta G_s' \Lambda G_s + F_{ss}$$

$$H_{sx} = \delta G_s' \Lambda G_x + F_{sx}$$

$$H_{xx} = \delta G_x' \Lambda G_x + F_{xx}$$

- Given Λ, Γ may be derived from

$$\Gamma = -H_{xx}^{-1} H_{sx}'$$

- The Riccati equation can often be solved using function iteration, but QZ decomposition is more reliable.

- If the state and action are one-dimensional, the Riccati equation takes a quadratic form

$$a\Lambda^2 + b\Lambda + c = 0$$

where

$$\begin{aligned} a &= \delta G_x^2 \\ b &= F_{xx} - \delta F_{xx} G_s^2 - \delta F_{ss} G_x^2 + 2\delta F_{sx} G_s G_x \\ c &= F_{sx}^2 - F_{ss} F_{xx} \end{aligned}$$

which can be solved using the quadratic formula.

- If in addition $F_{ss}F_{xx} = F_{sx}^2$, a condition often encountered in economic problems, then

$$\Lambda = \frac{1}{G_x^2}(F_{ss}G_x^2 - 2F_{sx}G_sG_x + F_{xx}G_s^2 - F_{xx}/\delta).$$

Model 1:

Linear-Quadratic Model

Jupyter notebook:

`dp/16 Linear-Quadratic Model.ipynb`

One-dimensional L-Q control problem

Consider the one-dimensional L-Q control problem with

$$
\begin{aligned}
F_0 &= 0.0 & F_s &= -0.8 & F_x &= -0.7 \\
F_{ss} &= -0.8 & F_{sx} &= 0.0 & F_{xx} &= -0.1 \\
G_0 &= 0.5 & G_s &= -0.1 & G_x &= 0.2 \\
\delta &= 0.9
\end{aligned}
$$

To solve the model using `LQmodel`, see the `dp/16 Linear-Quadratic Model` notebook

```
F0, Fs, Fx = 0.0, -1.0, -0.0
Fss, Fsx, Fxx = -1.0, 0.0, -0.1
G0, Gs, Gx  =  0.5, -0.2, 0.5
delta = 0.9

model = LQmodel(F0,Fs,Fx,Fss,Fsx,Fxx,G0,Gs,Gx,delta)
model.steady
```

```
{'p': array([[-0.4637]]),
 's': array([[-0.4528]]),
 'v': array([[1.3257]]),
 'x': array([[-2.0867]])}
```

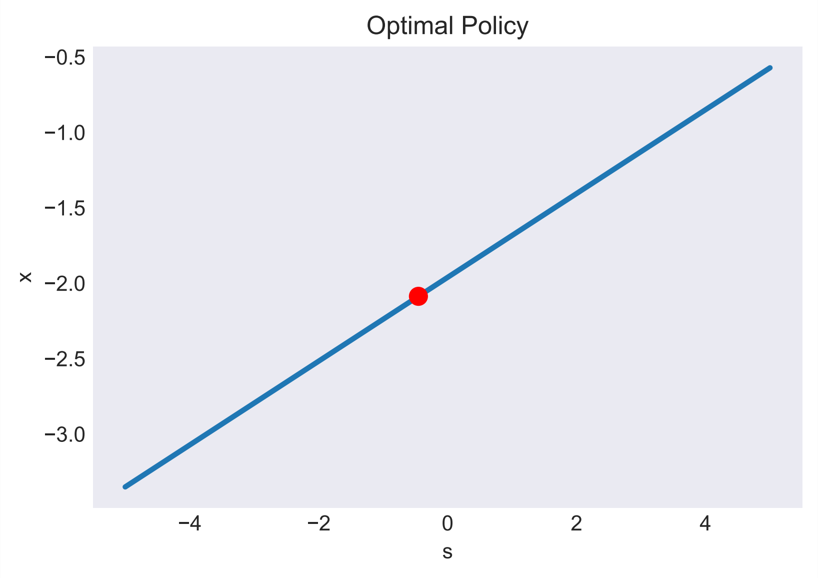

Figure 1: Optimal Policy

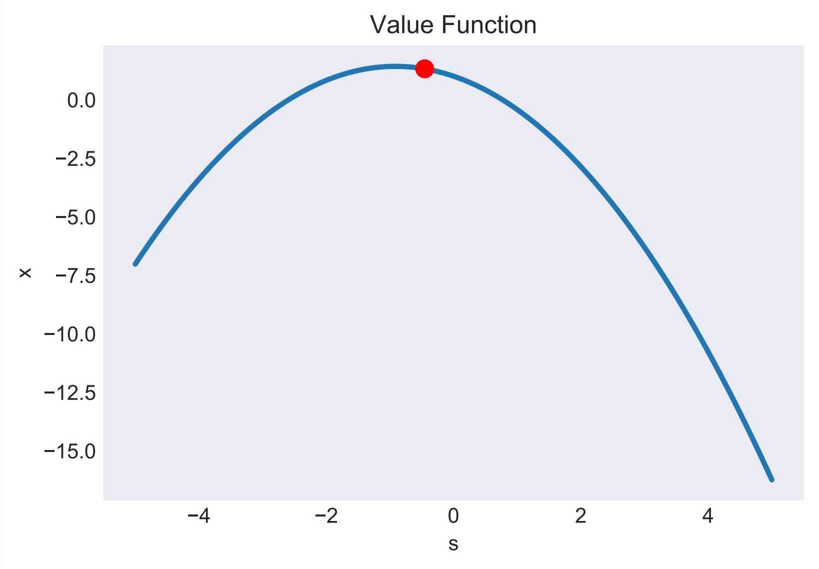

Figure 2: Value Function

Higher-dimensional L-Q control problem

Consider the higher-dimensional L-Q control problem with $F_0 = 3$ and $\delta = 0.95$.

$$F_s = \begin{bmatrix} 1 & 1 \end{bmatrix} \qquad F_x = \begin{bmatrix} 2 & 3 \end{bmatrix}$$

$$F_{ss} = \begin{bmatrix} -7 & -2 \\ -2 & -8 \end{bmatrix} \qquad F_{sx} = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix} \qquad F_{xx} = \begin{bmatrix} -2 & 0 \\ 0 & -2 \end{bmatrix}$$

$$G_0 = \begin{bmatrix} 1 \\ 1 \end{bmatrix} \qquad G_s = \begin{bmatrix} -1 & 1 \\ 1 & 0 \end{bmatrix} \qquad G_x = \begin{bmatrix} -1 & -1 \\ 2 & 3 \end{bmatrix}$$

To solve the model, execute the script

```
F0  = 3
Fs  = [1, 0]
Fx  = [1, 1]
Fss = [[-7, -2],[-2, -8]]
Fsx = [[0, 0], [0, 1]]
Fxx = [[-2, 0], [0, -2]]
G0  = [[1], [1]]
Gs  = [[-1, 1],[1, 0]]
Gx  = [[-1, -1],[2, 3]]
delta = 0.95

model2 = LQmodel(F0,Fs,Fx,Fss,Fsx,Fxx,G0,Gs,Gx,delta)
model2.steady
```

This should return, along with other information,

```
sstar =
  0.6436
 -0.3275
xstar =
  0.1272
 -0.7418
pstar =
 -2.1849
 -1.4849
vstar =
 24.9664
```

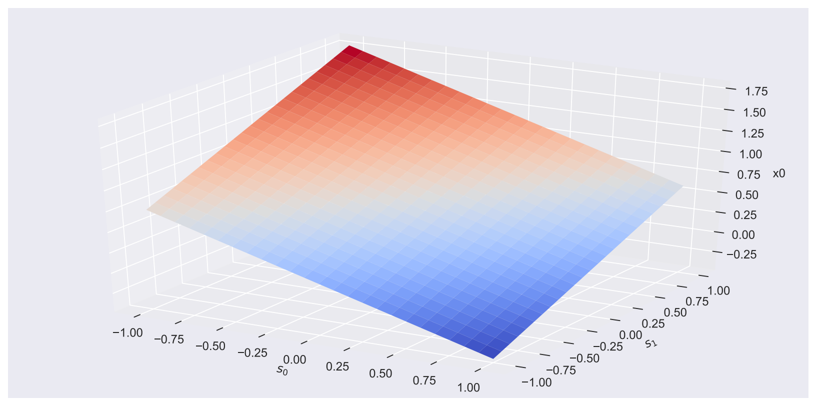

Figure 3: Optimal Policy x_1

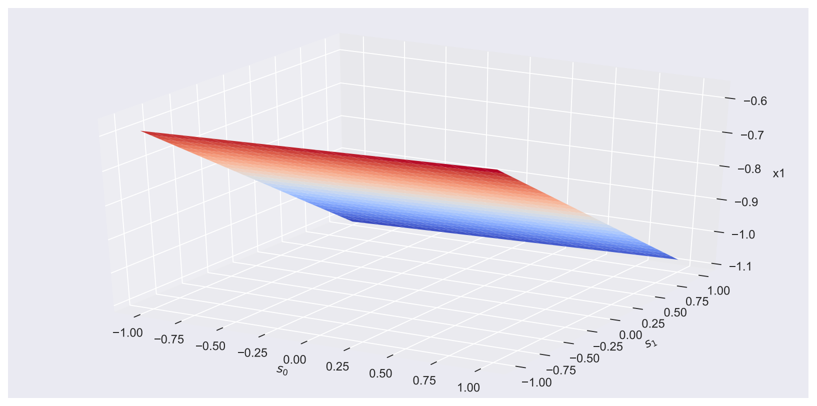

Figure 4: Optimal Policy x_2

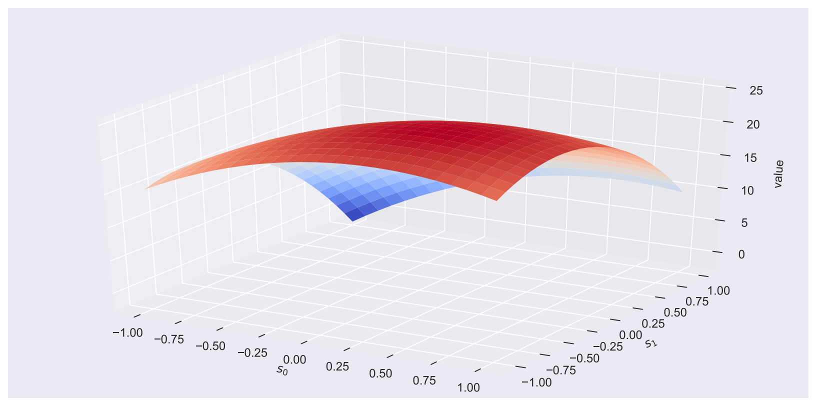

Figure 5: Value Function

Linear-Quadratic Approximation

- Because L-Q models are relatively easy to solve, many analysts use *linear-quadratic approximation* to compute approximate solutions to more general models.
- L-Q approximation calls for
    - replace reward function f with quadratic approximation
    - replace state transition function g with linear approximation
    - discard action constraints, if any exist
- Then accept the optimal policy of resulting L-Q control model as an approximate solution to the original model.

- Typically, g and f approximated using first- and second-order Taylor expansions about the deterministic steady-state state s^* and action x^*:

$$\begin{aligned} f(s,x) \quad &\approx f^* + f_s^*(s-s^*) + f_x^*(x-x^*) + \tfrac{1}{2}(s-s^*)' f_{ss}^*(s-s^*) \\ &\quad + (s-s^*)' f_{sx}^*(x-x^*) + \tfrac{1}{2}(x-x^*)' f_{xx}^*(x-x^*) \end{aligned}$$

$$g(s,x,\epsilon) \approx s^* + g_s^*(s-s^*) + g_x^*(x-x^*).$$

- Here, f^*, f_s^*, f_x^*, f_{ss}^*, f_{sx}^*, f_{xx}^*, g^*, g_s^*, and g_x^* are the values and partial derivatives of f and g evaluated at the deterministic steady-state.

- This model can be recast in canonical form using the following substitutions:

$$
\begin{aligned}
F_0 &= f^* - f_s^* s^* - f_x^* x^* + \tfrac{1}{2} s^{*\prime} f_{ss}^* s^* + s^{*\prime} f_{sx}^* x^* + \tfrac{1}{2} x^{*\prime} f_{xx}^* x^* \\
F_s &= f_s^* - s^{*\prime} f_{ss}^* - x^{*\prime} f_{xs}^* \\
F_n &= f_n^* - s^{*\prime} f_{sx}^* - x^{*\prime} f_{xx}^* \\
F_{ss} &= f_{ss}^* \\
F_{sx} &= f_{sx}^* \\
F_{xx} &= f_{xx}^* \\
G_0 &= g - g_s^* s^* - g_x^* x^* \\
G_s &= g_s^* \\
G_x &= g_x^*
\end{aligned}
$$

- L-Q approximation relies on Taylor series expansions that are valid only near the deterministic steady-state.
- L-Q approximation can work well near the steady-state if the model is deterministic or has low variance shock.
- However, L-Q approximation will perform poorly if shocks can drive the state variable far from the steady-state, where Taylor approximations deteriorate.
- L-Q approximation will perform especially poorly if the ignored constraints are occasionally binding.
- L-Q approximation is discouraged, except when assumptions of the L-Q model hold globally, or very nearly so.

One-Dimensional Continuous Action Models

Model 2:

Deterministic Optimal Economic Growth

Jupyter notebook:

`dp/06 Deterministic Optimal Economic Growth.ipynb`

Deterministic Optimal Economic Growth

- Consider an abstract economy that produces and consumes a single composite good.
- Each period t begins with a predetermined stock of wealth s_t, of which a quantity k_t is invested and the remainder $s_t - k_t$ is consumed, yielding a social benefit $\log(s_t - k_t)$.
- Wealth evolves according to $s_{t+1} = k_t^\beta$, where β is the aggregate production elasticity.
- What consumption-investment policy maximizes the present value of current and future social benefits?

Formulation

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

- One continuous state variable, wealth

$$s_t \in (0, \infty).$$

- One continuous action variable, investment

$$k_t \in [0, s_t].$$

- The reward is current social benefit

$$\log(s_t - k_t).$$

- State transitions are governed by

$$s_{t+1} = k_t^{\beta}$$

  where $0 < \beta < 1$.

The present value of current and future social benefits, given wealth s, satisfies the Bellman equation

$$V(s) = \max_{0 \leq k \leq s} \left\{ \log(s-k) + \delta V(k^\beta) \right\}.$$

- This Bellman equation has a well-known closed-form solution

$$V(s) = v^* + \frac{1}{1-\delta\beta}(\log(s) - \log(s^*))$$

with optimal policy

$$k = \delta\beta s,$$

where

$$s^* = (\delta\beta)^{\beta/(1-\beta)}$$

$$v^* = \log((1-\delta\beta)s^*)/(1-\delta).$$

- Knowledge of the closed-form solution will allow us to test the accuracy of approximate solutions obtained numerically.

Euler Conditions

- By assumptions, the constraints on k will never bind at an optimum, implying that the shadow price function $\lambda(s) = V'(s)$ must satisfy the Euler conditions

$$0 = -(s-k)^{-1} + \delta\beta\lambda(k^{\beta})k^{\beta-1}$$

$$\lambda(s) = (s-k)^{-1}.$$

- The Euler conditions imply that along the optimal path

$$\lambda_t = \delta\beta\lambda_{t+1}k_t^{\beta-1} = (s_t - k_t)^{-1}.$$

Steady-State

- The steady-state wealth s^*, investment k^*, and shadow price λ^* must satisfy the Euler and state stationarity conditions

$$\begin{aligned} \lambda^* &= \delta\beta\lambda^* {k^*}^{\beta-1} = (s^* - k^*)^{-1} \\ s^* &= {k^*}^{\beta}. \end{aligned}$$

- These conditions imply

$$\begin{aligned} s^* &= (\delta\beta)^{\beta/(1-\beta)} \\ k^* &= (\delta\beta)^{\frac{1}{1-\beta}} \\ \lambda^* &= \frac{(\delta\beta)^{\beta/(\beta-1)}}{1-\delta\beta}. \end{aligned}$$

Numerical Solution

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

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

- The n coefficients c_j 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 nonlinear collocation equations

$$\sum_{j=1}^{n} c_j \phi_j(s_i) = \max_{0 \leq k \leq s_i} \left\{ \log(s_i - k) + \delta \sum_{j=1}^{n} c_j \phi_j(k^\beta) \right\},$$

$i = 1, 2, \ldots, n$, for the n unknown coefficients c_j, $j = 1, 2, \ldots, n$.

Parametric Example

- Open and run CompEcon demo program `dp/06 Deterministic Optimal Economic Growth Model`.
- This demo solves the deterministic optimal economic growth model assuming $\beta = 0.5$ and $\delta = 0.9$.
- The value function is approximated by a linear combination of $n = 15$ Chebychev polynomial basis functions on $[0.2, 1.0]$.

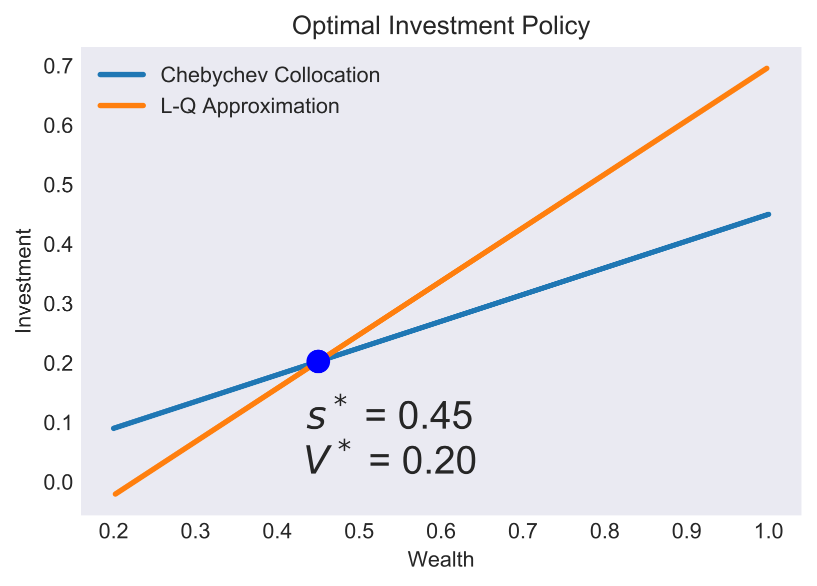

Figure 6: Optimal Investment Policy

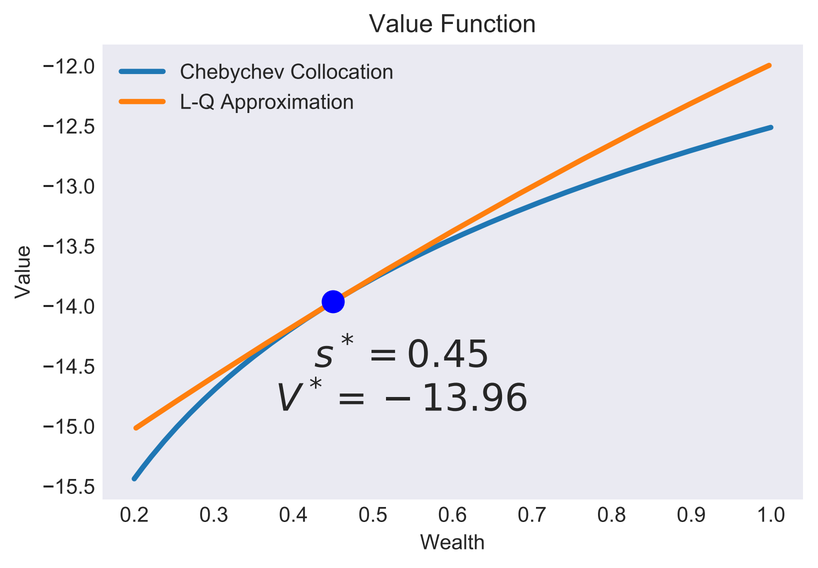

Figure 7: Value Function

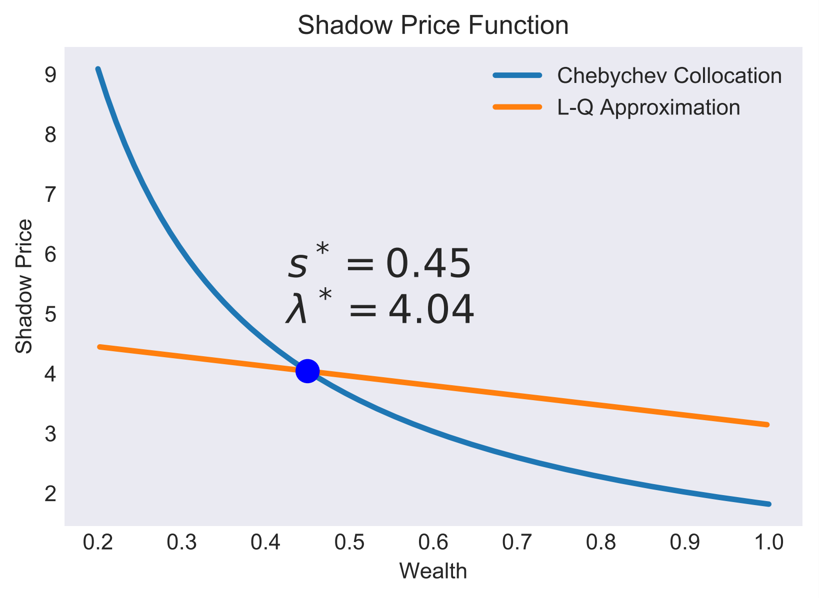

Figure 8: Shadow Price Function

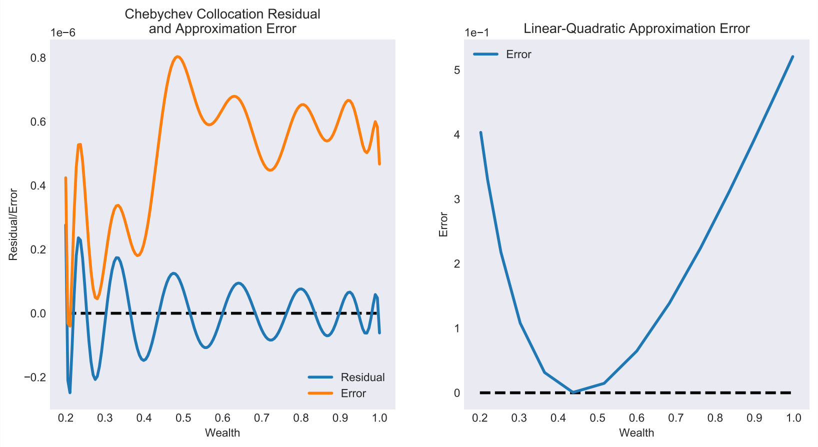

Figure 9: Chebychev Collocation Bellman Equation Residual and Value Function Approximation Error

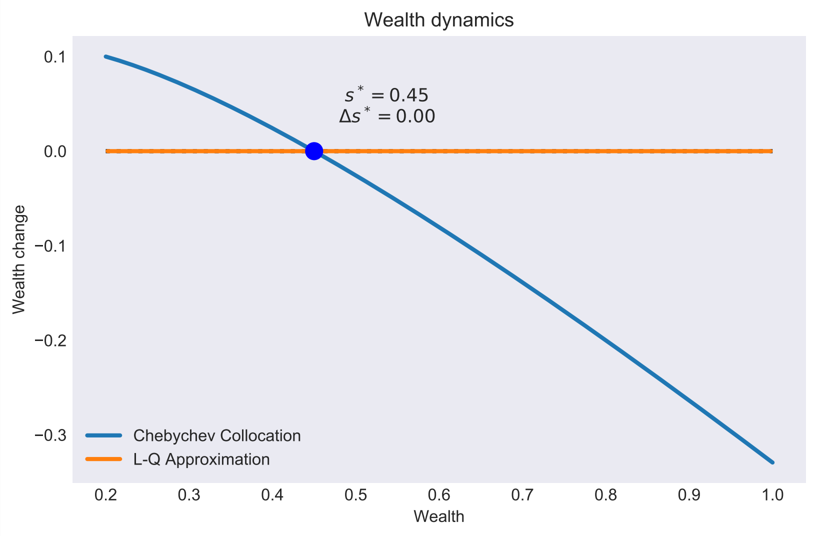

Figure 10: Linear-Quadratic Approximation Value Function Approximation Error

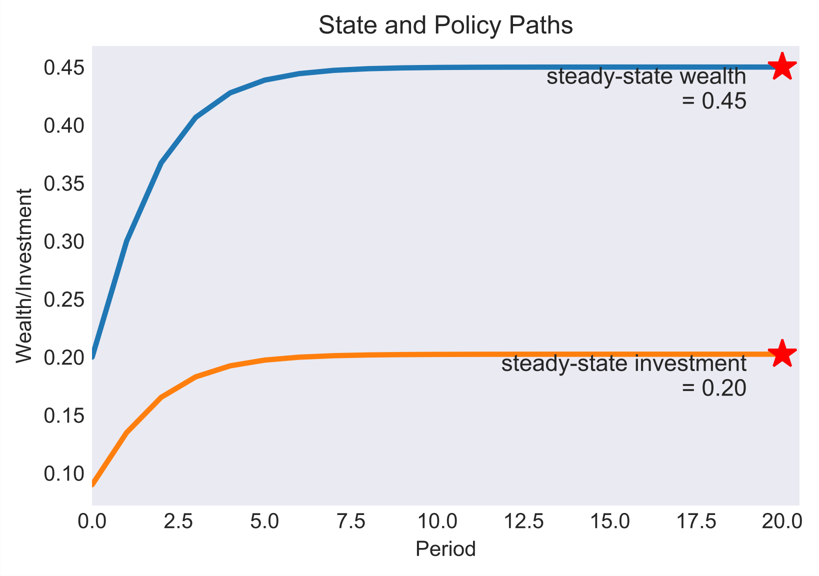

Figure 11: Simulated Wealth and Investment

Parametric Analysis

- The steady-state wealth and investment are 0.45 and 0.20, respectively.
- How do these values change if ...
  - production elasticity is 0.7? 0.34 and 0.21.
  - production elasticity is 0.3? 0.57 and 0.15.
  - planner discounts the future less ($\delta = 0.95$)? 0.47 and 0.23.
- How would you revise the model and code to allow for stochastic production shocks?

Model 3:

Stochastic Optimal Economic Growth

Jupyter notebook:

`dp/07 Stochastic Optimal Economic Growth.ipynb`

Stochastic Optimal Economic Growth

- Consider an abstract economy that produces and consumes a single composite good.
- Each period t begins with a predetermined stock of wealth s_t, of which a quantity k_t is invested and the remainder $s_t - k_t$ is consumed, yielding a social benefit $u(s_t - k_t)$.
- Wealth evolves according to $s_{t+1} = \gamma k_t + \epsilon_{t+1} h(k_t)$ where γ is the capital survival rate (1 minus the depreciation rate), h is the aggregate production function, and the ϵ are serially i.i.d. positive production shocks with mean 1.
- What consumption-investment policy maximizes the present value of current and expected future social benefits?

Formulation

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

- One continuous state variable, wealth

$$s_t \in (0, \infty).$$

- One continuous action variable, investment

$$k_t \in [0, s_t].$$

- The reward is current social benefit

$$u(s_t - k_t)$$

where $u' > 0$, $u'' < 0$, and $u'(0) = \infty$.

- State transitions are governed by

$$s_{t+1} = \gamma k_t + \epsilon_{t+1} h(k_t)$$

where $0 < \gamma < 1$, $h' > 0$, $h'' < 0$, and $h(0) = 0$.

Bellman Equation

The present value of current and expected future social benefits, given wealth s, satisfies the Bellman equation

$$V(s) = \max_{0 \leq k \leq s} \left\{ u(s-k) + \delta E_{\epsilon} V(\gamma k + \epsilon h(k)) \right\}.$$

Euler Conditions

- By assumptions, the constraints on k will never bind at an optimum, implying that the shadow price of wealth $\lambda(s)$ must satisfy the Euler conditions

$$0 = -u'(s-k) + \delta E_\epsilon \left[\lambda(\gamma k + \epsilon h(k))(\gamma + \epsilon h'(k))\right]$$

$$\lambda(s) = u'(s-k).$$

- The Euler conditions imply that along the optimal path

$$u'_t = \delta E_t \left[(\gamma + h'_{t+1})u'_{t+1}\right]$$

where $u'_t \equiv u'(s_t - k_t)$ is marginal social benefit in period t and $h'_{t+1} \equiv \epsilon_{t+1}h'(k_t)$ is the ex post marginal product of capital in period $t+1$.

Deterministic Steady-State

- The steady-state wealth s^*, investment k^*, and shadow price λ^* when the production shock ϵ fixed at its mean 1 must satisfy the Euler and state stationarity conditions

$$\begin{aligned} 0 &= -u'(s^* - k^*) + \delta\lambda^*(\gamma + h'(k^*)) \\ \lambda^* &= u'(s^* - k^*) \\ s^* &= \gamma k^* + h(k^*). \end{aligned}$$

- The deterministic steady-state conditions imply the *Golden Rule*

$$h'(k^*) = 1 - \gamma + \rho$$

where $\rho \equiv 1/\delta - 1$ is the *discount rate*.
- That is, in deterministic steady-state, the marginal product of capital equals the capital depreciation rate plus the discount rate.

Numerical Solution

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

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

- The n coefficients c_j 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 nonlinear collocation equations

$$\sum_{j=1}^{n} c_j \phi_j(s_i) = \max_{0 \le k \le s_i} \left\{ u(s_i - k) + \delta E_\epsilon \sum_{j=1}^{n} c_j \phi_j(\gamma k + \epsilon h(k)) \right\},$$

$i = 1, 2, \ldots, n$, for the n unknown coefficients c_j, $j = 1, 2, \ldots, n$.

Parametric Example

- Open and run CompEcon demo program `dp/07 Stochastic Optimal Economic Growth Model`.
- This demo solves the stochastic optimal economic growth model assuming
  - $u(c) = c^{1-\alpha}/(1-\alpha)$, $\alpha = 0.2$;
  - $h(k) = k^{\beta}$, $\beta = 0.5$;
  - ϵ are serially i.i.d. lognormal, mean 1, volatility $\sigma = 0.1$;
  - $\gamma = 0.9$, $\delta = 0.9$.
- The value function is approximated by a linear combination of $n = 10$ Chebychev polynomial basis functions on $[5, 10]$.
- The production shock ϵ is discretized using an $m = 3$ node Gaussian quadrature scheme.

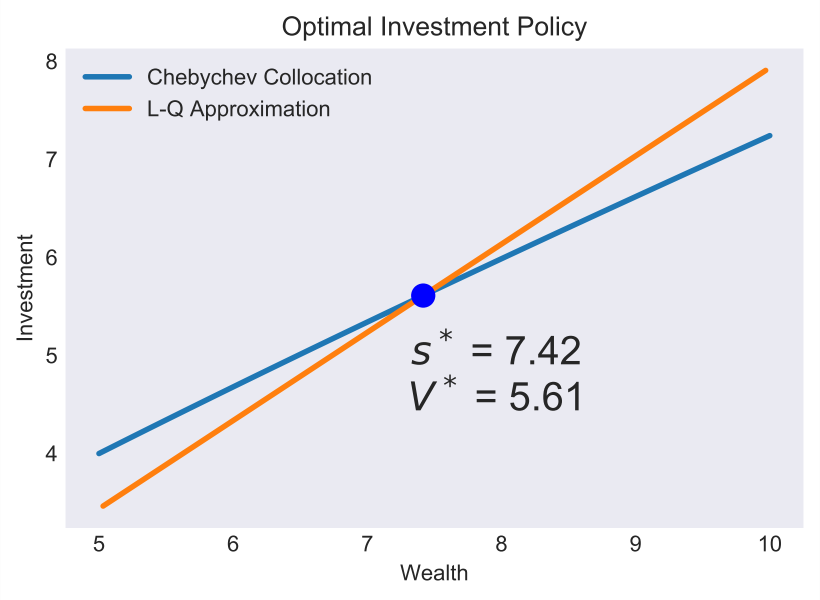

Figure 12: Optimal Investment Policy

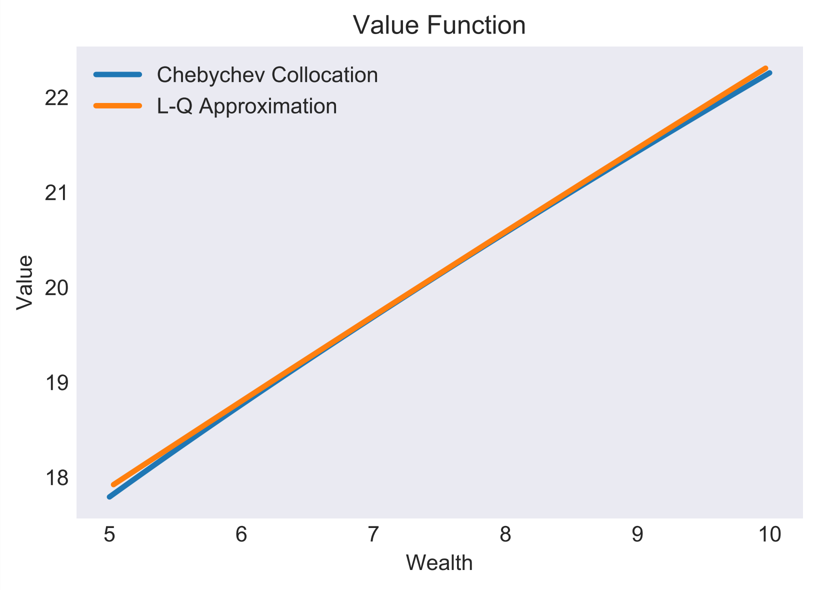

Figure 13: Value Function

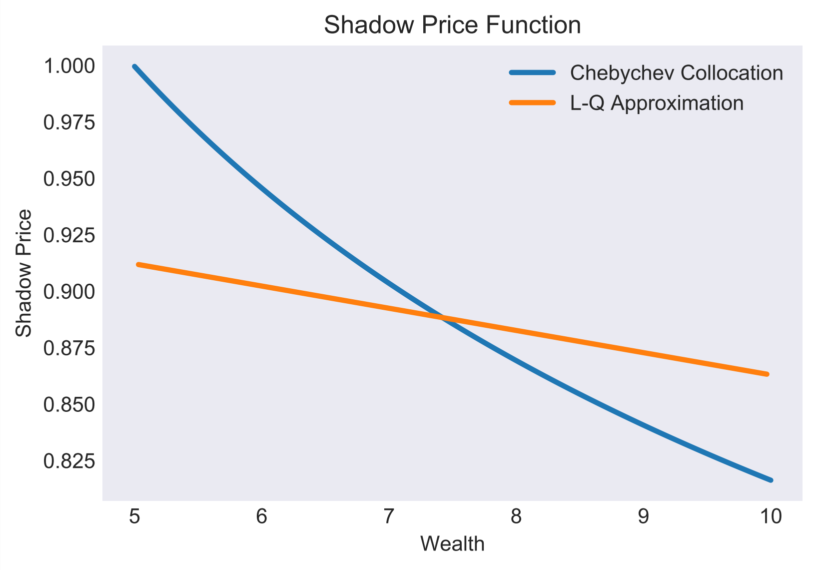

Figure 14: Shadow Price Function

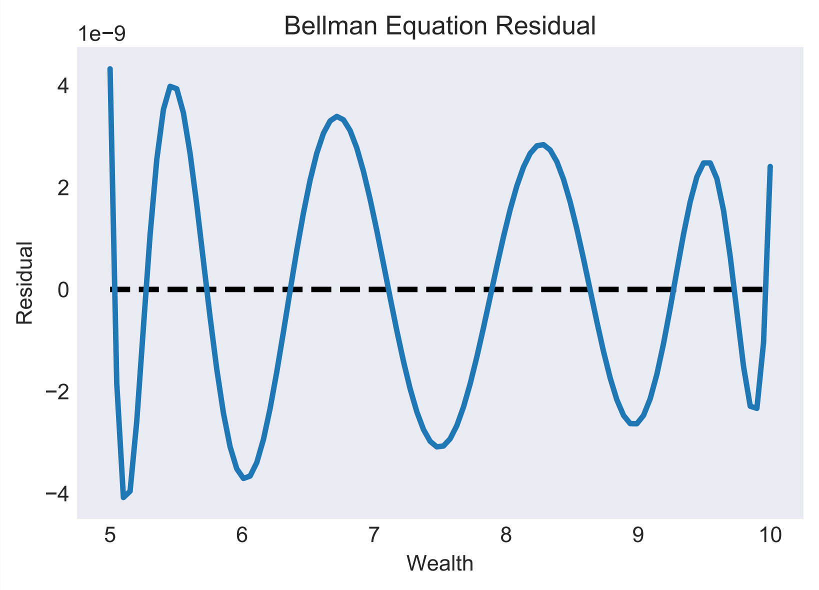

Figure 15: Bellman Equation Residual

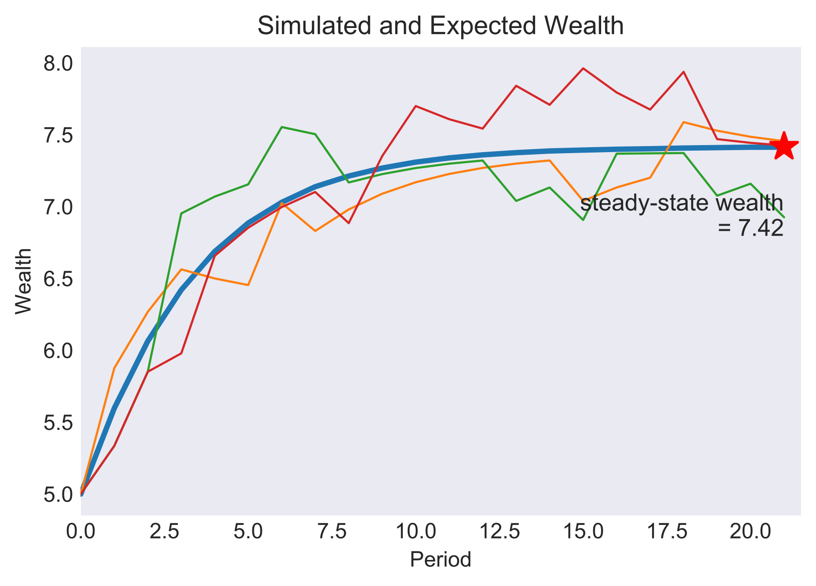

Figure 16: Simulated and Expected Wealth

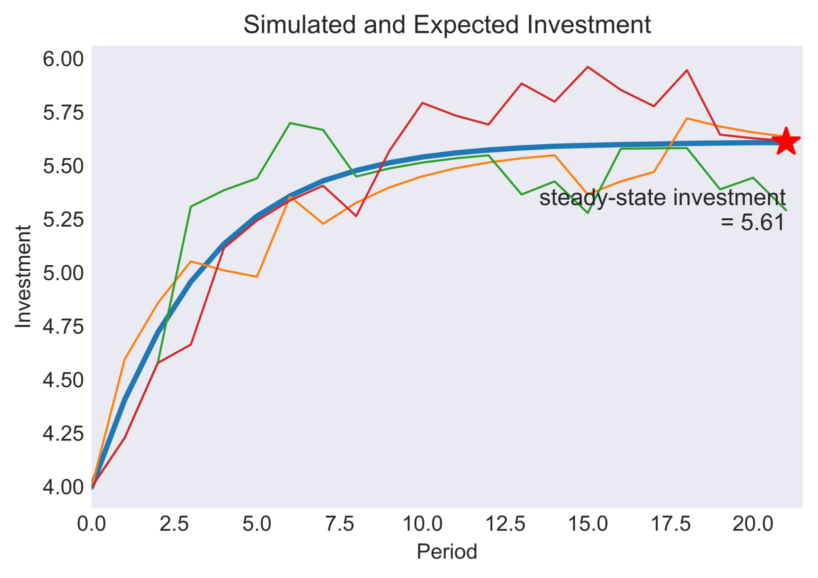

Figure 17: Simulated and Expected Investment

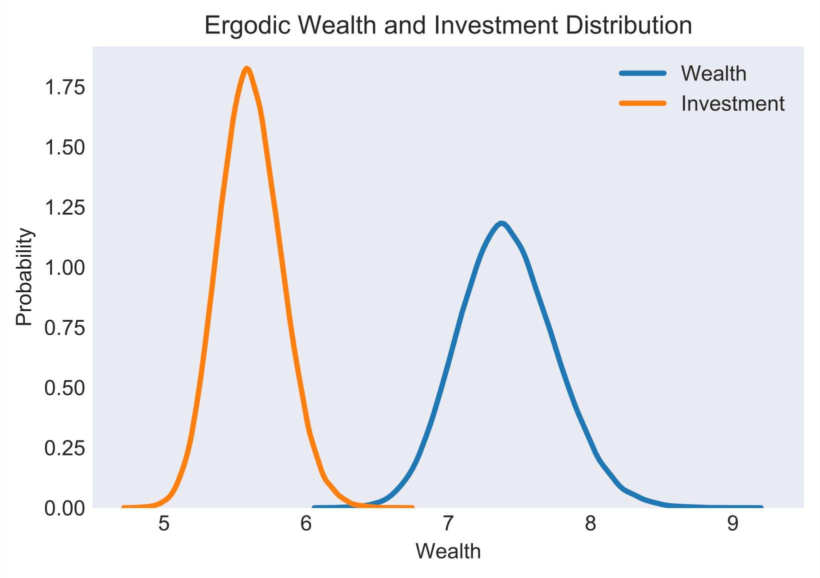

Figure 18: Ergodic Distribution of Wealth

- The ergodic distribution of wealth has mean 7.41 and standard deviation 0.34.
- How do these values change if ...
    - production volatility rises to 0.15? 7.42 and 0.52.
    - production volatility falls to 0.05? 7.42 and 0.17.
- How would you revise the model and code to allow for constant absolute risk aversion?

Model 4:

Public Renewable Resource Management

Jupyter notebook:

`dp/08 Public Renewable Resource Management.ipynb`

Public Renewable Resource Management

- A social planner wishes to maximize social surplus derived from harvesting a publicly-owned renewable resource.
- Each period t begins with a predetermined stock s_t of the resource, of which a quantity q_t is harvested at a constant unit cost k and sold at a market clearing price $p_t = p(q_t)$.
- The remainder $s_t - q_t$ is retained for reproduction, yielding a resource stock $s_{t+1} = g(s_t - q_t)$ the following period.
- What harvest policy maximizes the present value of current and future social surplus?

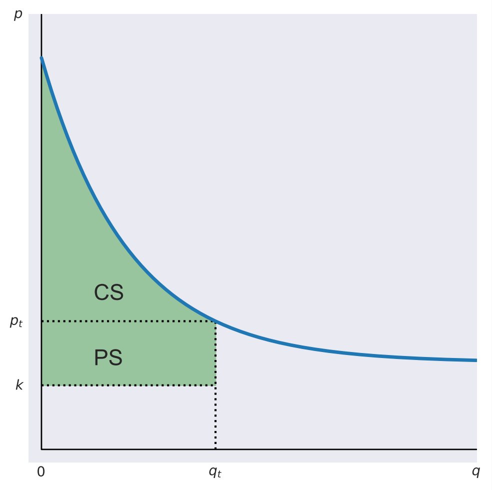

Figure 19: Social Surplus

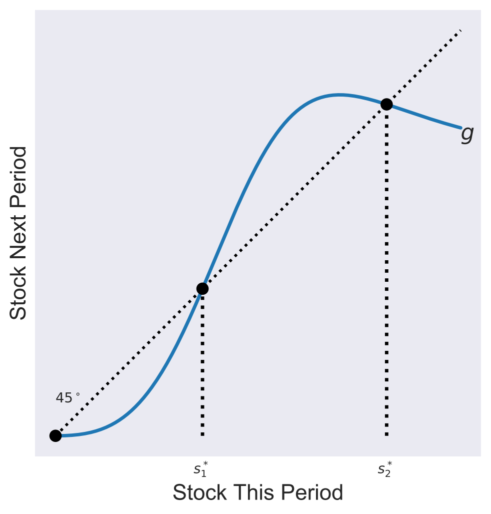

Figure 20: Biological Reproduction Function and Natural Steady-States

In the absence of human intervention, the natural system admits three steady-states:

- Extinction, $s^* = 0$
- Unsustainable, Unstable Steady-State s_1^*
- Sustainable, Stable Steady-State s_2^*

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

- One continuous state variable, resource stock

$$s_t \in [0, \infty).$$

- One continuous action variable, harvest

$$q_t \in [0, s_t].$$

- The reward is current social surplus

$$\int_0^{q_t} p(q)dq - kq_t$$

where $p > 0$, $p' < 0$, $p(0) = \infty$, and $k > 0$.
- State transitions are governed by

$$s_{t+1} = g(s_t - q_t)$$

where $g \geq 0$, $g(0) = 0$, and $g'(0) > 0$.

The present value of current and future social surplus, given a resource stock s, satisfies the Bellman equation

$$V(s) = \max_{0 \leq q \leq s} \left\{ \int_0^q p(\xi) d\xi - kq + \delta V(g(s-q)) \right\}.$$

- By assumptions, the constraint on q will never bind at an optimum, implying that the shadow price of the resource $\lambda(s)$ must satisfy the Euler conditions:

$$p(q) = k + \delta\lambda(g(s-q))g'(s-q)$$

$$\lambda(s) = \delta\lambda(g(s-q))g'(s-q).$$

- The Euler conditions imply that along the optimal path

$$p_t = k + \lambda_t$$

$$\lambda_t = \delta\lambda_{t+1}g'_t$$

where p_t is the market price and g'_t is the marginal yield of resource stock retained in period t.

- Thus, the market price of the harvested resource must cover both the shadow price of the unharvested resource and the marginal cost of harvesting it.
- Moreover, the current value of one unit of the resource equals the discounted value of its yield in the following period.

Steady-State

- The steady-state resource stock s^*, harvest q^*, and shadow price λ^* must satisfy the Euler and state stationarity conditions

$$p(q^*) = k + \delta\lambda^* g'(s^* - q^*)$$

$$\lambda^* = \delta\lambda^* g'(s^* - q^*)$$

$$s^* = g(s^* - q^*).$$

- These conditions imply $g'(s^* - q^*) = 1 + \rho$.
- That is, in steady-state, the marginal rate of growth of resource stock equals the discount rate.

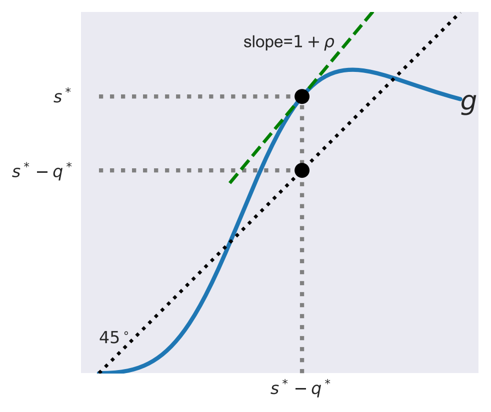

Figure 21: Optimal Stock, Harvest and Retention

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

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

- The n coefficients c_j 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 nonlinear collocation equations

$$\sum_{j=1}^{n} c_j \phi_j(s_i) = \max_{0 \le q \le s_i} \left\{ \int_0^q p(\xi) d\xi - kq + \delta \sum_{j=1}^{n} c_j \phi_j(g(s_i - q)) \right\},$$

$i = 1, 2, \ldots, n$, for the n unknown coefficients c_j, $j = 1, 2, \ldots, n$.

- Open and run CompEcon demo program `dp/08 Public Renewable Resource Model`.
- This demo solves the renewable resource model assuming
  - $p(q) = q^{-\gamma}$, $\gamma = 0.5$;
  - $g(s) = \alpha s - 0.5\beta s^2$, $\alpha = 4$, $\beta = 1$;
  - $k = 0.2$, $\delta = 0.9$.
- The value function is approximated by a linear combination of $n = 8$ Chebychev polynomial basis functions on $[6, 9]$.

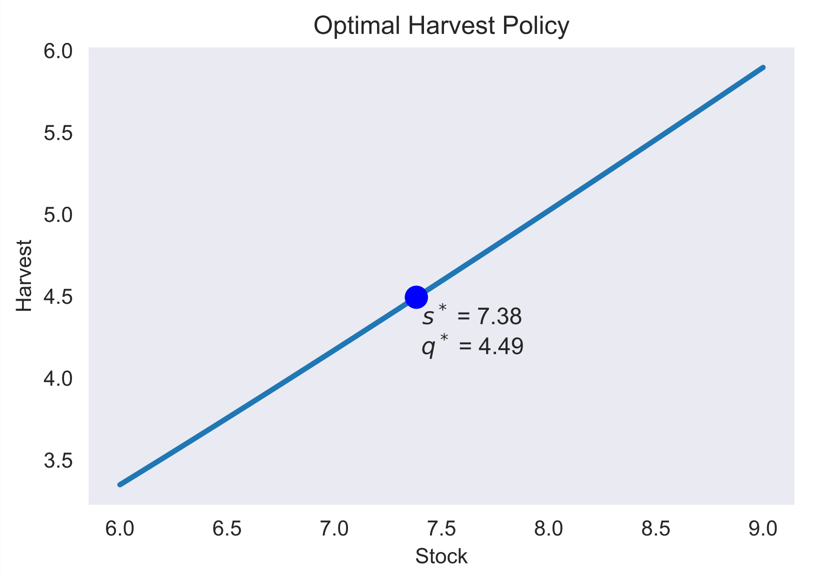

Figure 22: Optimal Harvest Policy

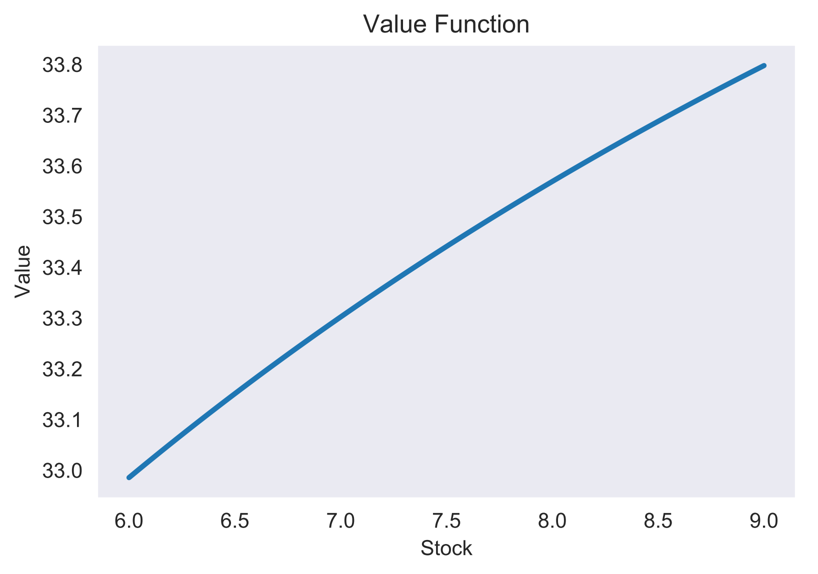

Figure 23: Value Function

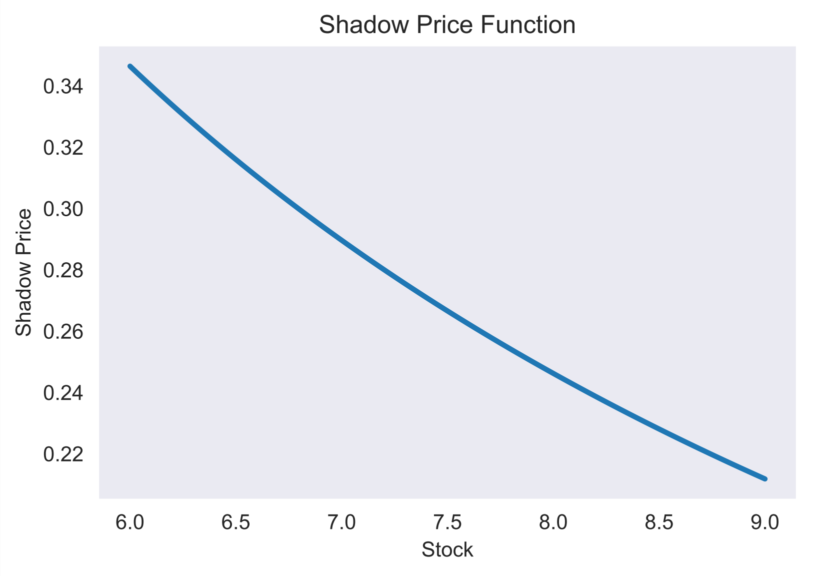

Figure 24: Shadow Price Function

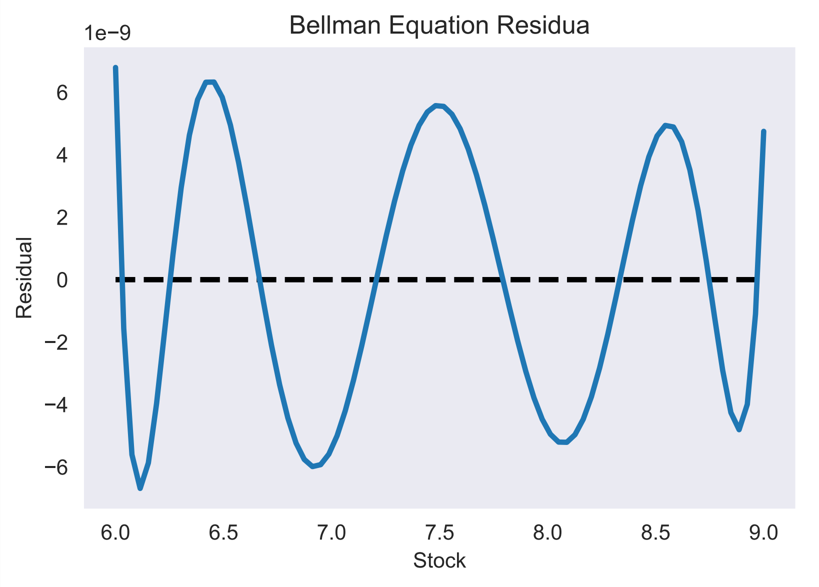

Figure 25: Bellman Equation Residual

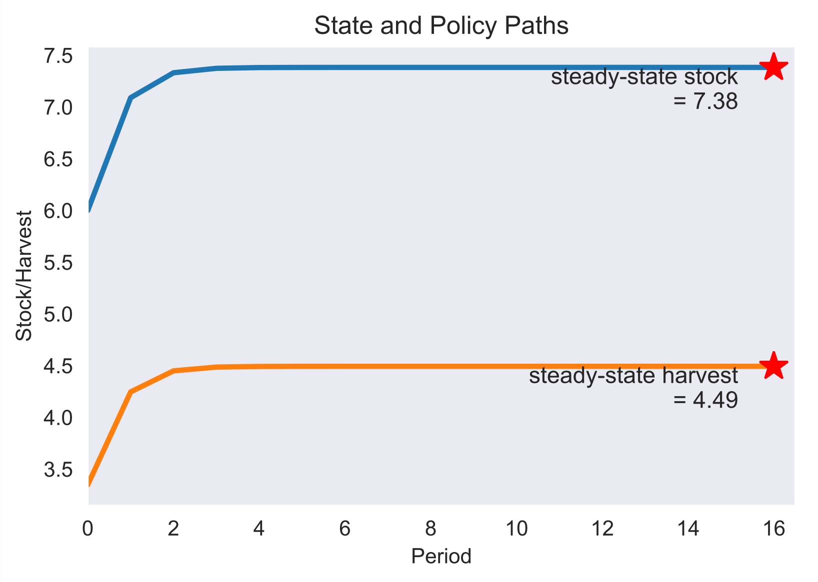

Figure 26: Simulated Stock and Harvest

- The steady-state stock and harvest are 7.38 and 4.49, respectively.
- How do these values change if …
  - planner discounts future less, $\delta = 0.95$? 7.45 & 4.50.
  - climate change cuts α and β in half? 2.77 & 0.99.
  - cost of harvest doubles? No change.
- How would you revise the model and code to allow for a logistic growth function?

Model 5:

Private Nonrenewable Resource Management

Jupyter notebook:

`dp/09 Private Nonrenewable Resource Management.ipynb`

- A mine owner begins each period t with a predetermined stock of ore s_t, of which he will extract and sell a quantity q_t at the market price $p_t = p(q_t)$.
- The total cost of extraction is given by

$$C_t = \int_{s_t - q_t}^{s_t} k(s)ds$$

  where $k(s)$ is the marginal cost of extracting one unit of ore when the stock is s ($k' < 0$).
- Given the current stock of ore is s_0, what extraction policy maximizes the value of the mine?

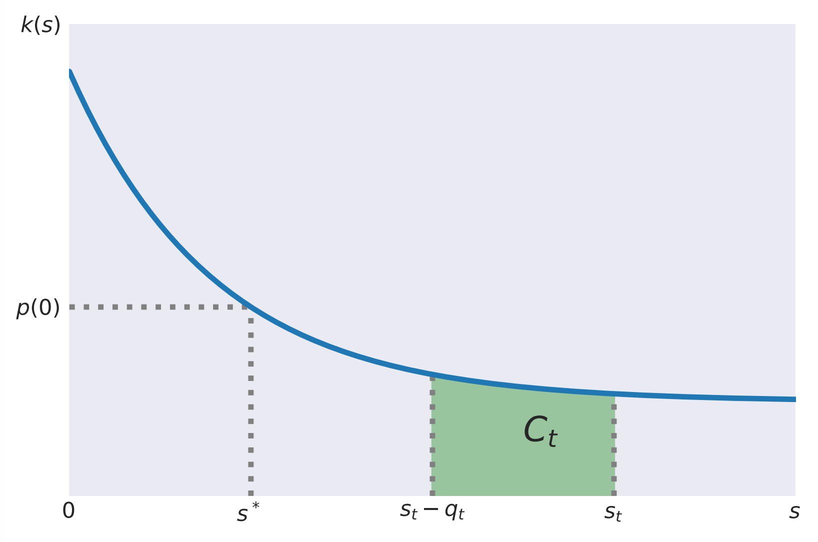

Figure 27: Cost of Extraction

Formulation

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

- One continuous state variable, ore stock

$$s_t \in [0, s_0].$$

- One continuous action variable, ore extracted and sold

$$q_t \in [0, s_t].$$

- The reward is current profit

$$p(q_t)q_t - \int_{s_t - q_t}^{s_t} k(s)ds.$$

  where $p > 0$, $p' < 0$, $k > 0$, $k' < 0$, and $k(s_0) < p(0) < k(0) < \infty$.
- State transitions are governed by

$$s_{t+1} = s_t - q_t.$$

The value of the mine, given it contains a stock of ore s, satisfies the Bellman equation

$$V(s) = \max_{0 \leq q \leq s} \left\{ p(q)q - \int_{s-q}^{s} k(\xi) d\xi + \delta V(s-q) \right\}.$$

- The Euler conditions take the form of a complementarity condition

$$\mu \equiv p(q) + p'(q)q - k(s-q) - \delta\lambda(s-q)$$

$$q \geq 0 \perp \mu \leq 0.$$

$$\lambda(s) = k(s-q) - k(s) + \delta\lambda(s-q)$$

where μ is the longrun marginal profit of extraction.

- It is optimal to abandon the mine if $k(s) > p(0)$, that is, if the marginal cost of extraction exceeds the maximum price anyone would pay.
- The abandonment point is thus s^*, where $k(s^*) = p(0)$.

- Until such time that the mine is abandoned,

$$p_t + p_t' q_t = \lambda_t + k_{t+1}$$

$$\lambda_t - \delta \lambda_{t+1} = k_{t+1} - k_t,$$

  where k_t is the marginal cost of extraction at the beginning of period t.
- That is, the marginal revenue of extracted ore will equal the shadow price of unextracted ore plus the marginal cost of extraction.
- Also, the present-valued shadow price of unextracted ore falls at the rate at which the marginal cost of extraction rises.

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

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

- The n coefficients c_j 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 nonlinear collocation equations

$$\sum_{j=1}^{n} c_j \phi_j(s_i) = \max_{0 \le q \le s_i} \left\{ p(q)q - \int_{s_i - q}^{s_i} k(\xi) d\xi + \delta \sum_{j=1}^{n} c_j \phi_j(s_i - q) \right\},$$

$i = 1, 2, \ldots, n$, for the n unknown coefficients c_j, $j = 1, 2, \ldots, n$.

- Open and run CompEcon demo program `dp/09 Private Non-Renewable Resource Model`.
- This demo solves the nonrenewable resource model assuming
    - $p(q) = a_1 - a_2 q$, $a_1 = 5$, $a_2 = 0.8$;
    - $k(\xi) = b_1 - b_2 \xi$, $b_1 = 7$, $b_2 = 1$;
    - $s_0 = 10$, $\delta = 0.9$.
- The value function is approximated by a linear combination of $n = 101$ cubic spline basis functions on $[0, 10]$.

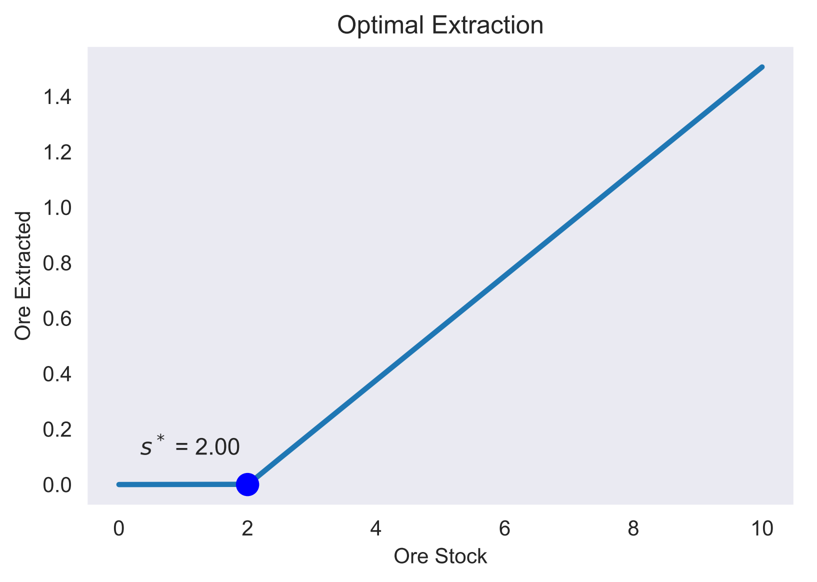

Figure 28: Optimal Ore Extraction Policy

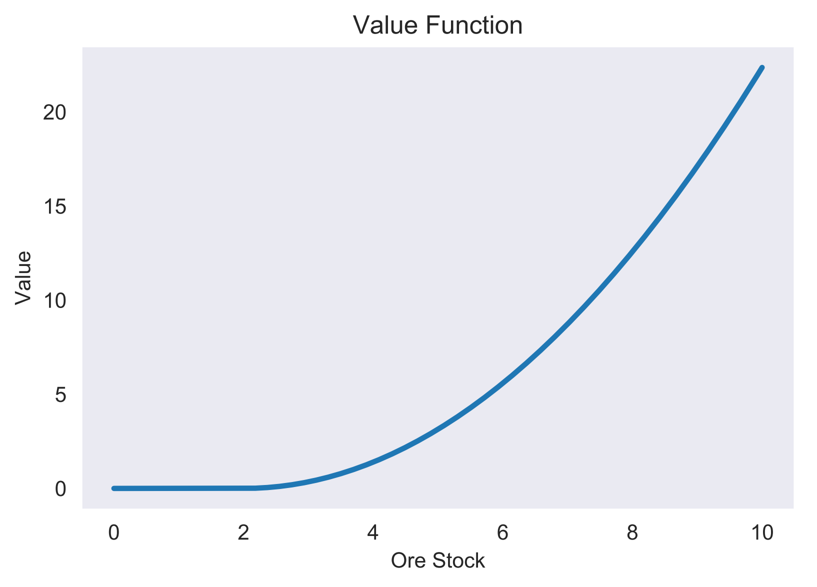

Figure 29: Value Function

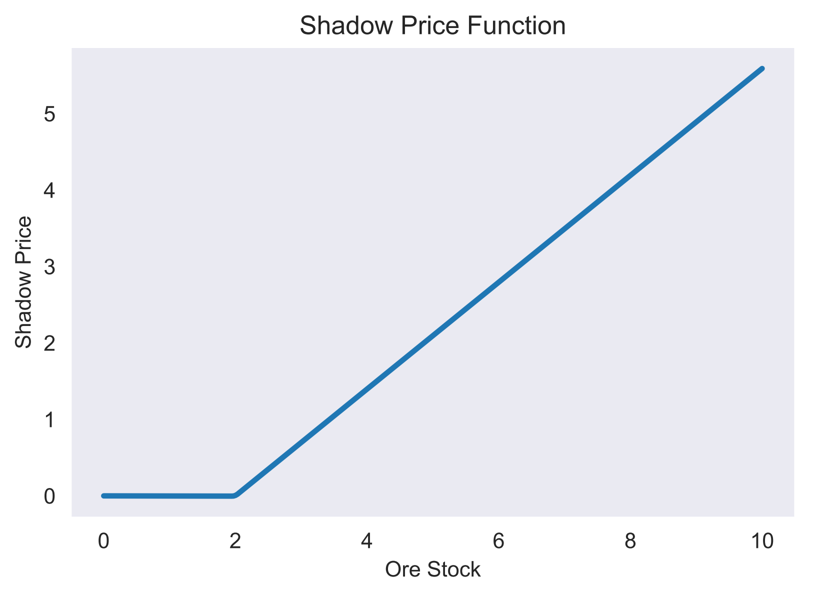

Figure 30: Shadow Price Function

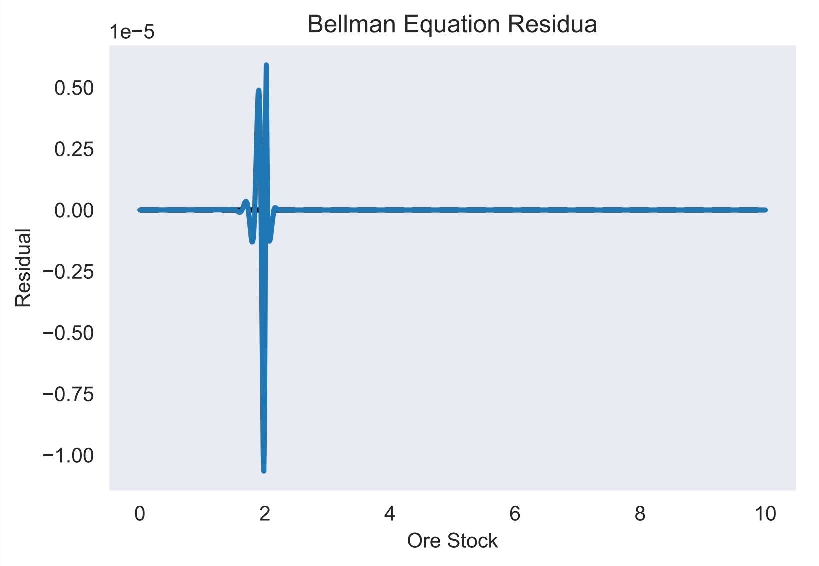

Figure 31: Bellman Equation Residual

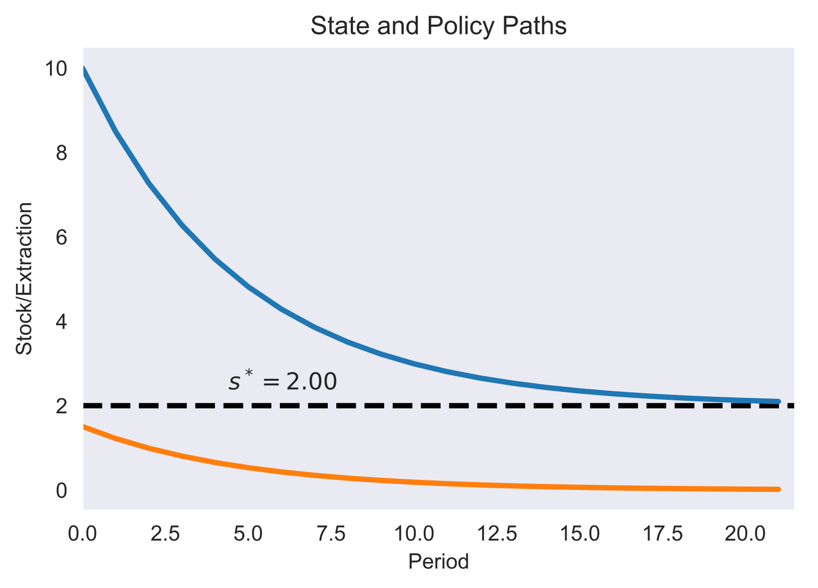

Figure 32: Simulated Stock and Extraction

- The abandonment point is 2.
- How does this value change if ...
  - owner discounts future less, $\delta = 0.95$? No change.
  - demand rises, i.e., a_1 rises to 7? 1.
  - government imposes \$2 tax on extraction?4.
- How would you revise the model and code if demand were stochastic?

Model 6:

Water Resource Management

Jupyter notebook:

`dp/10 Water Resource Management.ipynb`

Water Resource Management

- Water from a reservoir is used for irrigation and recreation.
- Irrigation in spring benefits farmers, but reduces the reservoir level in summer, damaging recreational users.
- Each year t begins in spring with a predetermined stock of water s_t in the reservoir, of which a quantity q_t is released for irrigation and the remainder $s_t - q_t$ retained, yielding farmer and recreational user benefits $F(q_t)$ and $U(s_t - q_t)$, respectively.
- Reservoir levels are replenished by serially i.i.d rainfalls ϵ during the winter.
- What irrigation policy maximizes the sum of current and expected future farmer and recreational user benefits?

Formulation

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

- One continuous state variable, reservoir level

$$s_t \in [0, \infty).$$

- One continuous action variable, quantity of water released for irrigation

$$q_t \in [0, s_t].$$

- The reward is total social benefits

$$F(q_t) + U(s_t - q_t)$$

  where $F' > 0$, $F'' < 0$, $U' > 0$, $U'' < 0$, and $F'(0) = U'(0) = \infty$.
- State transitions are governed by

$$s_{t+1} = s_t - q_t + \epsilon_{t+1}.$$

The social value of the reservoir, given that it contains s units of water at the beginning of the year, satisfies the Bellman equation

$$V(s) = \max_{0 \leq q \leq s} \left\{ F(q) + U(s-q) + \delta E_{\epsilon} V(s-q+\epsilon) \right\}.$$

- By assumptions, the constraints on q will never bind at an optimum, implying that the shadow price of water in the reservoir $\lambda(s)$ must satisfy the Euler conditions

$$0 = F'(q) - U'(s-q) - \delta E_\epsilon \lambda(s-q+\epsilon)$$

$$\lambda(s) = U'(s-q) + \delta E_\epsilon \lambda(s-q+\epsilon).$$

- The Euler conditions imply that along the optimal path

$$\lambda_t = F'_t = U'_t + \delta E_t \lambda_{t+1}$$

where F'_t and U'_t are the marginal farmer and recreational user benefits in year t, respectively.

- Thus, on the margin, the benefit received by farmers this year from releasing one unit of water must equal the marginal benefit received by recreational users this year from retaining the unit of water plus the benefits of having that unit available for either irrigation or recreation the following year.

- The deterministic steady-state reservoir level s^*, irrigation level q^*, and shadow price λ^* when rainfall ϵ is fixed at its mean $\bar{\epsilon}$ must satisfy the Euler and state stationarity conditions

$$F'(q^*) = \lambda^*$$

$$U'(s^* - q^*) = (1 - \delta)\lambda^*$$

$$q^* = \bar{\epsilon}.$$

- These conditions imply that the deterministic steady-state irrigation level and shadow price of water are not affected by the discount rate.
- The deterministic steady-state reservoir level, however, is affected by the discount rate.

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

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

- The n coefficients c_j 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 nonlinear collocation equations

$$\sum_{j=1}^{n} c_j \phi_j(s_i) = \max_{0 \le q \le s_i} \left\{ F(q) + U(s_i - q) + \delta E_\epsilon \sum_{j=1}^{n} c_j \phi_j(s_i - q + \epsilon) \right\},$$

$i = 1, 2, \ldots, n$, for the n unknown coefficients c_j, $j = 1, 2, \ldots, n$.

- Open and run CompEcon demo program `dp/10 Water Resource Management Model`.
- This demo solves the water management model assuming
  - $F(q) = \frac{a_1}{1+a_2} q^{1+a_2}$, $a_1 = 1$, $a_2 = -2$;
  - $U(s-q) = \frac{b_1}{1+b_2}(s-q)^{1+b_2}$, $b_1 = 2$, $b_2 = -3$;
  - ϵ are serially i.i.d. lognormal, mean 1, volatility $\sigma = 0.2$;
  - $\delta = 0.9$.
- The value function is approximated by a linear combination of $n = 15$ Chebychev polynomial basis functions on $[2, 8]$.
- Rainfall ϵ is discretized using an $m = 3$ node Gaussian quadrature scheme.

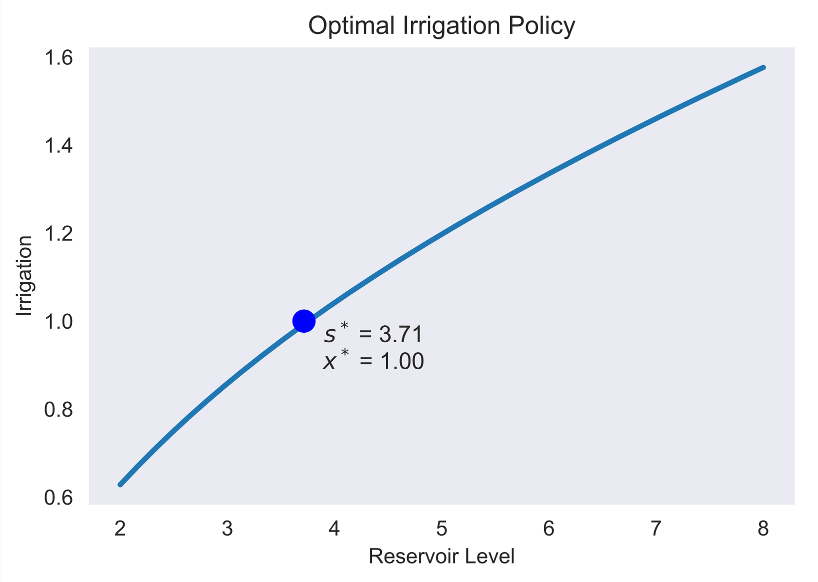

Figure 33: Optimal Irrigation Policy

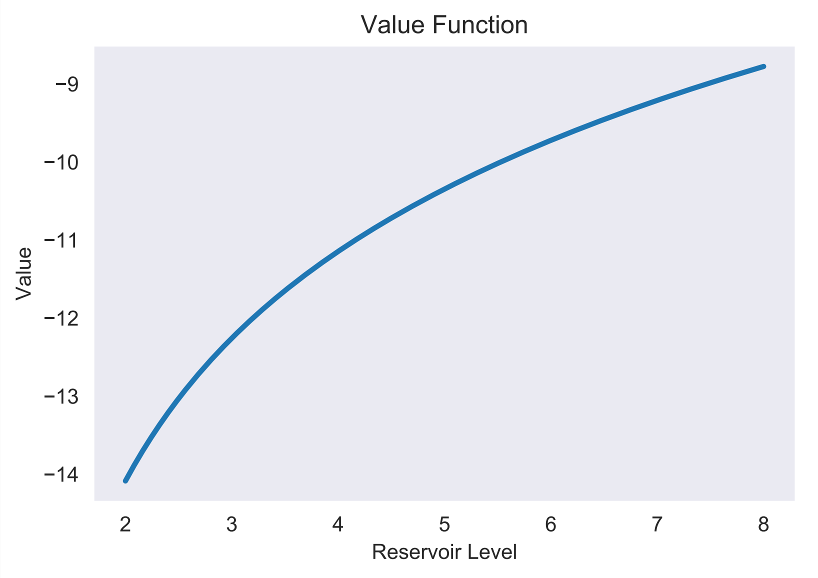

Figure 34: Value Function

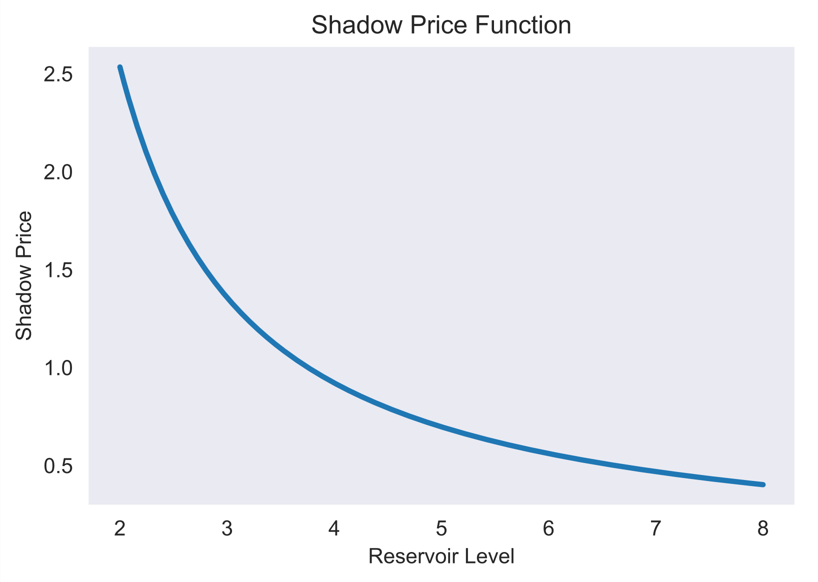

Figure 35: Shadow Price Function

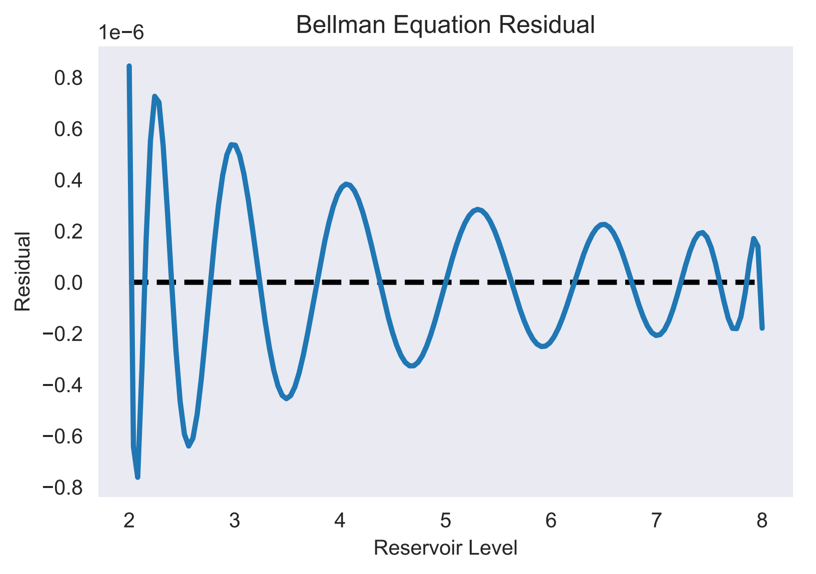

Figure 36: Bellman Equation Residual

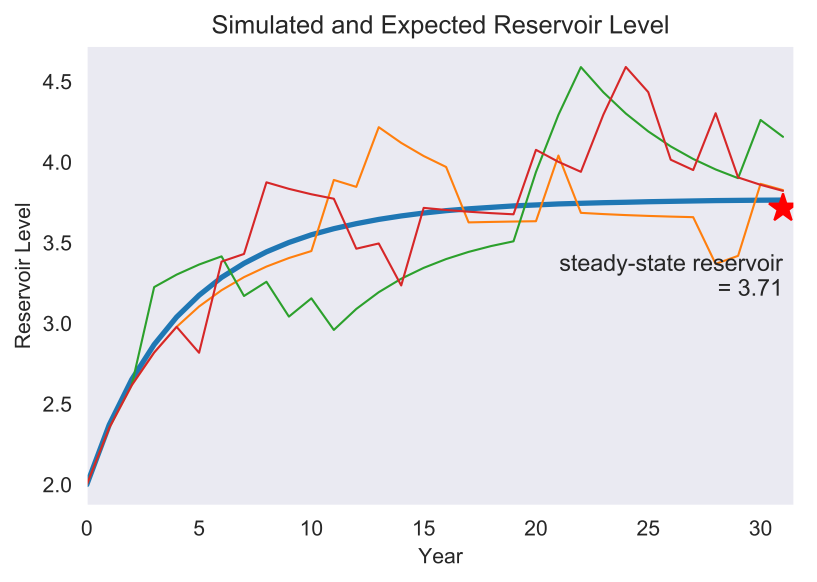

Figure 37: Simulated and Expected Reservoir Levels

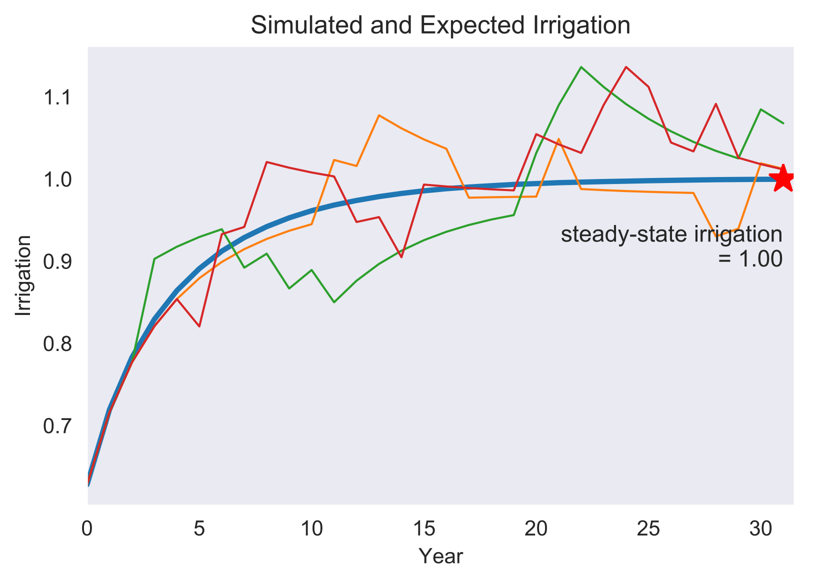

Figure 38: Simulated and Expected Irrigation

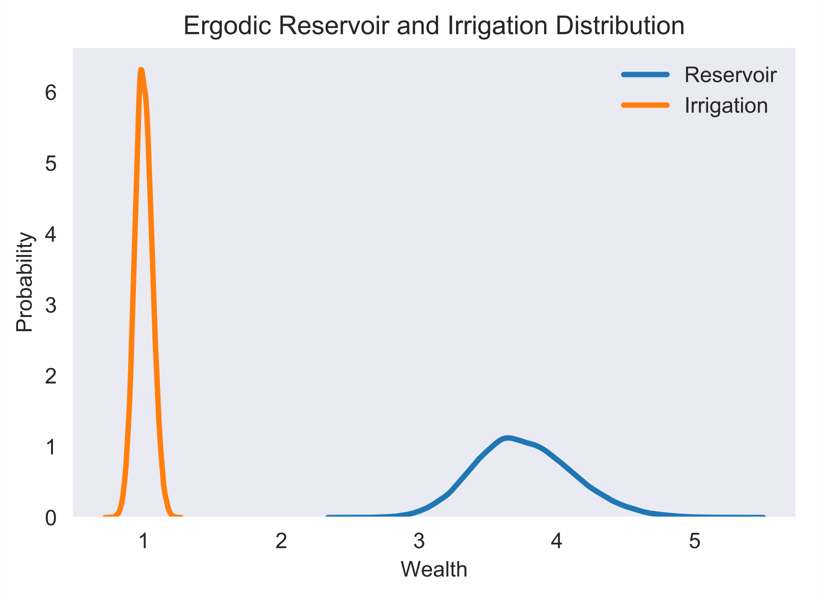

Figure 39: Ergodic Distribution of Reservoir Level

- The ergodic distribution of the reservoir level has mean 3.76 and standard deviation 0.36.
- How do these values change if ...
    - climate change raises mean rainfall by 20%?4.33 & 0.42.
    - climate change raises rainfall volatility to 25%?3.79 & 0.45.
    - climate change lowers rainfall volatility to 15%?3.74 & 0.27.
    - planner discounts future more, $\delta = 0.85$? 3.41 & 0.33.
    - planner doubles welfare weight on farmers? 3.21 & 0.34.
- How would you adapt the model and code to allow for additional fixed demand for residential use at a nearby town?

Higher-Dimensional Continuous Action Models

Model 7:

Monetary Policy

Jupyter notebook:

`dp/11 Monetary Policy.ipynb`

- A central bank wishes to manage the nominal interest rate so as to stabilize the gross domestic product (GDP) gap and inflation around specified targets.
- Each period t begins with a predetermined GDP gap s_{t1} and inflation rate s_{t2}, yielding a stabilization penalty

$$L(s_t) = \frac{1}{2}(s_t - \bar{s})'\Omega(s_t - \bar{s})$$

where s_t is a 2×1 vector containing the GDP gap and inflation rate, $\bar{s}$ is a 2×1 vector of targets, and Ω is a 2×2 constant positive definite matrix of preference weights.

(cont)

- The GDP gap and inflation rate evolve according to

$$s_{t+1} = \alpha + \beta s_t + \gamma x_t + \epsilon_{t+1}$$

  where α and γ are 2×1 constant vectors, β is a 2×2 constant matrix, and the ϵ are 2×1 serially i.i.d vectors with mean 0.
- What monetary policy minimizes the present value of current and expected future penalties, subject to the political constraint that nominal interest rate x_t be nonnegative?

Formulation

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

- Two continuous state variables, the GDP gap

$$s_{t1} \in (-\infty, \infty)$$

  and the inflation rate

$$s_{t2} \in (-\infty, \infty).$$

- One continuous action variable, the nominal interest rate

$$x_t \in [0, \infty).$$

- The "reward" is the negative of the weighted squared deviations

$$-L(s_t) = -\frac{1}{2}(s_t - \bar{s})'\Omega(s_t - \bar{s})$$

- State transitions are governed by

$$s_{t+1} = \alpha + \beta s_t + \gamma x_t + \epsilon_{t+1}$$

Bellman Equation

The sum of current and expected future rewards satisfies the Bellman equation

$$V(s) = \max_{x \geq 0} \{-L(s) + \delta E_{\epsilon} V(\alpha + \beta s + \gamma x + \epsilon)\}.$$

Conditions

- One cannot assume a priori that the nonnegativity constraint on the nominal interest rate will be nonbinding in all states.
- As such, the shadow price function $\lambda(s)$ is characterized by the Euler complementarity conditions

$$\mu \equiv \delta\gamma' E_\epsilon \lambda(g(s, x, \epsilon))$$

$$x \geq 0 \perp \mu \leq 0$$

$$\lambda(s) = -\Omega(s - \bar{s}) + \delta\beta' E_\epsilon \lambda(g(s, x, \epsilon))$$

where μ is the expected longrun marginal "reward".

- The Euler conditions imply that along the optimal path

$$x_t \geq 0 \perp \delta\gamma' E_t \lambda_{t+1} \leq 0$$

- That is, in any period, the nominal interest rate is reduced until either the expected longrun marginal reward or the nominal interest rate is driven to zero.

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

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

- The n coefficients c_j 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 nonlinear collocation equations

$$\sum_{j=1}^{n} c_j \phi_j(s_i) = \max_{x \geq 0} \left\{ -L(s_i) + \delta E_\epsilon \sum_{j=1}^{n} c_j \phi_j(\alpha + \beta s_i + \gamma x + \epsilon) \right\},$$

$i = 1, 2, \ldots, n$, for the n unknown coefficients c_j, $j = 1, 2, \ldots, n$.

- Open and run CompEcon demo program `dp/11 Monetary Policy Model`.
- This demo solves the monetary policy model assuming the ϵ are serially i.i.d. bivariate normal with mean 0 and variance matrix Σ, with $\delta = 0.9$ and

$$\bar{s} = \begin{bmatrix} 1 \\ 0 \end{bmatrix} \qquad \Omega = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

$$\alpha = \begin{bmatrix} 0.9 \\ -0.1 \end{bmatrix} \quad \beta = \begin{bmatrix} -0.5 & 0.2 \\ 0.3 & -0.4 \end{bmatrix}$$

$$\gamma = \begin{bmatrix} -0.1 \\ 0.0 \end{bmatrix} \quad \Sigma = \begin{bmatrix} 0.08 & 0.00 \\ 0.00 & 0.08 \end{bmatrix}.$$

- The value function is approximated by a linear combination of $n = 900 = 30 \times 30$ bivariate Chebychev polynomial basis functions on $[-2, -3] \times [2, 3]$.
- The shock vector ϵ is discretized using an $m = 9 = 3 \times 3$ node bivariate Gaussian quadrature scheme.

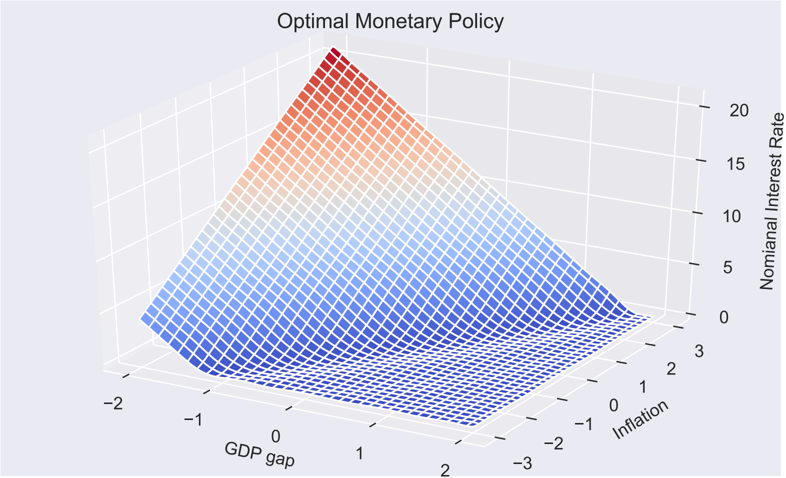

Figure 40: Optimal Monetary Policy

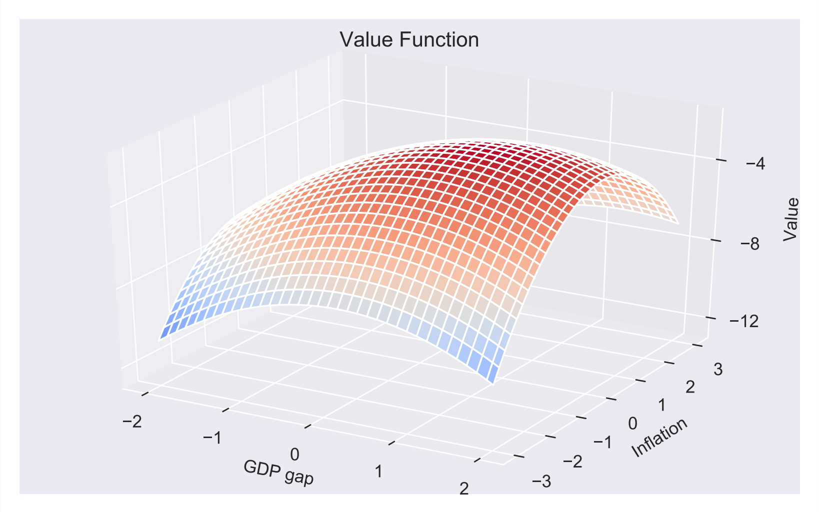

Figure 41: Value Function

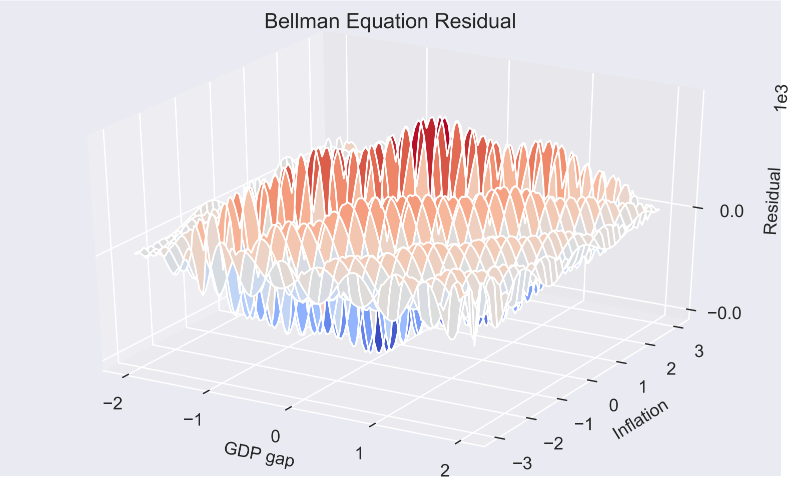

Figure 42: Bellman Equation Residual

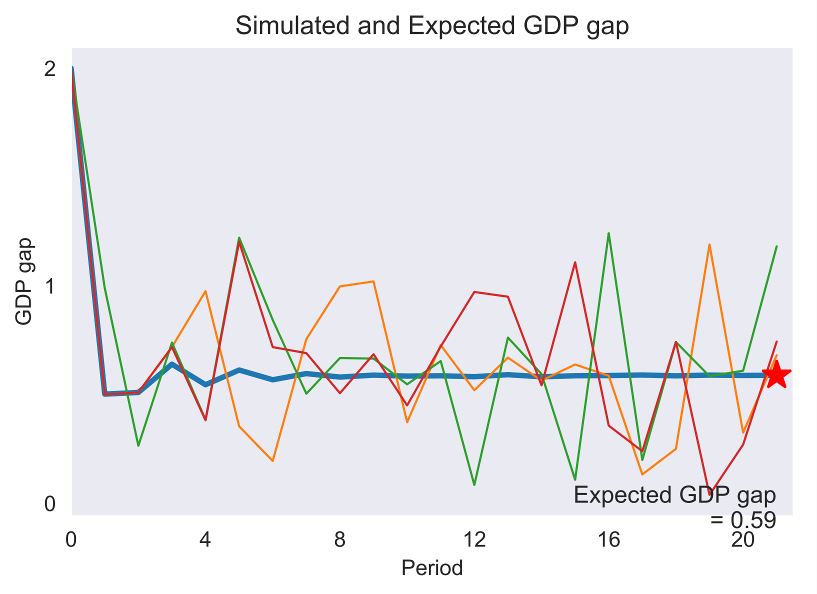

Figure 43: Simulated and Expected GDP Gap

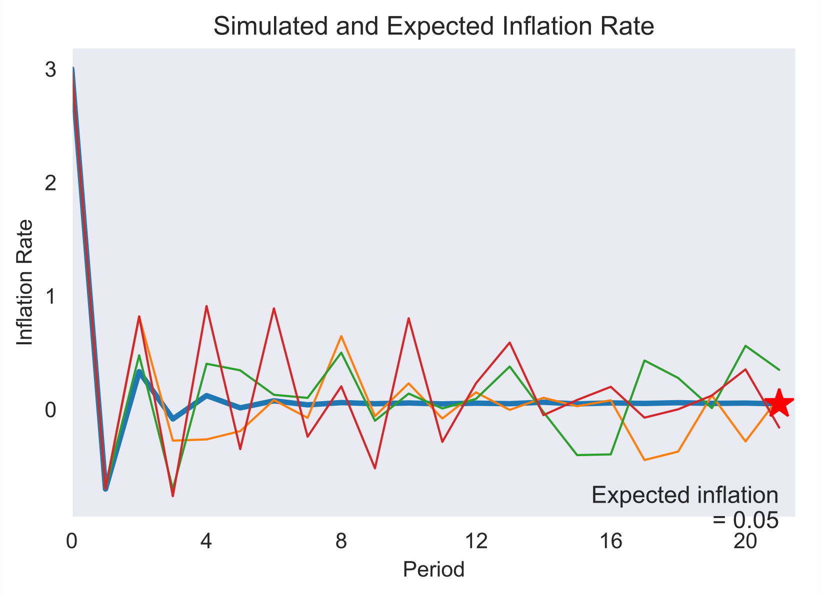

Figure 44: Simulated and Expected Inflation Rate

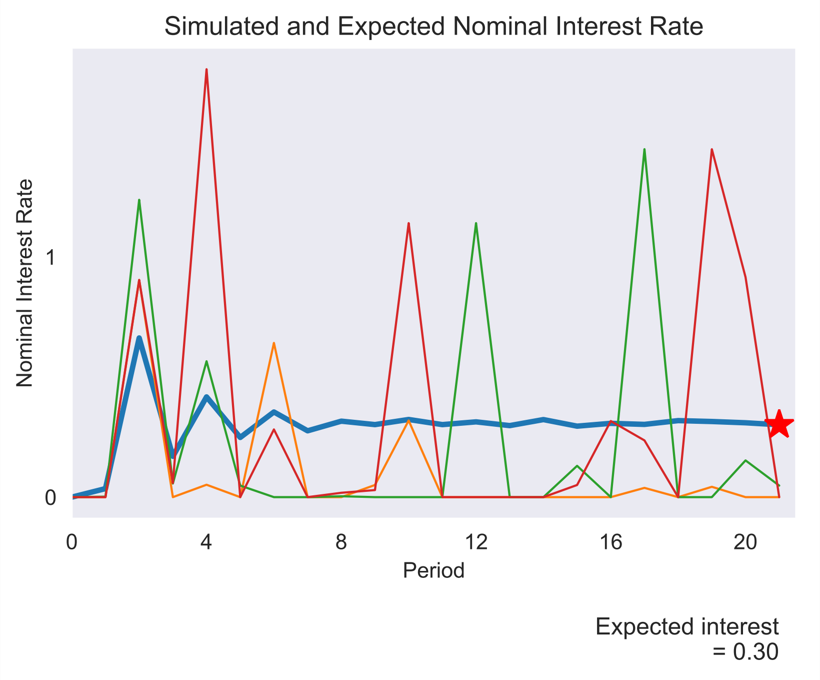
Simulated and Expected Nominal Interest Rate
Nominal Interest Rate
1
0
0
4
8
12
16
20
Period
Expected interest
= 0.30

- The ergodic distribution of the nominal interest rate has mean 0.27 and standard deviation 0.73.
- How do these values change if ...
  - variance of the shocks double?0.54 & 1.21.
  - variance of the shocks disappear?0.01 & 0.01.
  - GDP gap target is set to 0?8.47 & 1.68.
- How would you adapt the model and code to allow for penalties for excessive nominal interest rates?

References

Miranda, Mario J. and Paul L. Fackler (2002). *Applied Computational Economics and Finance*. MIT Press. ISBN: 0-262-13420-9.