A Comparison of Discrete and Parametric Methods for Continuous-State Dynamic Programming Problems

Hugo Benítez-Silva
George Hall
Günter Hitsch
Giorgio Pauletto
John Rust

<http://gemini.econ.yale.edu/jrust/sdp.html>

<http://www.econ.yale.edu/~econ681/dpa.pdf>

<http://www.econ.yale.edu/~econ681/dpa.ps>

Stanford Institute for Theoretical Economics 2000
Computational Methods for Dynamic Stochastic Models

Outline

- The Dynamic Programming Problem
- Solution Algorithms
  - Discrete Policy Iteration
  - Parametric Policy Iteration
  - Parameterized Expectations Method
  - Projection and Minimum Residual Methods
- Test Problems
  - Consumption/Saving Problem
  - Consumption/Saving Labor/Leisure Problem
  - Stochastic Growth with Leisure
  - Inventory Investment Model

Solving the DP Problem

Infinite horizon DP problem with Bellman's equation

$$\Gamma(V)(s) = V(s) = \max_{a \in A(s)} \left[u(s,a) + \beta \int V(s')\, p(s'|s,a) ds' \right]$$

and the optimal policy is

$$\alpha(s) = \operatorname*{argmax}_{a \in A(s)} \left[u(s,a) + \beta \int V(s')\, p(s'|s,a) ds' \right]$$

The Policy Iteration algorithm solves the nonlinear equation

$$(I - \Gamma)(V) = 0$$

using Newton's method

Policy Iteration Algorithm

1. **Policy Valuation Step:** given an initial guess of policy α compute $V_\alpha(s)$

$$V_\alpha(s) = u(s, \alpha(s)) + \beta E_\alpha V_\alpha(s)$$

where

$$E_\alpha V(s) = \int V(s')\ p(s'|s, \alpha(s))ds'$$

i.e. solve for V_α

$$(I - \beta E_\alpha)V_\alpha = u_\alpha$$

2. **Policy Improvement Step:** compute an improved policy α' using V_α

$$\alpha'(s) = \underset{a \in A(s)}{\text{argmax}} \left[u(s, a) + \beta \int V_\alpha(s')\ p(s'|s, a)ds' \right]$$

Discrete Policy Iteration (DPI)

Idea: Discretize the state space by a finite set $S_N = \{s_1, \ldots, s_N\}$

1. **Policy Valuation Step:** given α compute $V_{\alpha,N}$

$$V_{\alpha,N}(s) = u(s,\alpha(s)) + \beta E_{\alpha,N} V_{\alpha,N}(s)$$

$$\text{where} \quad E_{\alpha,N} V(s) = \sum_{i=1}^{N} V(s_i)\, p_N(s_i|s,\alpha(s)) ds'$$

   i.e. solve the $N \times N$ system of linear equations

$$(I - \beta E_{\alpha,N}) V_{\alpha,N} = u_\alpha \quad \text{where} \quad (E_{\alpha,N})_{ij} = p_N(s_i|s_j,\alpha(s_j))$$

2. **Policy Improvement Step:**

$$\alpha'(s) = \operatorname*{argmax}_{a \in A(s)} \left[u(s,a) + \beta \sum_{i=1}^{N} V_{\alpha,N}(s_i)\, p_N(s_i|s,a) \right]$$

Parametric Policy Iteration (PPI)

Idea: Parameterize the value function $V_\alpha(s) \approx \sum_{i=1}^{k} \theta_i \, \rho_i(s) \equiv V_\theta(s)$

1. **Policy Valuation Step:** given an initial guess of policy α compute $V_{\widehat{\theta}}(s)$
   Estimate θ by OLS using $M \geq k$ points

$$u(s, \alpha(s)) = \sum_{i=1}^{k} \theta_i \, \rho_i(s) - \beta \int \sum_{i=1}^{k} \theta_i \, \rho_i(s') \, p(s'|s,a) ds'$$

$$y = (P - \beta \, EP)\theta$$

$$\widehat{\theta} = (P - \beta \, EP) \backslash y$$

   Compute $\quad V_{\widehat{\theta}}(s) = \sum_{i=1}^{k} \widehat{\theta}_i \, \rho_i(s)$

2. **Policy Improvement Step:** compute an improved policy α'

$$\alpha'(s) = \operatorname*{argmax}_{a \in A(s)} \left[u(s,a) + \beta \int \sum_{i=1}^{k} \widehat{\theta}_i \, \rho_i(s') \, p(s'|s,a) ds' \right]$$

Minimum Residual Method on the Euler Equation

Idea: parameterize the Euler equation of the problem

$$\frac{\partial u(s,a)}{\partial a} + \beta \underbrace{\int \frac{\partial u(s',a')}{\partial s} \frac{\partial p(s'|s,a)}{\partial a}\, ds'}_{\approx \sum_{i=1}^{k} \theta_i\, \rho_i(s)} = 0$$

Choose θ that minimizes a weighted residual function

$$\int w(s)\; R(s, a(s)|\theta)\; ds = 0$$

here

$$\frac{\partial u(s,a)}{\partial a} + \beta \sum_{i=1}^{M} \left(w(s_i) \frac{\partial p(s_i|s,a)}{\partial a} \sum_{j=1}^{k} \theta_j \rho_j(s_i) \right) = 0$$

Choice of integration method, sample points s_i, basis functions ρ_j

Wright and Williams (1982), den Haan and Marcet (1990),
Judd (1992), Judd and Gaspar (1997), Christiano and Fisher (2000)

Projection and Minimum Residual Methods 1

Judd (1992), Judd and Gaspar (1997)

Idea: parameterize the value function (or some other function) and minimize a residual function

1. Parameterize $V_\theta(s) = \sum_{i=1}^{k} \theta_i p_i(s)$ such that $\lim_{k\to\infty} \inf_{\theta \in R^k} \|V_\theta - V\| = 0$
2. Define a residual operator $R(V)(s) = V(s) - \Gamma(V)(s)$

Projection and Minimum Residual Methods 2

3. Determine $\widehat{\theta}$ so that $R(V_{\widehat{\theta}})$ is minimal:

   (a) Projection methods:

   Solve $< R(V_{\widehat{\theta}}), \phi_i >_\mu = 0$

   Galerkin ($\phi_i = p_i$)   Collocation ($\phi_i = \delta_{s_i}$)

   (b) Minimum residual methods:

   Solve $\widehat{\theta} = \mathrm{argmin}_{\theta \in R^k} \|R(V_\theta)\|_p$

Methodological Choices for Solving Dynamic Economic Models	
Discrete Approximation	Parametric Approximation
Grids – uniform – random – low discrepancy	*Function to Parameterize* – value function – decision rule/expected value function
Integration – quadrature – Monte Carlo – low discrepancy	*Basis* – Chebychev – poly-log – neural networks
Smoothing – bilinear interp/simplicial – linear regression – probability weighted	*Fitting Technique* – linear regression – nonlinear regression
	Integration – quadrature – Monte Carlo – low discrepancy

Consumption/Saving Model 1

Phelps (1962), Levhari and Srinivasan (1969), Hakansson (1970)

$$\text{Phelps problem} \quad \begin{cases} \max\limits_{0 \le c_s \le w_s} E_t \left[\sum\limits_{s=t}^{\infty} \beta^{(s-t)} u(c_s) \right] \\ \quad \text{s.t.} \\ w_{t+1} = (1 + r_t)(w_t - c_t) \\ \quad r_t \sim \text{iid}(F) \end{cases}$$

Bellman's equation

$$V(w) = \max_{0 \le c \le w} \left[\log c + \beta \int_0^{\infty} V(R(w-c)) F(dR) \right]$$

C/S Model 2

- Analytical solution:
  1. log utility, optimal consumption is $(1-\beta)w$ independent of F
  2. Constant relative risk aversion CRRA $u(c) = (c^{1-\gamma})/(1-\gamma)$, and constant absolute risk aversion CARA $u(c) = -\exp\{-\gamma c\}$ $\rightarrow$ closed-form solutions
- Numerical solution: DPI, PPI

C/S Model 3 (DPI)

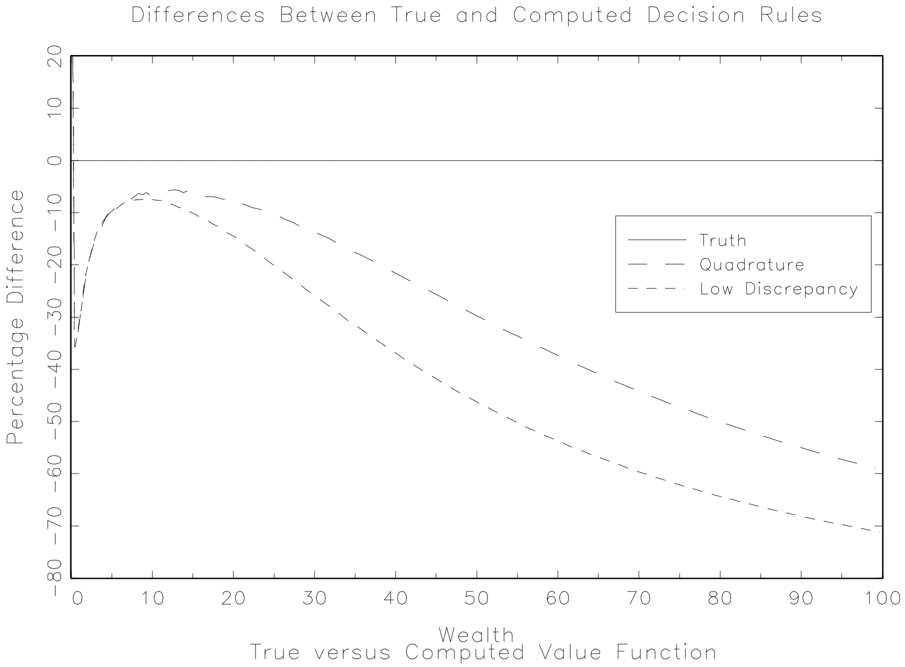

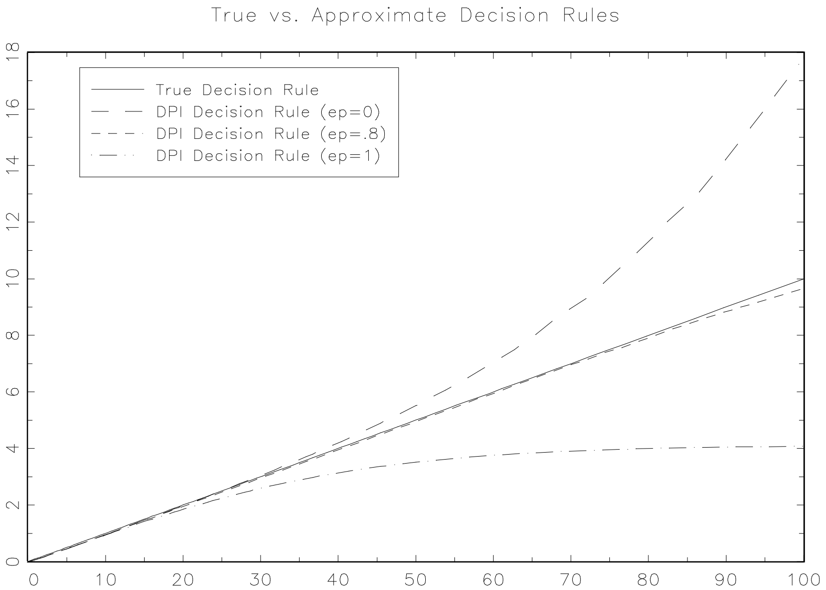

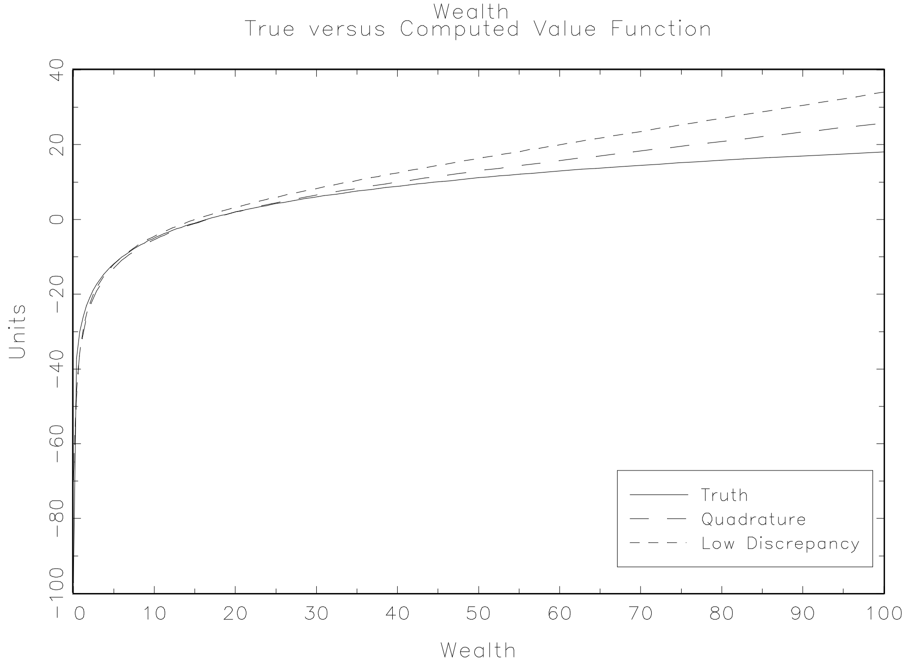

C/S Model 4 (PPI)

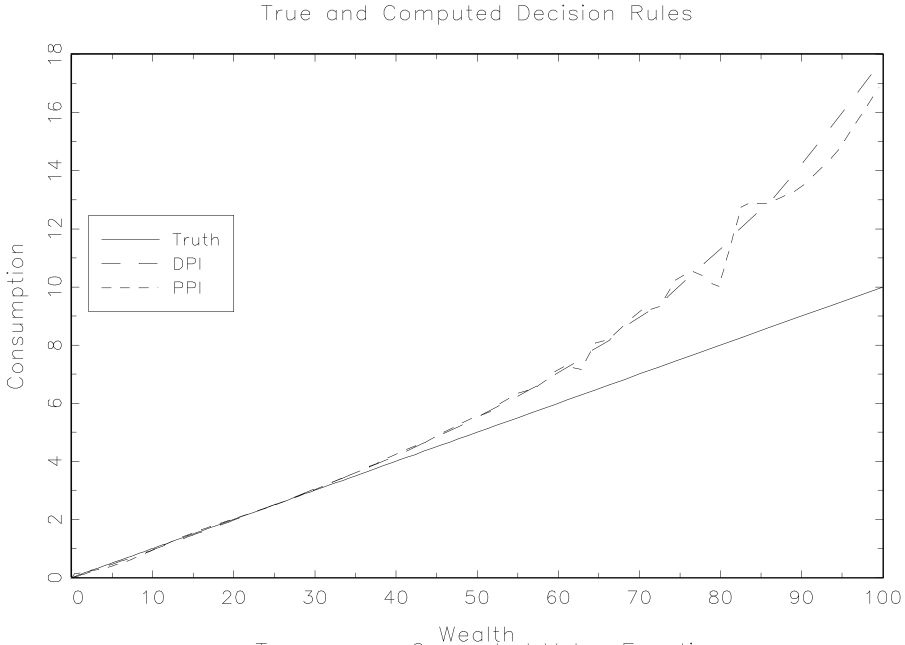

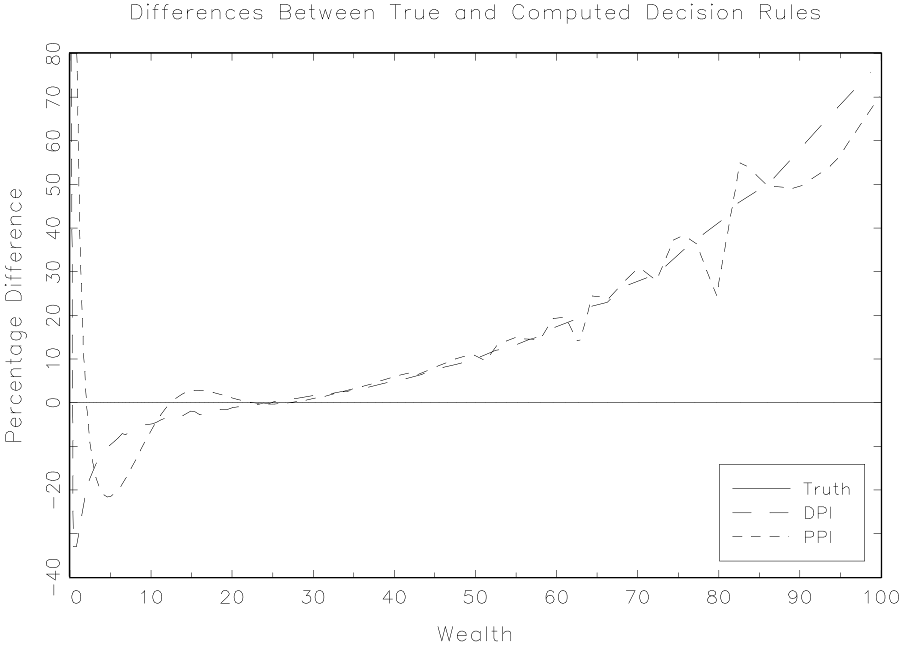

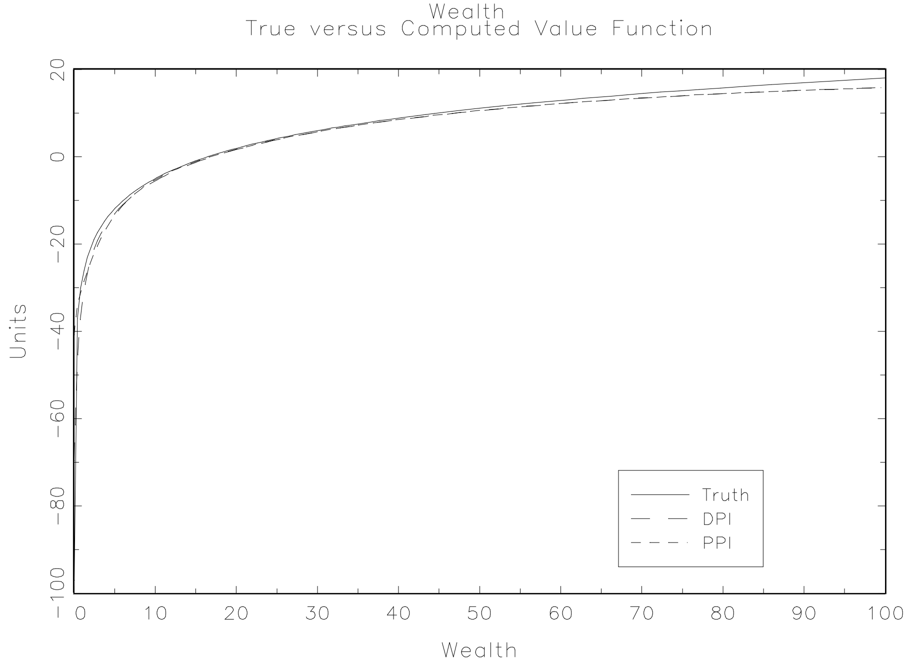

Phelps pb finite horizon 1

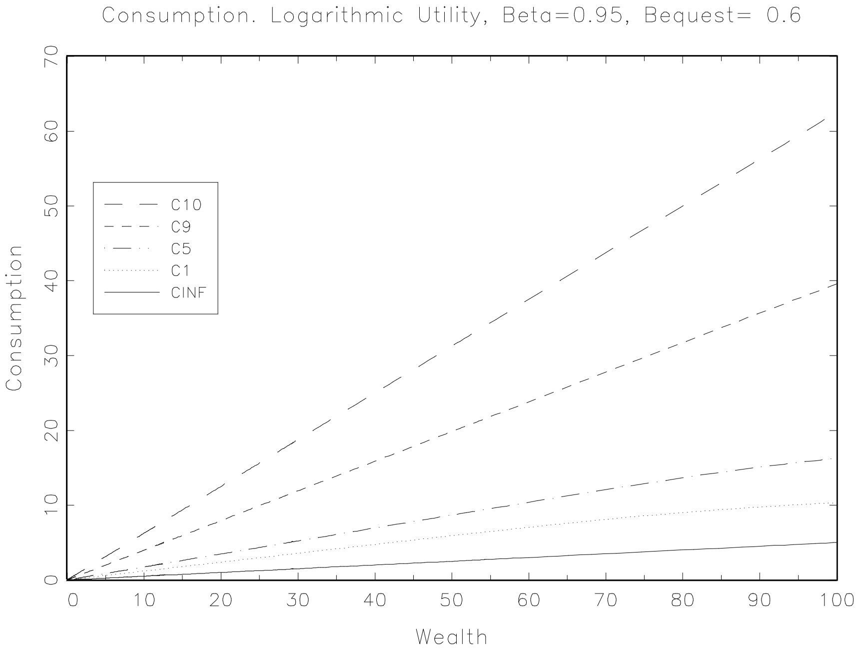

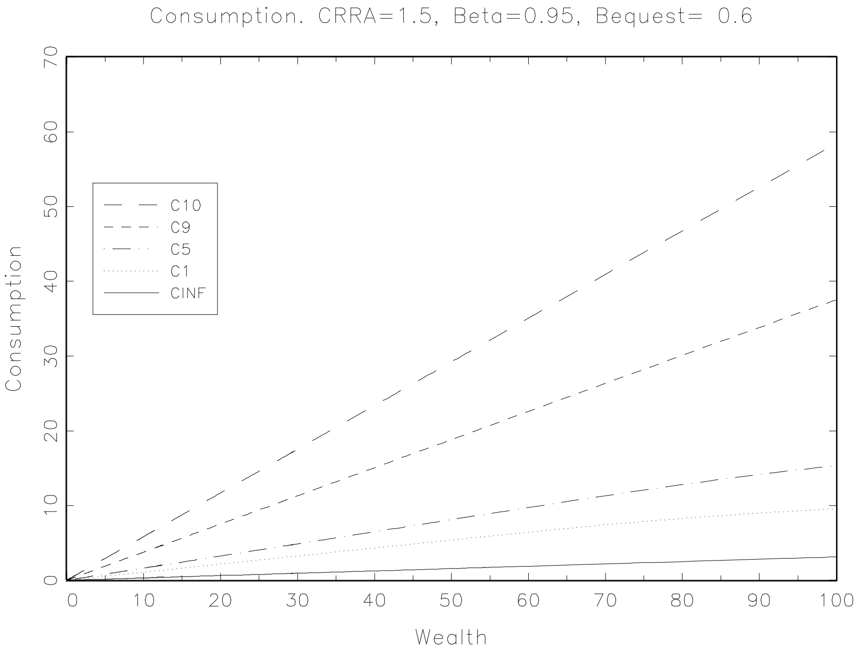

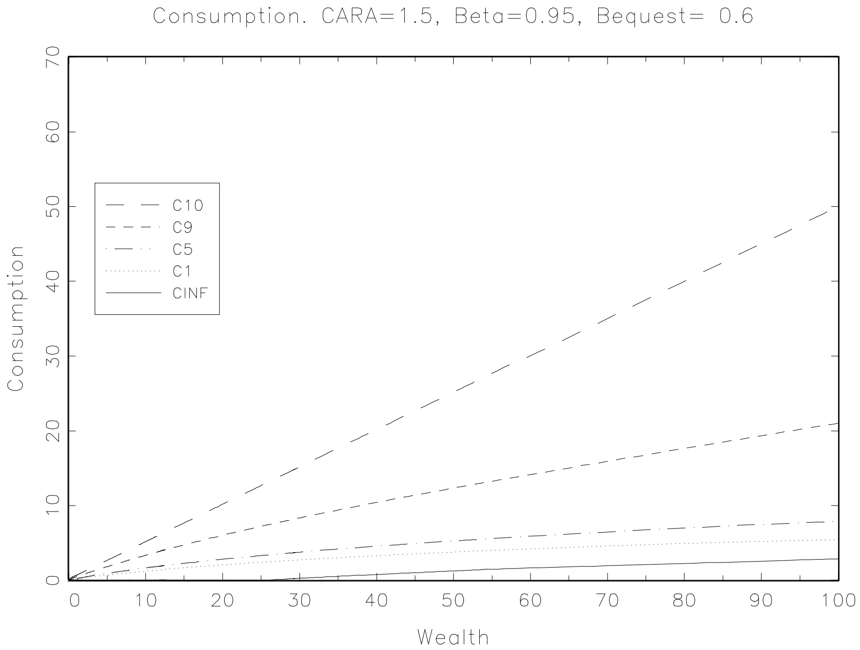

Phelps pb finite horizon 2

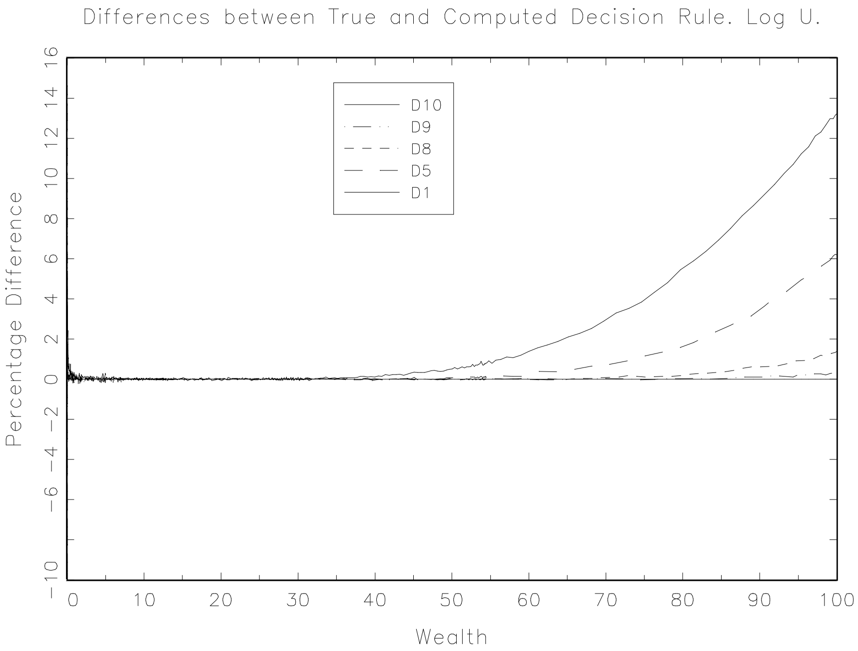

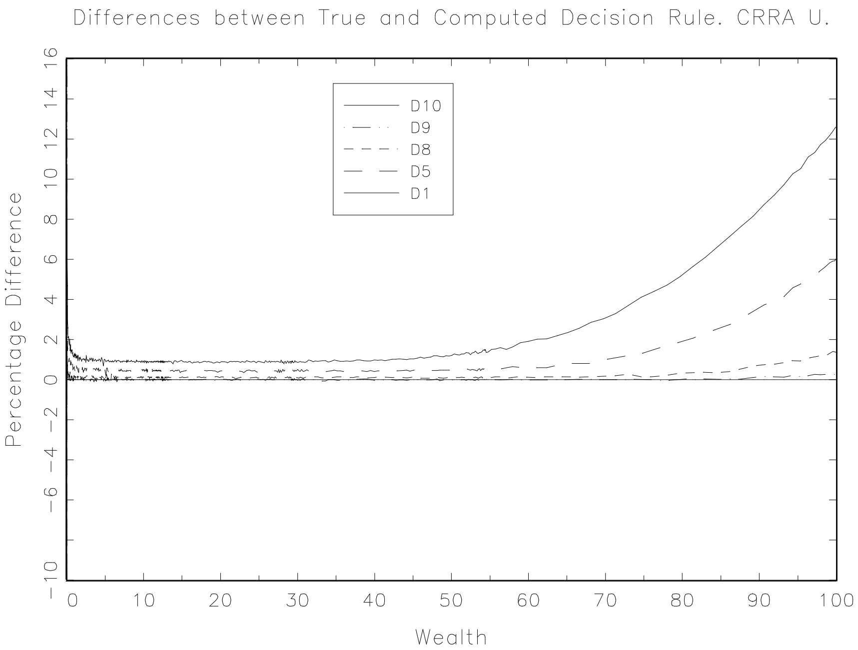

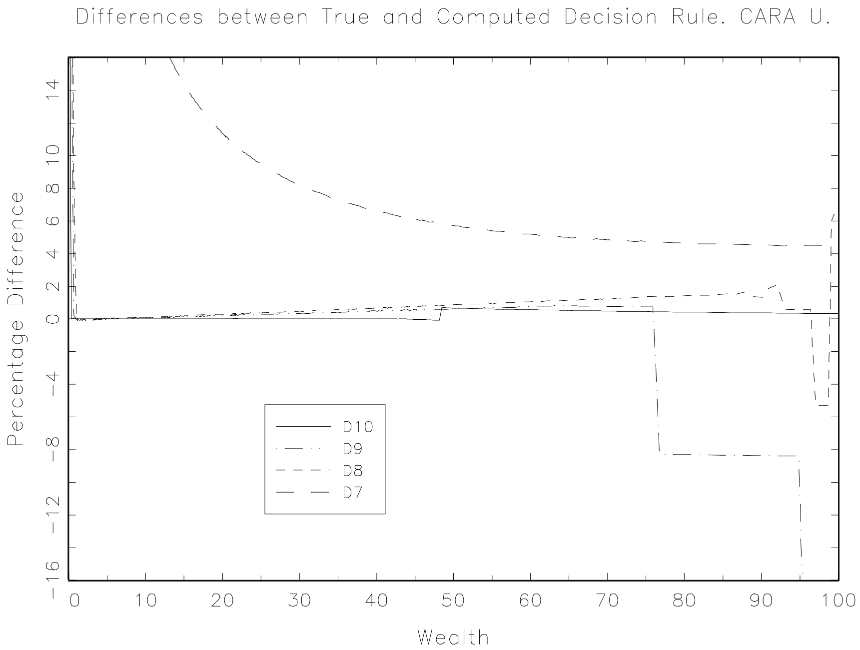

Consumption/Saving Labor/Leisure Model 1

$$\text{Including Labor/Leisure} \left\{ \begin{aligned} &\max_{c_s, l_s} E_t \left[\sum_{s=t}^{T} \beta^{(s-t)} u(c_s, l_s) \right] \\ &\quad \text{s.t.} \\ w_{t+1} &= (1 + r_t)(w_t + (1 - l_t)\omega_t - c_t) \\ \log \omega_t &= (1 - \rho)\alpha_t + \rho \log \omega_{t-1} + \epsilon_t \\ \epsilon_t &\sim N(0, \sigma_t^2) \\ r_t &\sim \text{iid}(F) \end{aligned} \right.$$

Introducing Labour/Leisure

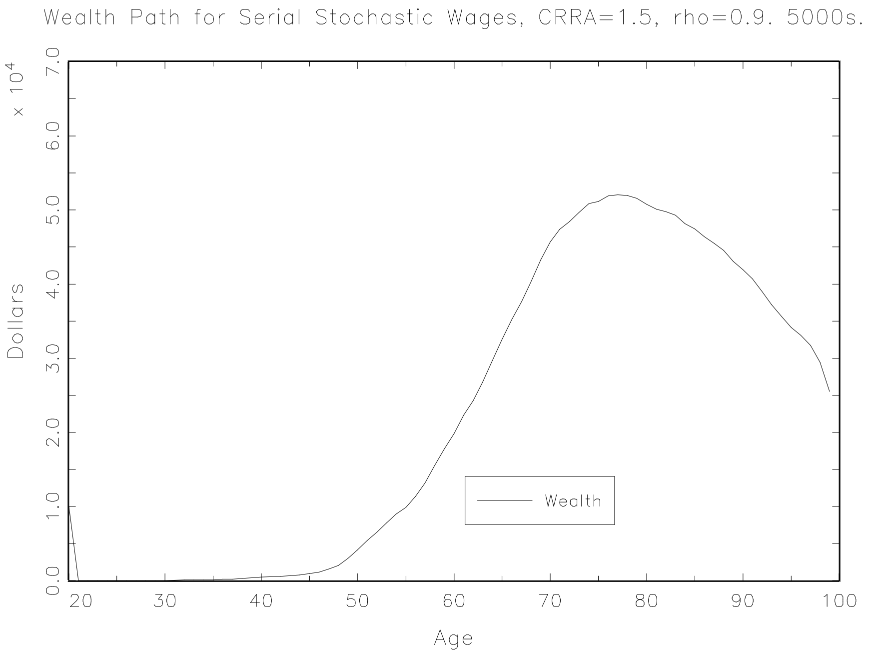

Stochastic Growth Model with Leisure

Taylor and Uhlig (1990), Judd (1992), Judd and Gaspar (1997), Santos (1999)

$$
\begin{aligned}
\max_{l_t, c_t, i_t} E_0 \sum_{t=0}^{\infty} \beta^t \lambda \log c_t + (1-\lambda) \log l_t \\
\text{s.t.} \\
c_t + i_t &= z_t A k^{\alpha} (1 - l_t)^{(1-\alpha)} \\
k_{t+1} &= (1-\delta) k_t + i_t \\
\log z_{t+1} &= \rho \log z_t + \epsilon_{t+1} \text{ with } \epsilon_t \sim N(0, \sigma^2)
\end{aligned}
$$

Growth Model using DPI 1

True V

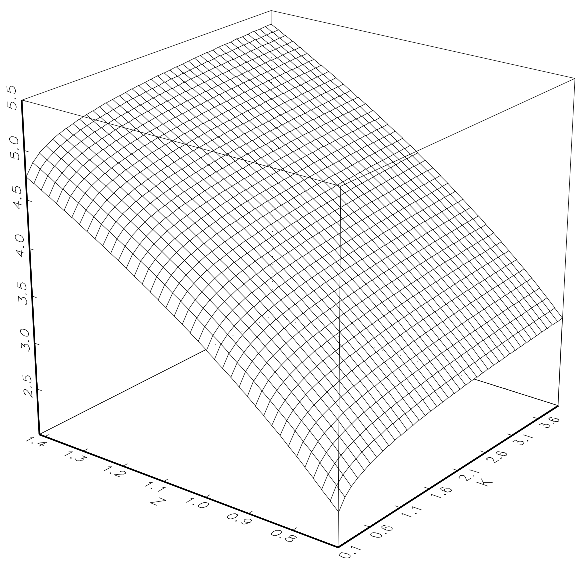

Computed V

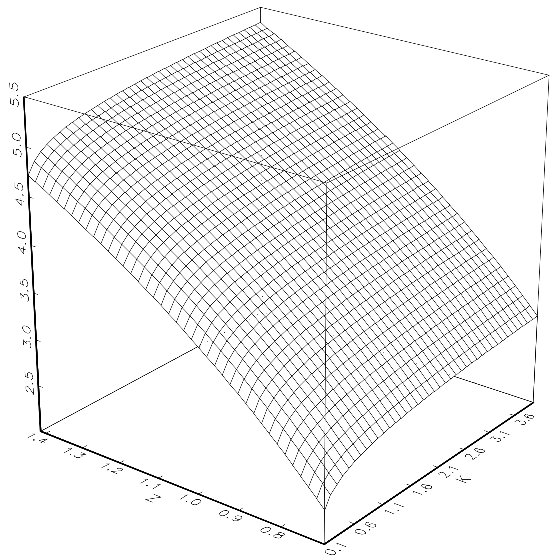

Difference between True and Computed V

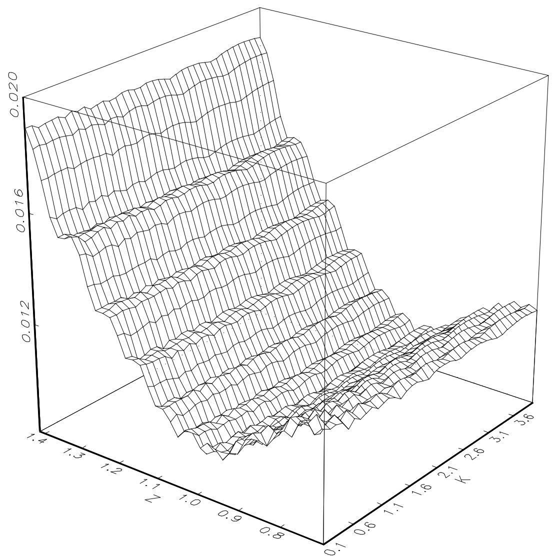

Growth Model using DPI 2

True C

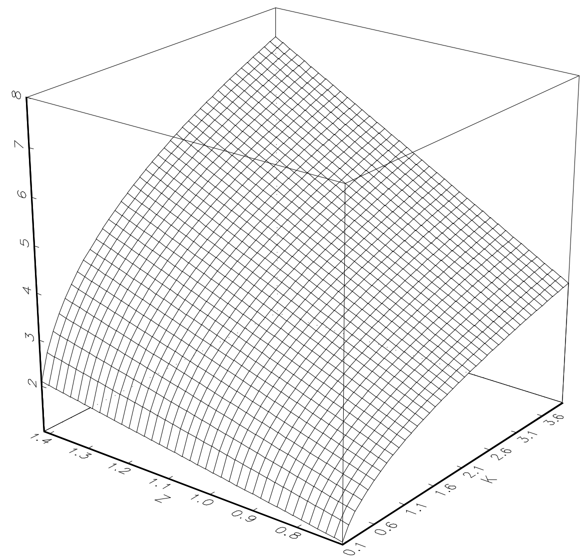

Computed C

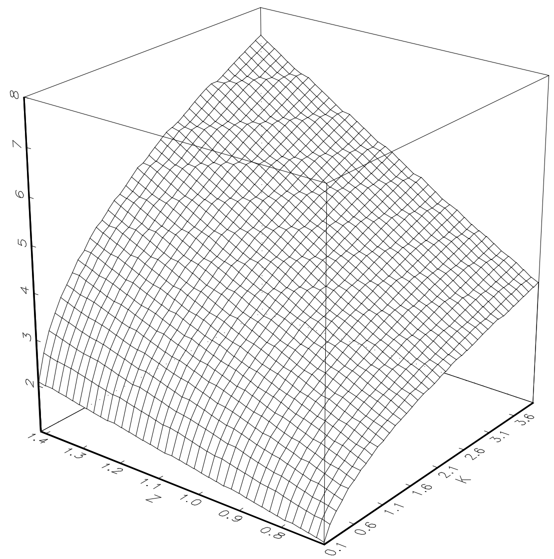

Difference between True and Computed C

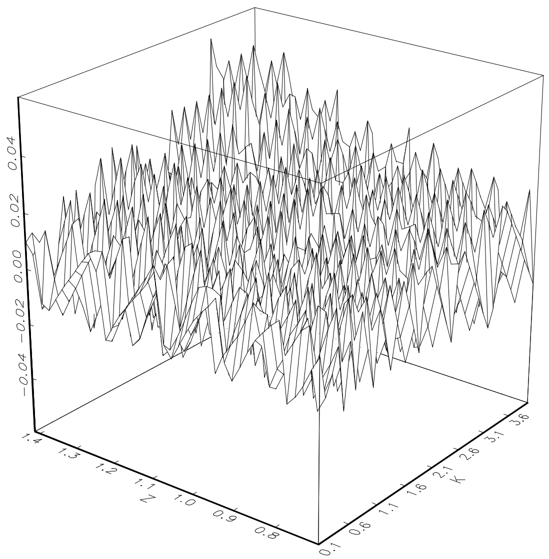

Growth Model using DPI 3

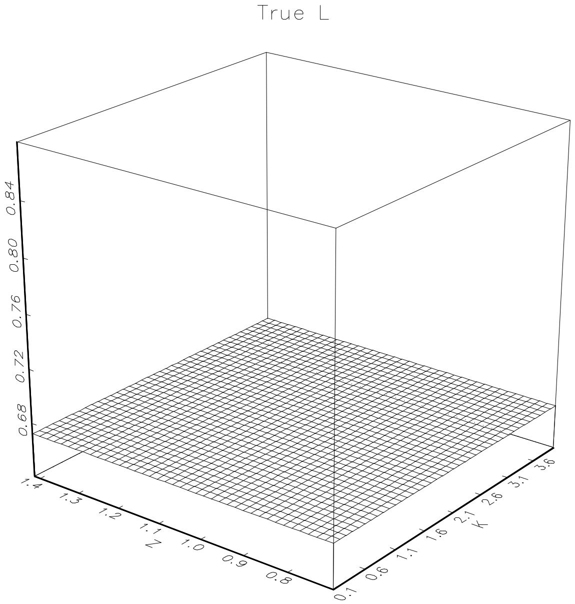

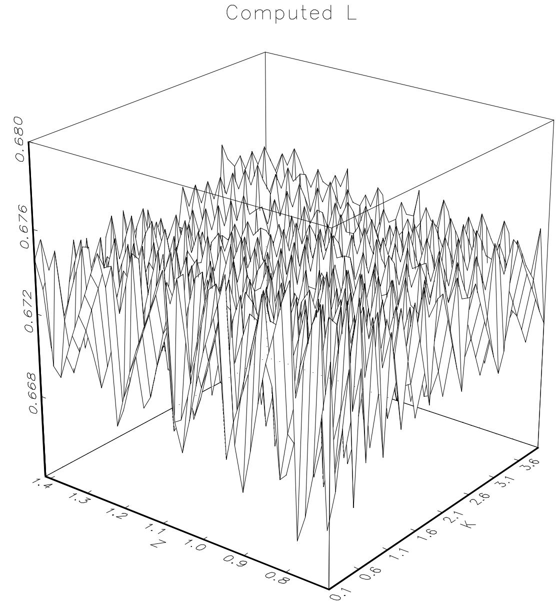

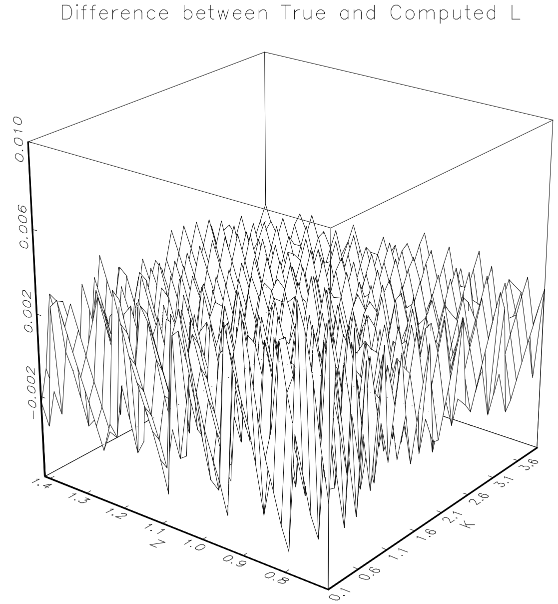

Growth Model using DPI 4

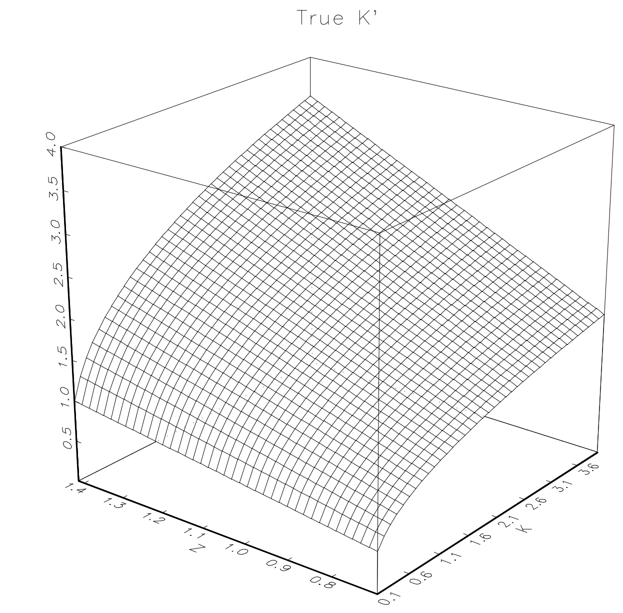

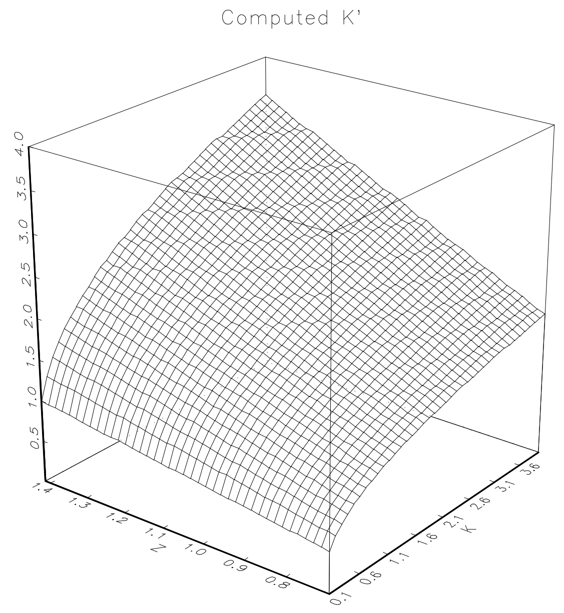

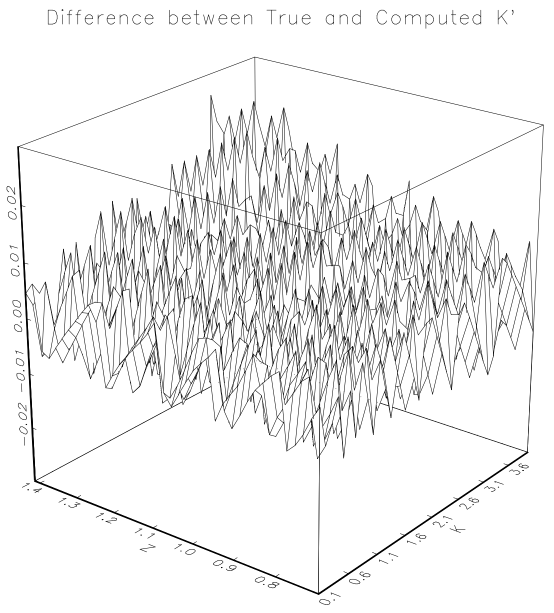

Growth Model using PPI 5

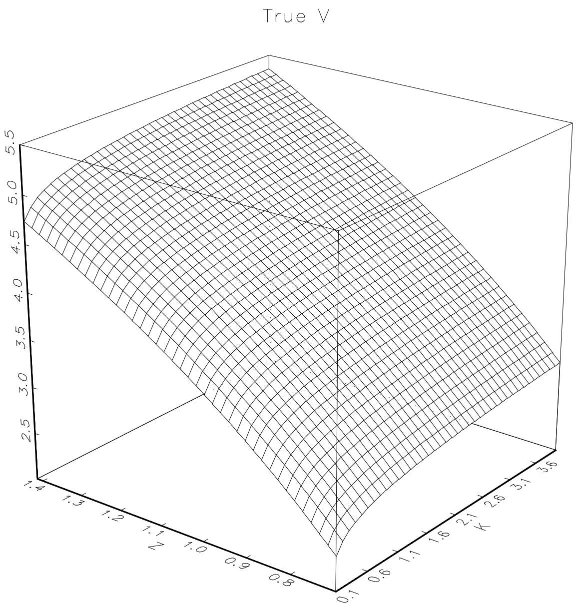

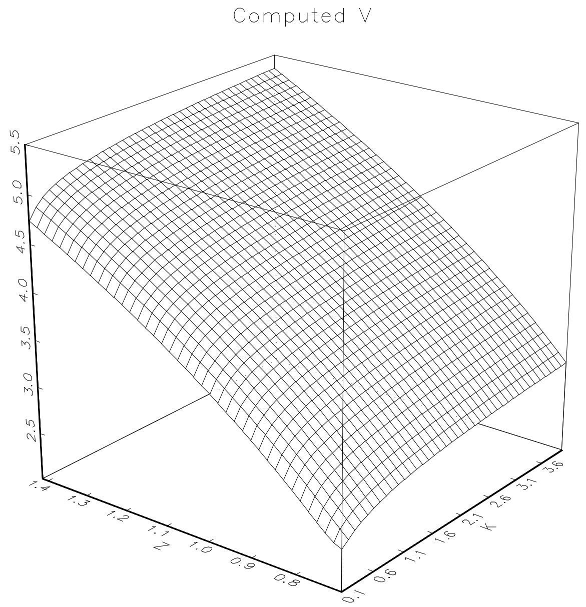

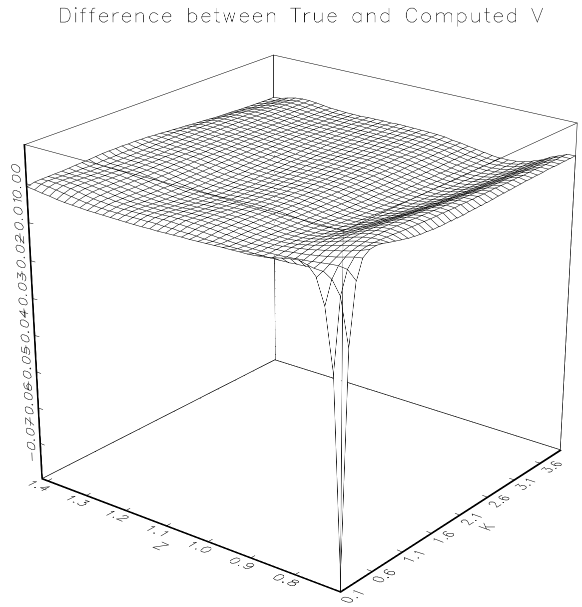

Growth Model using PPI 6

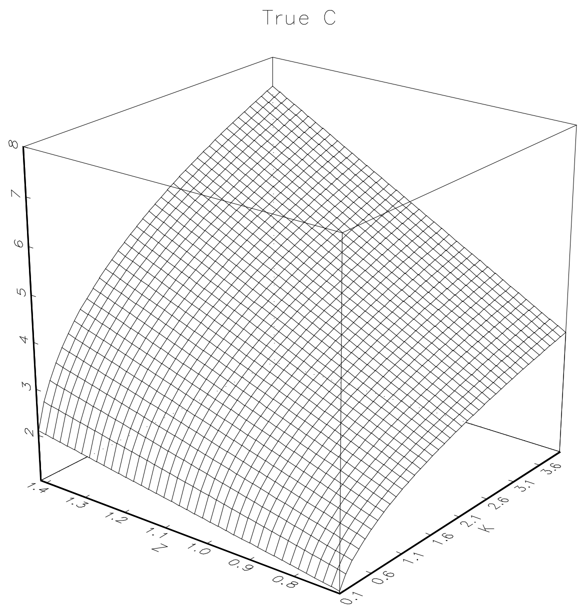

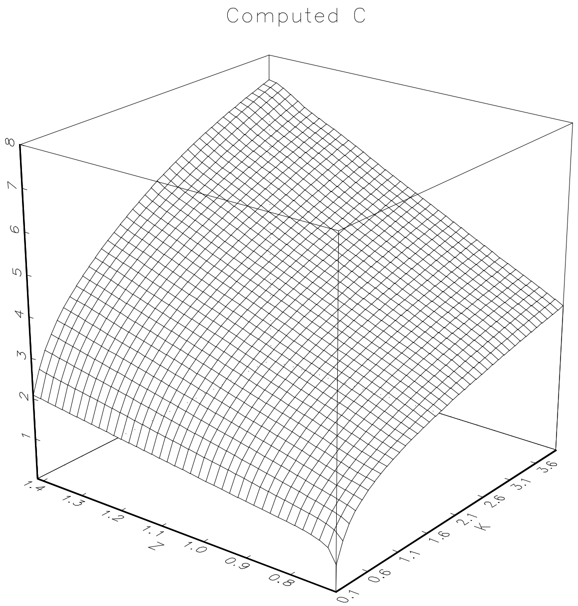

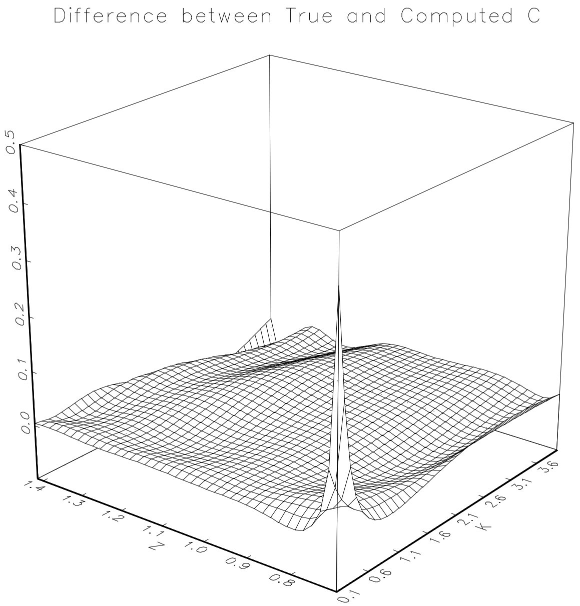

Growth Model using PPI 7

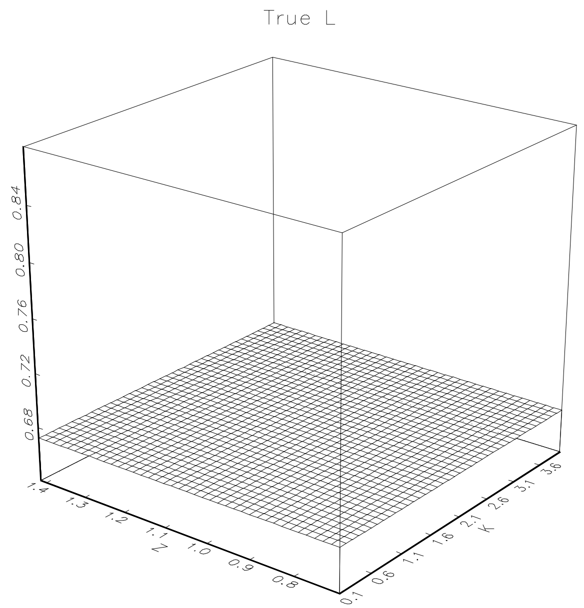

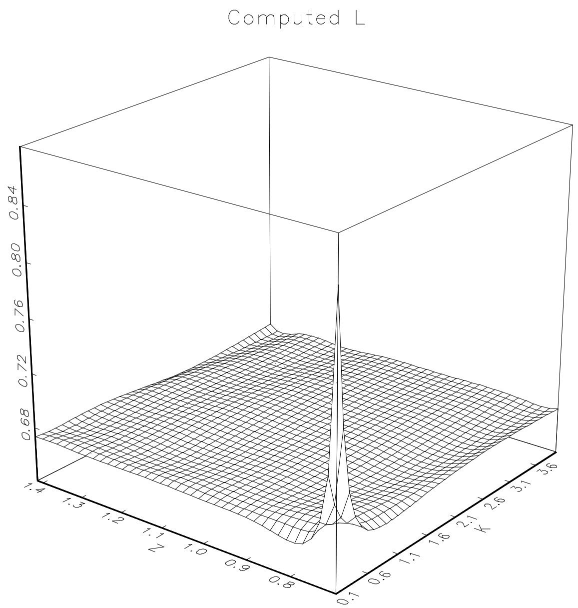

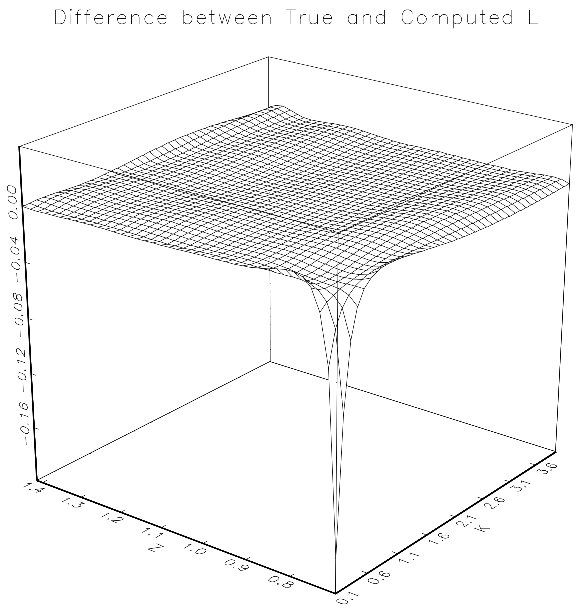

Growth Model using PPI 8

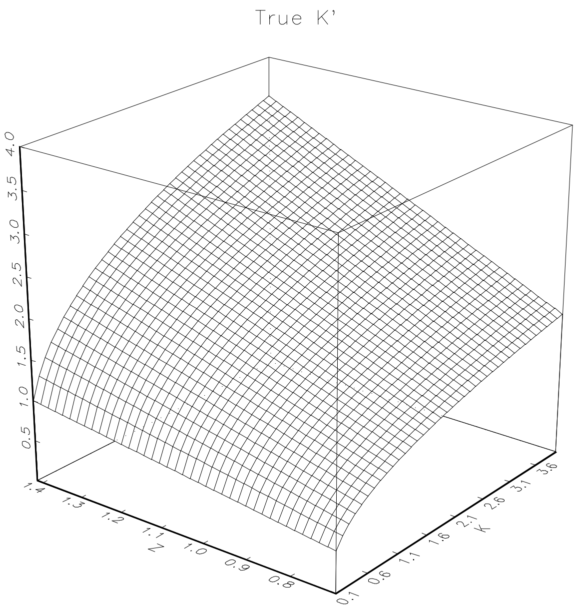

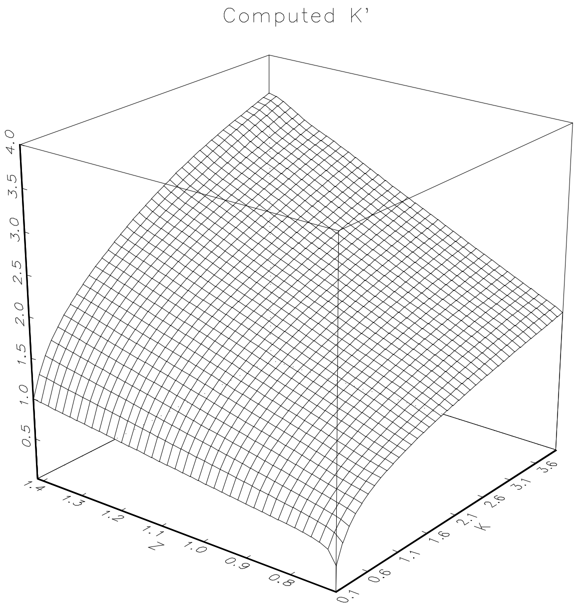

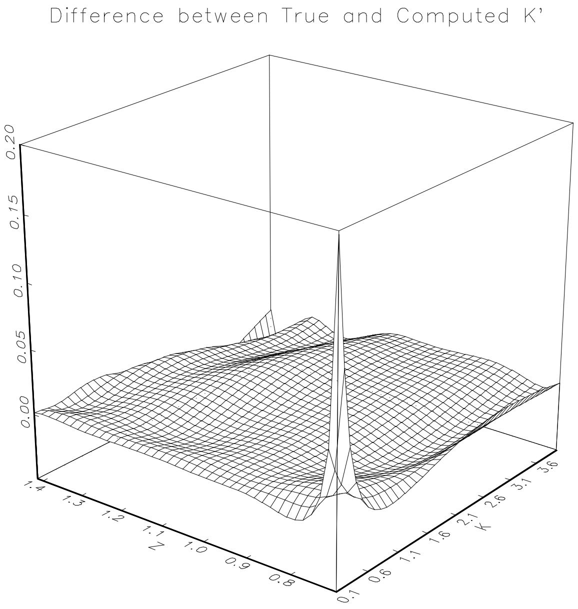

Growth Model using PEA

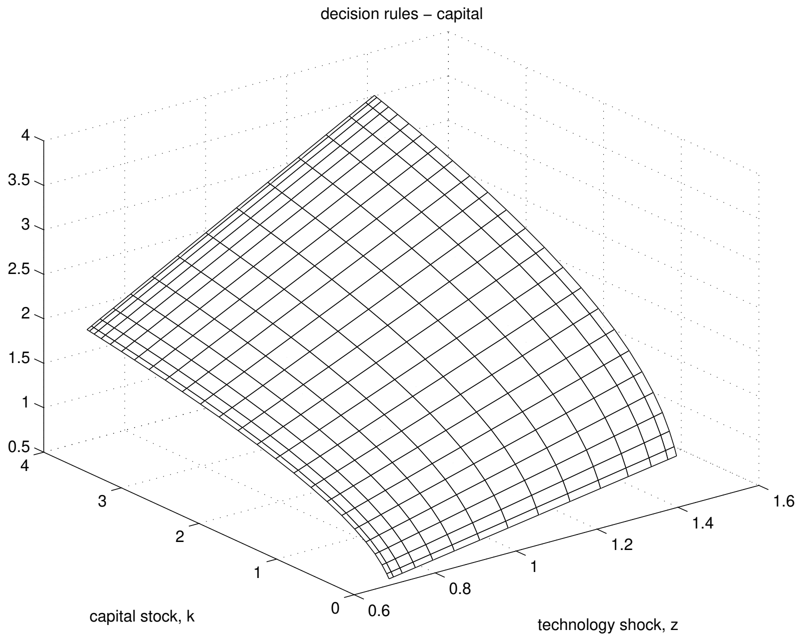

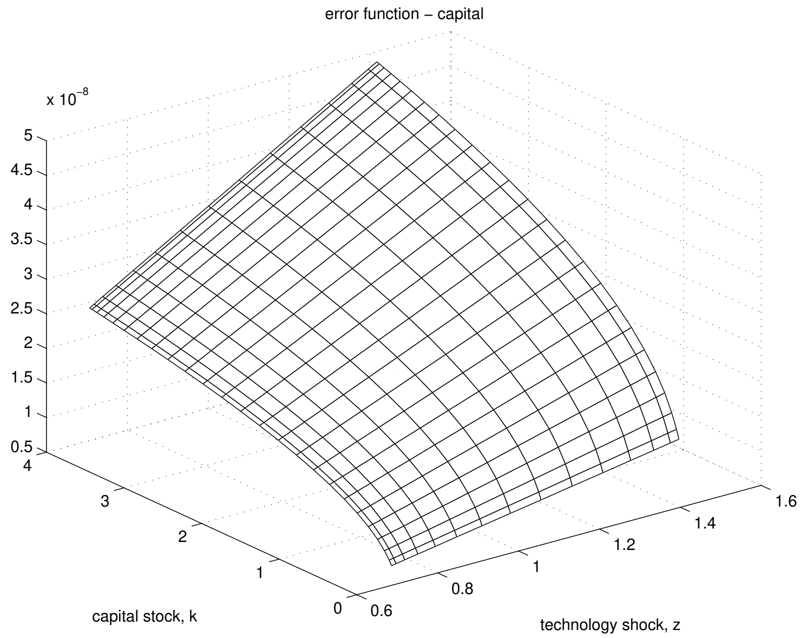

Inventory Investment Model 1

- Intermediary buys and sells a durable commodity in discrete time
- State variables (p_t, q_t) where p_t follows a Markov process with transition density $g(p_{t+1}, p_t)$
- Intermediary observes (p_t, q_t), orders q_t^o, sets uniform sales price $p_t^s = \alpha_0 + \alpha_1 p_t + \epsilon_t$ with $\alpha_0, \alpha_1 > 0$ and $E\{\epsilon_t | p_t\} = 0$
- After receiving q_t^o and setting p_t^s the intermediary observes q_t^d with distribution $H(q_t^d | p_t)$ with possible mass point at $q^d = 0$ i.e.

$$H(q^d|p) = H(0|p) + (1 - H(0|p)) \int_0^{q^d} h(q'|p) dq'$$

- The realized sales are $q_t^s = \min\left[q_t + q_t^o, q_t^d\right]$
- Law of motion for inventory holdings is $q_{t+1} = q_t + q_t^o - q_t^s$

Inventory Investment Model 2

- There may be unfilled demand $q_t^s < q_t^d$ with probability $1 - H(q + q^o|p)$
- Or stockouts $q_t^d > q_t^s$ i.e. $q_{t+1} = 0$
- The expected sales revenue is

$$ES(p, q, q^o) = p^s E\{q^s|q, q^o)$$
$$= [1 - H(0|p)] \left[\int_0^{q+q^o} q^d h(q^d|p) dq^d + (1 - H(q + q^o))(q + q^o) \right]$$

- The intermediary's cost for ordering inventory is

$$c^o(p, q^o) = \begin{cases} K + pq^o & \text{if} \quad q^o > 0 \\ 0 & \text{otherwise} \end{cases}$$

Inventory Investment Model 3

- The profit is given by $\pi(p_t, q_t^s, q_t, q_t^o) = p_t^s q_t^s - c^o(p_t, q_t^o)$
- The optimal rule is

$$V(p_t, q_t) = \max_{q^o} E\left\{\sum_{j=t}^{\infty} \beta^{(j-t)} \pi(p_j, p_j^s, q_j^s, q_j^o) | p_t, q_t\right\}$$

- The value function is the unique solution to Bellman's equation

$$V(p, q) = \max_{0 \leq q^o \leq \bar{q} - q} \left[ES(p, q) - c^o(p, q^o) + \frac{1}{1+r} EV(p, q)\right]$$

- The optimal purchasing rule is

$$q^o(p, q) = \inf \operatorname*{argmax}_{q^o} \left[ES(p, q) - c^o(p, q^o) + \frac{1}{1+r} EV(p, q)\right]$$

Inventory Investment Model 4

- Hall and Rust (2000) prove that the optimal rule is a generalized (S, s) policy

$$q^o(p,q) = \begin{cases} 0 & \text{if} \quad q \geq s(p) \\ S(p) - q & \text{otherwise} \end{cases} \quad \text{with} \quad S(p) \geq s(p)$$

- Numerical solution of the DP problem
- Parameters calibrated to data

$$r = 2 \times 10{-}4 \quad \alpha_0 = 0.9 \quad \alpha_1 = 1.06$$

$$p^t \quad \text{follows} \quad \log(p_{t+1}) = 0.06 + 0.98 \log(p_t) + \epsilon_t$$

$$\text{with} \quad \epsilon_t \sim iid\, \mathcal{N}(0, 3.94 \times 10^{-4})$$

Inventory Investment Model 5

- Demand

$$q_t^d = \begin{cases} 0 & \text{with prob } 0.5 \\ \log \mathcal{N}(5.5 - 0.7\log(p_t), 1.4) & \text{with prob } 0.5 \end{cases}$$

- Cost

$$c^o(q) = \begin{cases} 0 & \text{if} \quad q = 0 \\ \$75 & \text{if} \quad q > 0 \end{cases}$$

Inventory Investment Model 6

- Solution via PPI
  1. Approximate $V(p,q)$ with a linear combination of basis fcn: Chebyshev polynomials degree 3 in p and q
  2. Discretize the state space of (p,q) pairs: $15 \times 15 = 225$ at Chebyshev zeros
  3. Policy improvement step: Compute the optimal decision rule $q^o(p,q)$ at each (p,q) pair
  4. Policy valuation step: Compute $V_i(p,q) = E\left\{\sum_{j=0}^{\infty}\left(\frac{1}{1+r}\right)^j \pi(p_j, p_j^s, q_j^s, q_i^o(p_j,q_j) + q_j^s)|p,q\right\}$
  5. Regress the $V_i(p,q)$ on the set of p's and q's to compute a new approximation of the value function
  6. Iterate over i on steps 3–5 until the coefficients on the parameterized approximation converge

Inventory Investment Model 7

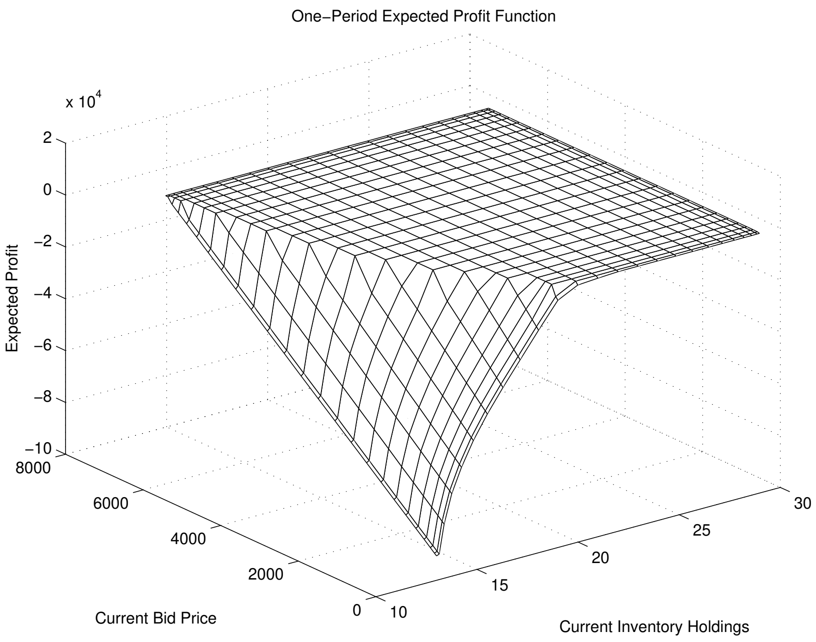

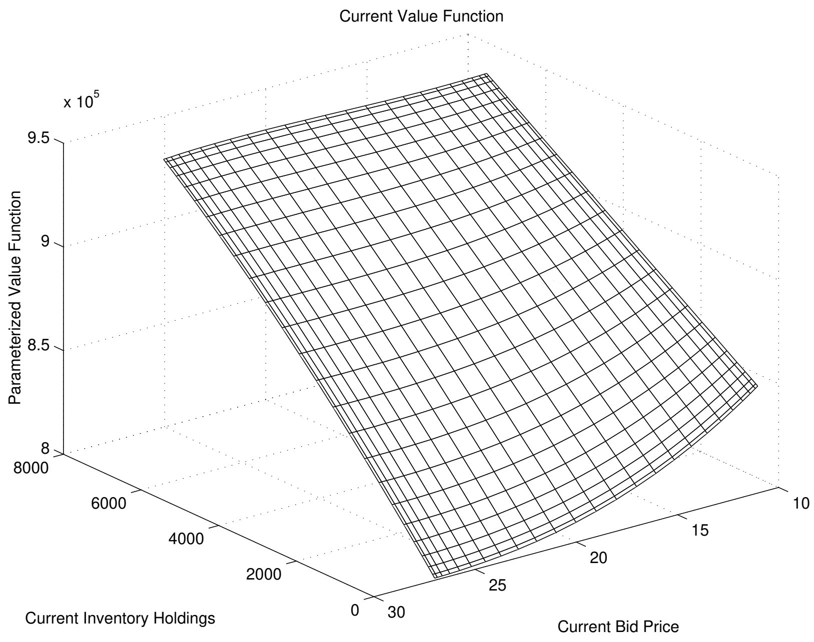

Inventory Investment Model 8

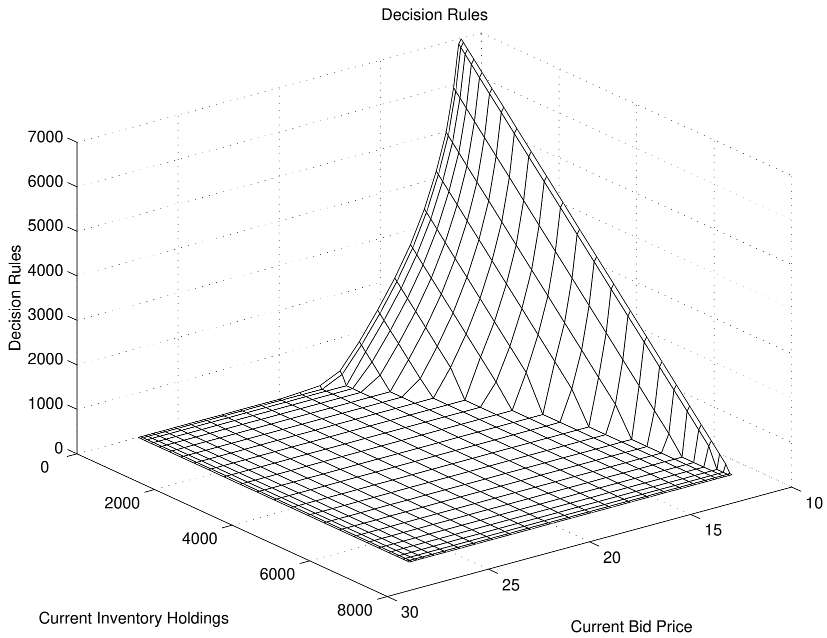

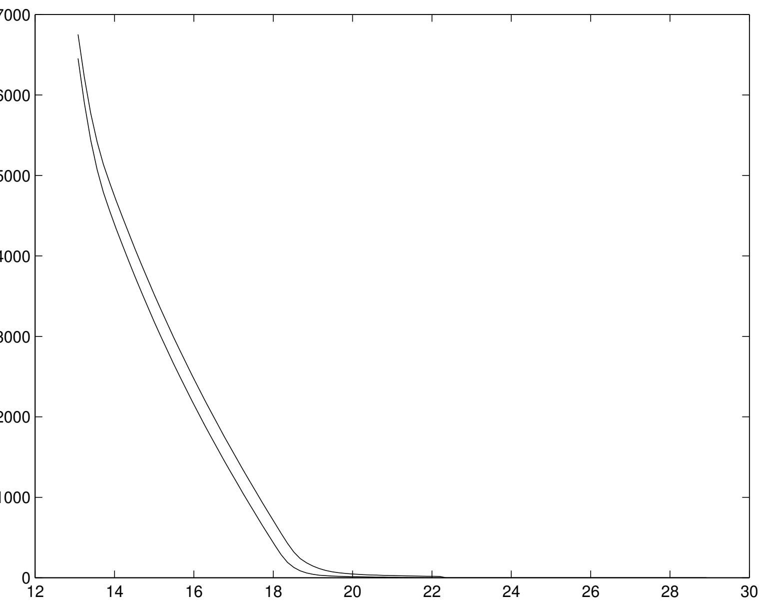

Model of Product Introduction 1

- Firms experiment the exit/stay decision in a market
- Advertising can be altered to learn about the true product quality
- Dynamic setup

$$\max_{x_t} E_\tau \left(\sum_{\tau=t}^{\infty} \beta^{(\tau-t)} \chi_{\tau+1} (Q_\tau (p_\tau - c) - a_\tau - k) \right)$$

  where the state variables $x_t = (\chi_t, b_t, y_t, z_t)$ and the control variables are χ_{t+1}, p_t, a_t
- Bellman's equation

$$V(x_t) = \max \left[0, \sup_{\{\chi_{t+1}=1, p_t, a_t\}} E_t \Big(Q_t (p_t - c) - a_t - k + \beta V(x_{t+1}) \Big) \right]$$

Model of Product Introduction 2

χ_t	indicator in/out of market
b_t	belief of product quality (μ_t, σ_t^2)
y_t	vector of endogenous var's in demand (goodwill, lagged demand, . . .)
z_t	vector of exogenous var's in demand (competitors pricing, advertisement, . . .)
c	unit production cost
k	fixed cost
a_t	adverstisement expenses
p_t	price of product

Model of Product Introduction 3 DPI