Dynamic Programming with State-Dependent Discounting

John Stachurski Junnan Zhang

December 17, 2020

Summary

We study dynamic programming with a discount *function* $\beta(z)$ instead of a constant discount factor β

- Give an "eventual discounting" condition on $\beta(\cdot)$ analogous to $\beta < 1$ and recover classical results
- Show how to test the condition in applications
- Consider extensions related to recursive preferences and unbounded rewards

Outline

Outline

Classical Theory of Dynamic Programming

- A typical dynamic program

$$\max_{\{x_t\}_{t=1}^{\infty}} \mathbb{E} \left\{ \sum_{t=0}^{\infty} \beta^t \, r(x_t, z_t, x_{t+1}) \right\}$$

$$\text{s.t. } \; x_{t+1} \in \Gamma(x_t, z_t)$$

$$z_{t+1} \sim Q(z_t, \cdot)$$

$$x_0, z_0 \text{ given}$$

- We want to find
  1. The value function $v^*(x, z)$
  2. An optimal policy σ^* that maps (x, z) to next-period state

Classical Theory of Dynamic Programming

- The Bellman equation is

$$
\begin{aligned}
v(x,z) &= (Tv)(x,z) \\
&= \sup_{x' \in \Gamma(x,z)} \left\{ r(x,z,x') + \beta \mathbb{E}_z v(x',z') \right\}
\end{aligned}
$$

- A crucial condition is $\beta < 1$ so that the Contraction Mapping Theorem can be applied

Classical Theory of Dynamic Programming

- Under standard assumptions, we have
  1. The value function exists and is the unique fixed point of T
  2. There exists an optimal policy
  3. A policy is optimal if and only if it solves the Bellman equation
  4. Value iteration and policy iteration work
- Under additional assumptions on the primitives
  1. The value function v^* is increasing, strictly concave, and continuously differentiable
  2. The optimal policy σ^* is single-valued and continuous

Classical Theory of Dynamic Programming

Example: Optimal Growth Model

- The agent solves

$$\max_{\{C_t, K_t\}_{t=0}^{\infty}} \mathbb{E} \left\{ \sum_{t=0}^{\infty} \beta^t u(C_t) \right\}$$

$$\text{s.t.} \quad C_t = f(K_t, Z_t) - K_{t+1} \geq 0$$

$$Z_{t+1} \sim Q(Z_t, \cdot)$$

- The Bellman equation is

$$v(k, z) = \sup_{0 \leq k' \leq f(k,z)} \left\{ u(f(k, z) - k') + \beta \mathbb{E}_z v(k', z') \right\}$$

Outline

Why State-Dependent Discount Factors?

State-dependent discounting has been adopted to explain a series empirical phenomena:

Issues	Related work
Equity premium puzzle	Albuquerque et al. (2016) Schorfheide et al. (2018)
Extreme wealth distribution	Krusell and Smith (1998) Hubmer et al. (2020)
Zero lower bound	Christiano et al. (2011) Hills and Nakata (2018)
Macroeconomic volatility	Primiceri et al. (2006) Justiniano and Primiceri (2008)

Dynamic Programming with State-Dependent Discounting

- The agent solves

$$\max_{\{x_t\}_{t=1}^{\infty}} \mathbb{E} \left\{ \sum_{t=0}^{\infty} \left(\prod_{i=0}^{t-1} \beta(z_i) \right) r(x_t, z_t, x_{t+1}) \right\}$$

- The Bellman equation is

$$v(x, z) = \sup_{x' \in \Gamma(x,z)} \left\{ r(x, z, x') + \beta(z) \mathbb{E}_z v(x', z') \right\}$$

Dynamic Programming with State-Dependent Discounting

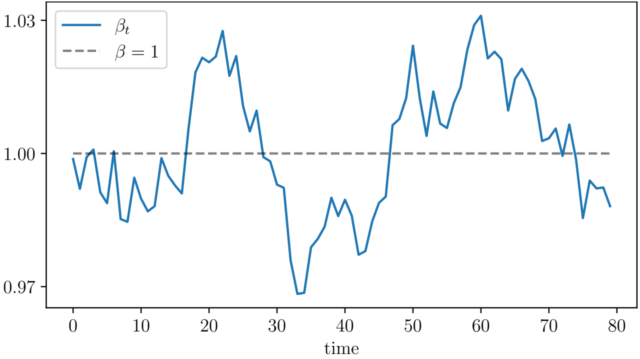

Figure 1: Simulated time path for $\{\beta_t\}$ in Hills et al. (2019)

Dynamic Programming with State-Dependent Discounting

$$(Tv)(x, z) = \sup_{x' \in \Gamma(x,z)} \left\{ r(x, z, x') + \beta(z) \mathbb{E}_z v(x', z') \right\}$$

- Since $\beta(z)$ can exceed 1 for some z, T is no longer a contraction
- What conditions should $\beta(z)$ satisfy to ensure that all classical results still hold?

Outline

Outline

Framework

- A dynamic program consists of
  1. Endogenous state space X
  2. Exogenous state space Z with Markov transition kernel Q
  3. *Feasible correspondence* $\Gamma(x, z) \subset X$
  4. *Continuation aggregator* $H(x, z, x', v)$
- The Bellman equation is

$$v(x, z) = (Tv)(x, z) := \sup_{x' \in \Gamma(x,z)} H(x, z, x', v)$$

Framework

Example: Optimal Growth Model

- $X = \mathbb{R}_+$
- Feasible correspondence $\Gamma(x, z) = [0, f(x, z)]$
- Continuation aggregator

$$H(x, z, x', v) = u(f(x, z) - x') + \beta(z)\mathbb{E}_z v(x', z')$$

Framework

- The set of feasible policies $\Sigma := \{\sigma : \sigma(x, z) \in \Gamma(x, z)\}$
- Any $\sigma \in \Sigma$ corresponds to a *policy operator* T_σ

$$(T_\sigma v)(x, z) := H(x, z, \sigma(x, z), v)$$

- The *σ-value function* is defined by

$$v_\sigma(x, z) := \lim_{n \to \infty} (T_\sigma^n v_0)(x, z)$$

- The value function is defined by

$$v^*(x, z) = \sup_{\sigma \in \Sigma} v_\sigma(x, z)$$

where the maximum is achieved by σ^*

Outline

Assumptions

Definition (Eventual Discounting)

Let $\beta_t = \beta(Z_t)$ where $\{Z_t\}$ is a Markov process on Z with transition kernel Q. We call (β, Q) *eventually discounting* if for some $n \in \mathbb{N}$,

$$r_n^\beta := \sup_{z \in Z} \mathbb{E}_z \prod_{t=0}^{n-1} \beta_t < 1.$$

- Long-run average of the discount process $(r_n^\beta)^{1/n} < 1$
- Examples:
  1. If $\beta_t \equiv b$ for all t, then eventual discounting holds *iff* $b < 1$
  2. If $\{Z_t\}$ is IID, then $r_n^\beta = (\mathbb{E}\beta_t)^n$ and eventual discounting holds *iff* $\mathbb{E}\beta_t < 1$

Assumptions

- One assumption on the aggregator: there exists β such that

$$|H(x, z, x', v) - H(x, z, x', w)| \leq \beta(z)\mathbb{E}_z|v(x', z') - w(x', z')|$$

for all (x, z) and $z' \in \Gamma(x, z)$

- The constant discount case

$$|H(x, z, x', v) - H(x, z, x', w)| \leq b \, \|v - w\|, \quad b < 1$$

Outline

Key Results

Theorem 2.1

If (β, Q) is eventually discounting and certain regularity conditions are satisfied, then all the classical results hold.

The Bellman operator T is *eventually contracting*: there exists an $n \in \mathbb{N}$ such that T^n is a contraction mapping

Necessity of the Eventual Discounting Condition

Necessity

Let $\{\pi_t\}$ be a sequence of rewards such that $0 < a \leq \pi_t \leq b$. If Z is compact and β is continuous, then

$$\mathbb{E} \sum_{t \geq 0} \prod_{i=0}^{t-1} \beta_i \, \pi_t < \infty$$

if and only if eventual discounting holds.

Outline

Eventual Discounting

Connection to Spectral Radii

Let L_β be the *discount operator* defined by

$$(L_\beta h)(z) := \beta(z)\mathbb{E}_z h(z').$$

The spectral radius $r(L_\beta) := \lim_{n\to\infty} \|L_\beta^n\|^{1/n}$.

Proposition 4.1

The spectral radius is $r(L_\beta) = \lim_{n\to\infty}(r_n^\beta)^{1/n}$ and eventual discounting holds ($r_n^\beta < 1$ for some n) *iff* $r(L_\beta) < 1$.

Eventual Discounting

Finite Exogenous State Space

- When Z is finite, $r(L_\beta)$ is easy to calculate
- The transition kernel Q becomes a transition matrix of values Q_{ij}
- Let $\beta_i = \beta(z_i)$. The linear operator L_β becomes a matrix

$$L_\beta := (\beta_i Q_{ij})_{1 \leq i,j \leq N}$$

- $r(L_\beta)$ is defined by the largest absolute value of its eigenvalues

Eventual Discounting

Finite Exogenous State Space

In Christiano et al. (2011), β_t stays at $\beta^h > 1$ with probability p and shifts permanently to $\beta^\ell < 1$ with probability $1 - p$. Thus,

$$L_\beta = \begin{pmatrix} \beta^\ell & 0 \\ (1-p)\beta^h & p\beta^h \end{pmatrix}$$

and the eigenvalues are β^ℓ and $p\beta^h$.

Eventual Discounting

Stationary Spectral Radius

Recall that $r(L_\beta) = \lim_{n\to\infty}(r_n^\beta)^{1/n}$ where

$$r_n^\beta := \sup_{z \in Z} \mathbb{E}_z \prod_{t=0}^{n-1} \beta_t$$

Can we replace r_n^β with $s_n^\beta := \mathbb{E} \prod_{t=0}^{n-1} \beta_t$?

Proposition 4.2

If Z is finite and the exogenous state process $\{Z_t\}$ is irreducible, then $r(L_\beta)$ satisfies the stationary representation

$$r(L_\beta) = s^\beta := \lim_{n\to\infty} (s_n^\beta)^{1/n}$$

Autoregressive Specifications

AR(1) in Levels

In many applications, β_t follows

$$\beta_{t+1} = \rho\beta_t + (1-\rho)\mu + \sigma_\epsilon \epsilon_{t+1}, \quad 0 < \rho < 1$$

We compute $r(L_\beta)$ after following their discretization processes:

Hubmer et al. (2020)	$\rho = 0.992$, $\mu = 0.944$, $\sigma_\epsilon = 0.0006$
	$r(L_\beta) = 0.9469$
Hills et al. (2019)	$\rho = 0.85$, $\mu = 0.99875$, $\sigma_\epsilon = 0.0062$
	$r(L_\beta) = 0.9996$
Nakata (2016)	$\rho = 0.85$, $\mu = 0.995$, $\sigma_\epsilon = 0.00393$
	$r(L_\beta) = 0.9953$

Autoregressive Specifications

AR(1) in Levels

- We discretize the process and plot $r(L_\beta)$ as a function of persistence ρ and volatility σ_ϵ
- Main findings
  1. Larger $\rho, \sigma_\epsilon \implies$ larger $r(L_\beta)$
  2. Larger $\rho, \sigma_\epsilon \implies$ larger effect of increasing σ_ϵ, ρ
- Intuition: $\mathbb{E}\beta_t\beta_{t+1} = \mu^2 + \sigma_\epsilon^2 \frac{\rho}{1-\rho^2}$

Autoregressive Specifications

AR(1) in Levels

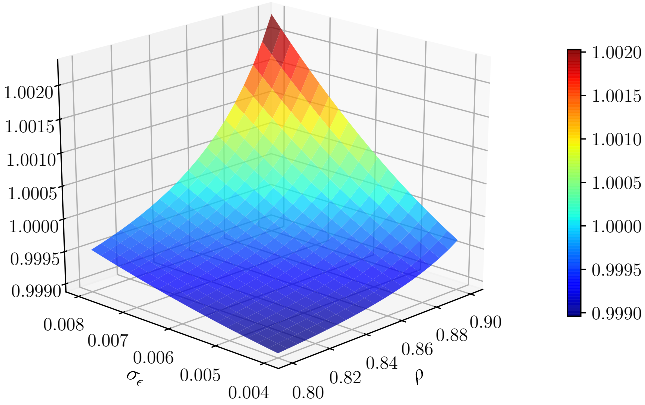

Figure 2: $r(L_\beta)$ as a function of ρ and σ_ϵ; $\mu = 0.944$

Autoregressive Specifications

AR(1) in Logs

- Let $\{\beta_t\}$ be AR(1) in logs

$$\log(\beta_{t+1}) = \rho \log(\beta_t) + (1-\rho)\log b + \sigma_\epsilon \epsilon_{t+1}$$

- $\prod_{t=0}^{n-1} \beta_t = e^{\sum_t \log(\beta_t)}$
- The spectral radius $r(L_\beta)$ after discretization can be closely approximated by

$$s^\beta = \lim_{n\to\infty} \left(\mathbb{E} \prod_{t=0}^{n-1} \beta_t \right)^{1/n} = b \exp\left(\frac{\sigma_\epsilon^2}{2(1-\rho)^2} \right)$$

Outline

Outline

Epstein-Zin Preferences

- The lifetime utility is defined recursively by

$$U_t = \left\{ C_t^{1-1/\psi} + \beta_t \left[\mathbb{E}_t U_{t+1}^{1-\gamma} \right]^{\frac{1-1/\psi}{1-\gamma}} \right\}^{\frac{1}{1-1/\psi}}$$

- $U_t := U(C_t, C_{t+1}, \dots)$
- The agent maximizes lifetime utility by choosing consumption $\{C_t\}$ subject to $X_{t+1} = R_t(X_t - C_t) \geq 0$
- Assume $\gamma > 1$ and $\psi > 1$

Epstein-Zin Preferences

- Let $\beta_t = \beta(Z_t)$ and $R_t = R(Z_t)$
- Define the aggregator by

$$H(x, z, c, v) = \left\{ c^{1-1/\psi} + \beta(z) \left[\mathbb{E}_z v \left(R(z)(x-c), z' \right)^{1-\gamma} \right]^{\frac{1-1/\psi}{1-\gamma}} \right\}^{\frac{1}{1-1/\psi}}$$

s.t. $c \in \Gamma(x, z) = [0, x]$

- The recursive utility of following policy σ is a fixed point of T_σ:

$$v_\sigma(x, z) = (T_\sigma v_\sigma)(x, z) = H(x, z, \sigma(x, z), v_\sigma)$$

Epstein-Zin Preferences

- The eventual discounting condition is

$$\sup_{z \in Z} \mathbb{E}_z \prod_{t=0}^{n-1} \beta_t^{1/(1-1/\psi)} R_t < 1, \quad \text{for some } n \in \mathbb{N}$$

- Equivalently, $r(L_R) < 1$ where

$$(L_R h)(z) := \beta(z)^{1/(1-1/\psi)} R(z) \int h(z') Q(z, dz')$$

Epstein-Zin Preferences

If the eventual discounting condition is satisfied, then

1. $v_\sigma := \lim_{n \to \infty} T_\sigma^n \mathbf{0}$ is well defined and is a fixed point of T_σ
2. $\bar{v} := \lim_{n \to \infty} T^n \mathbf{0}$ is well defined and is the value function: $\bar{v} = v^* := \sup_\sigma v_\sigma$
3. v^* is homogeneous of degree one in x
4. There exists an optimal consumption policy σ^* that is homogeneous of degree one in x
5. The principle of optimality holds

Epstein-Zin Preferences

Eventual Discounting

- Consider again $\{\beta_t\}$ that is AR(1) in logs

$$\log(\beta_{t+1}) = \rho \log(\beta_t) + (1-\rho)\log b + \sigma_\epsilon \epsilon_{t+1}$$

- If $R(z) \equiv R$, $r(L_R)$ can be approximated by

$$s = R \exp\left(\frac{1}{1-1/\psi}\log b + \frac{1}{(1-1/\psi)^2}\frac{\sigma_\epsilon^2}{2(1-\rho)^2}\right)$$

- Basu and Bundick (2017), de Groot et al. (2018)

Epstein-Zin Preferences

Monte Carlo

- Approximate the spectral radius by $\lim_{n \to \infty} s_n^{1/n}$, where

$$s_n = \mathbb{E} \prod_{t=0}^{n-1} \beta_t^{1/(1-1/\psi)} R_t$$

- Monte Carlo: generate m independent simulated paths of $\{\beta_t, R_t\}$ and calculate

$$\hat{s}_n = \frac{1}{m} \sum_{i=1}^{m} \prod_{t=0}^{n-1} \beta_{i,t}^{1/(1-1/\psi)} R_{i,t}$$

- Albuquerque et al. (2016)

Outline

Unbounded Rewards

- Consider aggregators of the form

$$H(x, z, x', v) = u(x, z, x') + \beta(z)\mathbb{E}_z v(x', z')$$

- u is continuous but not necessarily bounded
- Two approaches
  1. Homogeneous functions: Alvarez and Stokey (1998)
  2. Local contractions: Rincón-Zapatero and Rodríguez-Palmero (2003), Matkowski and Nowak (2011)

Unbounded Rewards

Homogeneous Functions

- Two assumptions
  1. $\Gamma(\cdot, z)$ homogeneous of degree one, $u(\cdot, z, \cdot)$ homogeneous of degree $\theta \in (0, 1]$, and other standard conditions
  2. There exists a bounded measurable function α such that $|x'| \leq \alpha(z)|x|$ for all $x' \in \Gamma(x, z)$ and there exists $n \in \mathbb{N}$ such that

$$\sup_{z \in Z} \mathbb{E}_z \prod_{t=0}^{n-1} \beta(Z_t) \alpha^{\theta}(Z_t) < 1$$

- When β and α are both constant, the second condition regresses to $\alpha^{\theta}\beta < 1$ in Alvarez and Stokey (1998)

Unbounded Rewards

Homogeneous Functions

Proposition 5.1

If the assumptions hold, previous results hold on $H(S;\theta)$, the space of bounded continuous functions that are homogeneous of degree θ with norm defined by

$$\|f\| := \sup_{z \in Z} \sup_{x \in X, \, \|x\|=1} |f(x, z)|.$$

Unbounded Rewards

Local Contractions

- Let Z be compact and assume $X = \bigcup_j \operatorname{int} K_j$ where K_j is a sequence of increasing compact sets and
- Let $c > 1$ and $\{m_j\}$ be an unbounded sequence of increasing positive real numbers. Let $C_m(S)$ be the space of all continuous f such that

$$\|f\| := \sum_{j=1}^{\infty} \frac{\|f\|_j}{m_j c^j} < \infty.$$

where $\|f\|_j$ is the sup of f on K_j

Unbounded Rewards

Local Contractions

Proposition 5.2

Under certain regularity conditions, if $r(L_\beta) < 1$ and $\Gamma(x, z) \subset K_j$ for all $x \in K_j$, then previous results hold on $C_m(S)$ for some $\{m_j\}$.

References

Albuquerque, R., M. Eichenbaum, V. X. Luo, and S. Rebelo (2016): "Valuation risk and asset pricing," *The Journal of Finance*, 71, 2861–2904.

Alvarez, F. and N. L. Stokey (1998): "Dynamic programming with homogeneous functions," *Journal of Economic Theory*, 82, 167–189.

Basu, S. and B. Bundick (2017): "Uncertainty shocks in a model of effective demand," *Econometrica*, 85, 937–958.

Christiano, L., M. Eichenbaum, and S. Rebelo (2011): "When is the government spending multiplier large?" *Journal of Political Economy*, 119, 78–121.

de Groot, O., A. W. Richter, and N. A. Throckmorton (2018): "Uncertainty shocks in a model of effective demand: Comment," *Econometrica*, 86, 1513–1526.

Hills, T. S. and T. Nakata (2018): "Fiscal multipliers at the zero lower bound: the role of policy inertia," *Journal of Money, Credit and Banking*, 50, 155–172.

References

Hills, T. S., T. Nakata, and S. Schmidt (2019): "Effective lower bound risk," *European Economic Review*, 120, 103321.

Hubmer, J., P. Krusell, and A. A. Smith (2020): "Sources of US wealth inequality: Past, present, and future," *NBER Macroeconomics Annual 2020, volume 35*.

Justiniano, A. and G. E. Primiceri (2008): "The time-varying volatility of macroeconomic fluctuations," *American Economic Review*, 98, 604–41.

Krusell, P. and A. A. Smith (1998): "Income and wealth heterogeneity in the macroeconomy," *Journal of Political Economy*, 106, 867–896.

Matkowski, J. and A. S. Nowak (2011): "On discounted dynamic programming with unbounded returns," *Economic Theory*, 46, 455–474.

Nakata, T. (2016): "Optimal fiscal and monetary policy with occasionally binding zero bound constraints," *Journal of Economic Dynamics and Control*, 73, 220–240.

Primiceri, G. E., E. Schaumburg, and A. Tambalotti (2006): "Intertemporal disturbances," Tech. rep., National Bureau of Economic Research.

References

Rincón-Zapatero, J. P. and C. Rodríguez-Palmero (2003): "Existence and uniqueness of solutions to the Bellman equation in the unbounded case," *Econometrica*, 71, 1519–1555.

Schorfheide, F., D. Song, and A. Yaron (2018): "Identifying Long-Run Risks: A Bayesian Mixed-Frequency Approach," *Econometrica*, 86, 617–654.