# 6.231 DYNAMIC PROGRAMMING 

## LECTURE 4

## LECTURE OUTLINE

- Review of approximation in value space
- Approximate VI and PI
- Projected Bellman equations
- Matrix form of the projected equation
- Simulation-based implementation
- LSTD and LSPE methods
- Optimistic versions
- Multistep projected Bellman equations
- Bias-variance tradeoff


## REVIEW

## DISCOUNTED MDP

- System: Controlled Markov chain with states $i=1, \ldots, n$, and finite control set $U(i)$ at state $i$ - Transition probabilities: $p_{i j}(u)$

- Cost of a policy $\pi=\left\{\mu_{0}, \mu_{1}, \ldots\right\}$ starting at state $i$ :

$$
J_{\pi}(i)=\lim _{N \rightarrow \infty} E\left\{\sum_{k=0}^{N} \alpha^{k} g\left(i_{k}, \mu_{k}\left(i_{k}\right), i_{k+1}\right) \mid i_{0}=i\right\}
$$

with $\alpha \in[0,1)$

- Shorthand notation for DP mappings

$$
\begin{aligned}
& (T J)(i)=\min _{u \in U(i)} \sum_{j=1}^{n} p_{i j}(u)(g(i, u, j)+\alpha J(j)), \quad i=1, \ldots, n, \\
& \left(T_{\mu} J\right)(i)=\sum_{j=1}^{n} p_{i j}(\mu(i))(g(i, \mu(i), j)+\alpha J(j)), \quad i=1, \ldots, n
\end{aligned}
$$

## "SHORTHAND" THEORY - A SUMMARY

- Bellman's equation: $J^{*}=T J^{*}, J_{\mu}=T_{\mu} J_{\mu}$ or

$$
J^{*}(i)=\min _{u \in U(i)} \sum_{j=1}^{n} p_{i j}(u)\left(g(i, u, j)+\alpha J^{*}(j)\right), \quad \forall i
$$

$$
J_{\mu}(i)=\sum_{j=1}^{n} p_{i j}(\mu(i))\left(g(i, \mu(i), j)+\alpha J_{\mu}(j)\right), \quad \forall i
$$

- Optimality condition:
$\mu:$ optimal $<==>\quad T_{\mu} J^{*}=T J^{*}$
i.e.,
$\mu(i) \in \arg \min _{u \in U(i)} \sum_{j=1}^{n} p_{i j}(u)\left(g(i, u, j)+\alpha J^{*}(j)\right), \quad \forall i$


## THE TWO MAIN ALGORITHMS: VI AND PI

- Value iteration: For any $J \in \Re^{n}$

$$
J^{*}(i)=\lim _{k \rightarrow \infty}\left(T^{k} J\right)(i), \quad \forall i=1, \ldots, n
$$

- Policy iteration: Given $\mu^{k}$
- Policy evaluation: Find $J_{\mu^{k}}$ by solving

$$
\begin{aligned}
J_{\mu^{k}}(i) & =\sum_{j=1}^{n} p_{i j}\left(\mu^{k}(i)\right)\left(g\left(i, \mu^{k}(i), j\right)+\alpha J_{\mu^{k}}(j)\right), \quad i=1, \ldots, n \\
& \text { or } J_{\mu^{k}}=T_{\mu^{k}} J_{\mu^{k}} \\
- & \text { Policy improvement: Let } \mu^{k+1} \text { be such that }
\end{aligned}
$$

$$
\begin{aligned}
& \mu^{k+1}(i) \in \arg \min _{u \in U(i)} \sum_{j=1}^{n} p_{i j}(u)\left(g(i, u, j)+\alpha J_{\mu^{k}}(j)\right), \quad \forall i \\
& \quad \text { or } T_{\mu^{k+1}} J_{\mu^{k}}=T J_{\mu^{k}}
\end{aligned}
$$

- Policy evaluation is equivalent to solving an $n \times n$ linear system of equations
- For large $n$, exact PI is out of the question (even though it terminates finitely)


## APPROXIMATION IN VALUE SPACE

- Approximate $J^{*}$ or $J_{\mu}$ from a parametric class $\tilde{J}(i ; r)$, where $i$ is the current state and $r=\left(r_{1}, \ldots, r_{s}\right)$ is a vector of "tunable" scalars weights
- Think $n$ : HUGE, $s$ : (Relatively) SMALL
- Many types of approximation architectures [i.e., parametric classes $\tilde{J}(i ; r)]$ to select from
- Any $r \in \Re^{s}$ defines a (suboptimal) one-step lookahead policy
$\tilde{\mu}(i)=\arg \min _{u \in U(i)} \sum_{j=1}^{n} p_{i j}(u)(g(i, u, j)+\alpha \tilde{J}(j ; r)), \quad \forall i$
- We want to find a "good" $r$
- We will focus mostly on linear architectures

$$
\tilde{J}(r)=\Phi r
$$

where $\Phi$ is an $n \times s$ matrix whose columns are viewed as basis functions

## LINEAR APPROXIMATION ARCHITECTURES

- We have

$$
\tilde{J}(i ; r)=\phi(i)^{\prime} r, \quad i=1, \ldots, n
$$

where $\phi(i)^{\prime}, i=1, \ldots, n$ is the $i$ th row of $\Phi$, or

$$
\tilde{J}(r)=\Phi r=\sum_{j=1}^{s} \Phi_{j} r_{j}
$$

where $\Phi_{j}$ is the $j$ th column of $\Phi$


- This is approximation on the subspace

$$
S=\left\{\Phi r \mid r \in \Re^{s}\right\}
$$

spanned by the columns of $\Phi$ (basis functions)

- Many examples of feature types: Polynomial approximation, radial basis functions, etc
- Instead of computing $J_{\mu}$ or $J^{*}$, which is hugedimensional, we compute the low-dimensional $r=$ ( $r_{1}, \ldots, r_{s}$ ) using low-dimensional calculations


## APPROXIMATE VALUE ITERATION

## APPROXIMATE (FITTED) VI

- Approximates sequentially $J_{k}(i)=\left(T^{k} J_{0}\right)(i)$, $k=1,2, \ldots$, with $\tilde{J}_{k}\left(i ; r_{k}\right)$
- The starting function $J_{0}$ is given (e.g., $J_{0} \equiv 0$ ) - Approximate (Fitted) Value Iteration: A sequential "fit" to produce $\tilde{J}_{k+1}$ from $\tilde{J}_{k}$, i.e., $\tilde{J}_{k+1} \approx$ $T \tilde{J}_{k}$ or (for a single policy $\mu$ ) $\tilde{J}_{k+1} \approx T_{\mu} \tilde{J}_{k}$


Fitted Value Iteration

- After a large enough number $N$ of steps, $\tilde{J}_{N}\left(i ; r_{N}\right)$ is used as approximation $\tilde{J}(i ; r)$ to $J^{*}(i)$
- Possibly use (approximate) projection $\Pi$ with respect to some projection norm,

$$
\tilde{J}_{k+1} \approx \Pi T \tilde{J}_{k}
$$

## WEIGHTED EUCLIDEAN PROJECTIONS

- Consider a weighted Euclidean norm

$$
\|J\|_{\xi}=\sqrt{\sum_{i=1}^{n} \xi_{i}(J(i))^{2}}
$$

where $\xi=\left(\xi_{1}, \ldots, \xi_{n}\right)$ is a positive distribution ( $\xi_{i}>0$ for all $i$ ).

- Let $\Pi$ denote the projection operation onto

$$
S=\left\{\Phi r \mid r \in \Re^{s}\right\}
$$

with respect to this norm, i.e., for any $J \in \Re^{n}$,

$$
\Pi J=\Phi r^{*}
$$

where

$$
r^{*}=\arg \min _{r \in \Re^{s}}\|\Phi r-J\|_{\xi}^{2}
$$

- Recall that weighted Euclidean projection can be implemented by simulation and least squares, i.e., sampling $J(i)_{k}$ according to $\xi$ and solving

$$
\min _{r \in \Re^{s}} \sum_{t=1}\left(\phi\left(i_{t}\right)^{\prime} r-J\left(i_{t}\right)\right)^{2}
$$

## FITTED VI - NAIVE IMPLEMENTATION

- Select/sample a "small" subset $I_{k}$ of representative states
- For each $i \in I_{k}$, given $\tilde{J}_{k}$, compute

$$
\left(T \tilde{J}_{k}\right)(i)=\min _{u \in U(i)} \sum_{j=1}^{n} p_{i j}(u)\left(g(i, u, j)+\alpha \tilde{J}_{k}(j ; r)\right)
$$

- "Fit" the function $\tilde{J}_{k+1}\left(i ; r_{k+1}\right)$ to the "small" set of values $\left(T \tilde{J}_{k}\right)(i), i \in I_{k}$ (for example use some form of approximate projection)
- Simulation can be used for "model-free" implementation
- Error Bound: If the fit is uniformly accurate within $\delta>0$, i.e.,

$$
\max _{i}\left|\tilde{J}_{k+1}(i)-T \tilde{J}_{k}(i)\right| \leq \delta,
$$

then

$$
\lim \sup _{k \rightarrow \infty} \max _{i=1, \ldots, n}\left(\tilde{J}_{k}\left(i, r_{k}\right)-J^{*}(i)\right) \leq \frac{2 \alpha \delta}{(1-\alpha)^{2}}
$$

- But there is a potential problem!


## AN EXAMPLE OF FAILURE

- Consider two-state discounted MDP with states 1 and 2, and a single policy.
- Deterministic transitions: $1 \rightarrow 2$ and $2 \rightarrow 2$
- Transition costs $\equiv 0$, so $J^{*}(1)=J^{*}(2)=0$.
- Consider (exact) fitted VI scheme that approximates cost functions within $S=\{(r, 2 r) \mid r \in \Re\}$ with a weighted least squares fit; here $\Phi=\binom{1}{2}$
- Given $\tilde{J}_{k}=\left(r_{k}, 2 r_{k}\right)$, we find $\tilde{J}_{k+1}=\left(r_{k+1}, 2 r_{k+1}\right)$, where $\tilde{J}_{k+1}=\Pi_{\xi}\left(T \tilde{J}_{k}\right)$, with weights $\xi=\left(\xi_{1}, \xi_{2}\right)$ : $r_{k+1}=\arg \min _{r}\left[\xi_{1}\left(r-\left(T \tilde{J}_{k}\right)(1)\right)^{2}+\xi_{2}\left(2 r-\left(T \tilde{J}_{k}\right)(2)\right)^{2}\right]$
- With straightforward calculation
$r_{k+1}=\alpha \beta r_{k}, \quad$ where $\beta=2\left(\xi_{1}+2 \xi_{2}\right) /\left(\xi_{1}+4 \xi_{2}\right)>1$
- So if $\alpha>1 / \beta$ (e.g., $\xi_{1}=\xi_{2}=1$ ), the sequence $\left\{r_{k}\right\}$ diverges and so does $\left\{\tilde{J}_{k}\right\}$.
- Difficulty is that $T$ is a contraction, but $\Pi_{\xi} T$ ( $=$ least squares fit composed with $T$ ) is not.


## NORM MISMATCH PROBLEM

- For the method to converge, we need $\Pi_{\xi} T$ to be a contraction; the contraction property of $T$ is not enough


Fitted Value Iteration with Projection

- We need a vector of weights $\xi$ such that $T$ is a contraction with respect to the weighted Euclidean norm $\|\cdot\|_{\xi}$
- Then we can show that $\Pi_{\xi} T$ is a contraction with respect to $\|\cdot\|_{\xi}$
- We will come back to this issue


## APPROXIMATE POLICY ITERATION

## APPROXIMATE PI



Approximate Policy Evaluation

Policy Improvement

- Evaluation of typical policy $\mu$ : Linear cost function approximation $\tilde{J}_{\mu}(r)=\Phi r$, where $\Phi$ is full rank $n \times s$ matrix with columns the basis functions, and $i$ th row denoted $\phi(i)^{\prime}$.
- Policy "improvement" to generate $\bar{\mu}$ :

$$
\bar{\mu}(i)=\arg \min _{u \in U(i)} \sum_{j=1}^{n} p_{i j}(u)\left(g(i, u, j)+\alpha \phi(j)^{\prime} r\right)
$$

- Error Bound (same as approximate VI): If

$$
\max _{i}\left|\tilde{J}_{\mu^{k}}\left(i, r_{k}\right)-J_{\mu^{k}}(i)\right| \leq \delta, \quad k=0,1, \ldots
$$

the sequence $\left\{\mu^{k}\right\}$ satisfies

$$
\limsup _{k \rightarrow \infty} \max _{i}\left(J_{\mu^{k}}(i)-J^{15}(i)\right) \leq \frac{2 \alpha \delta}{(1-\alpha)^{2}}
$$

## POLICY EVALUATION

- Let's consider approximate evaluation of the cost of the current policy by using simulation.
- Direct policy evaluation - Cost samples generated by simulation, and optimization by least squares
- Indirect policy evaluation - solving the projected equation $\Phi r=\Pi T_{\mu}(\Phi r)$ where $\Pi$ is projection w/ respect to a suitable weighted Euclidean norm


Direct Method: Projection of cost vector $J_{\mu}$


Indirect Method: Solving a projected form of Bellman's equation

- Recall that projection can be implemented by simulation and least squares


## PI WITH INDIRECT POLICY EVALUATION



Approximate Policy Evaluation

Policy Improvement

- Given the current policy $\mu$ :
- We solve the projected Bellman's equation

$$
\Phi r=\Pi T_{\mu}(\Phi r)
$$

- We approximate the solution $J_{\mu}$ of Bellman's equation

$$
J=T_{\mu} J
$$

with the projected equation solution $\tilde{J}_{\mu}(r)$

## KEY QUESTIONS AND RESULTS

- Does the projected equation have a solution?
- Under what conditions is the mapping $\Pi T_{\mu}$ a contraction, so $\Pi T_{\mu}$ has unique fixed point?
- Assumption: The Markov chain corresponding to $\mu$ has a single recurrent class and no transient states, i.e., it has steady-state probabilities that are positive

$$
\xi_{j}=\lim _{N \rightarrow \infty} \frac{1}{N} \sum_{k=1}^{N} P\left(i_{k}=j \mid i_{0}=i\right)>0
$$

Note that $\xi_{j}$ is the long-term frequency of state $j$. - Proposition: (Norm Matching Property) Assume that the projection $\Pi$ is with respect to $\|\cdot\|_{\xi}$, where $\xi=\left(\xi_{1}, \ldots, \xi_{n}\right)$ is the steady-state probability vector. Then:
(a) $\Pi T_{\mu}$ is contraction of modulus $\alpha$ with respect to $\|\cdot\|_{\xi}$.
(b) The unique fixed point $\Phi r^{*}$ of $\Pi T_{\mu}$ satisfies

$$
\left\|J_{\mu}-\Phi r^{*}\right\|_{\xi} \leq \frac{1}{\sqrt{18}} \sqrt{1-\alpha^{2}}\left\|J_{\mu}-\Pi J_{\mu}\right\|_{\xi}
$$

## PRELIMINARIES: PROJECTION PROPERTIES

- Important property of the projection $\Pi$ on $S$ with weighted Euclidean norm $\|\cdot\|_{\xi}$. For all $J \in$ $\Re^{n}, \Phi r \in S$, the Pythagorean Theorem holds:

$$
\|J-\Phi r\|_{\xi}^{2}=\|J-\Pi J\|_{\xi}^{2}+\|\Pi J-\Phi r\|_{\xi}^{2}
$$



- The Pythagorean Theorem implies that the projection is nonexpansive, i.e.,

$$
\|\Pi J-\Pi \bar{J}\|_{\xi} \leq\|J-\bar{J}\|_{\xi}, \quad \text { for all } J, \bar{J} \in \Re^{n}
$$

To see this, note that

$$
\begin{aligned}
\|\Pi(J-\bar{J})\|_{\xi}^{2} & \leq\|\Pi(J-\bar{J})\|_{\xi}^{2}+\|(I-\Pi)(J-\bar{J})\|_{\xi}^{2} \\
& =\|J-\bar{J}\|_{\xi}^{2}
\end{aligned}
$$

## PROOF OF CONTRACTION PROPERTY

- Lemma: If $P$ is the transition matrix of $\mu$,

$$
\|P z\|_{\xi} \leq\|z\|_{\xi}, \quad z \in \Re^{n}
$$

Proof: Let $p_{i j}$ be the components of $P$. For all $z \in \Re^{n}$, we have

$$
\begin{aligned}
\|P z\|_{\xi}^{2} & =\sum_{i=1}^{n} \xi_{i}\left(\sum_{j=1}^{n} p_{i j} z_{j}\right)^{2} \leq \sum_{i=1}^{n} \xi_{i} \sum_{j=1}^{n} p_{i j} z_{j}^{2} \\
& =\sum_{j=1}^{n} \sum_{i=1}^{n} \xi_{i} p_{i j} z_{j}^{2}=\sum_{j=1}^{n} \xi_{j} z_{j}^{2}=\|z\|_{\xi}^{2},
\end{aligned}
$$

where the inequality follows from the convexity of the quadratic function, and the next to last equality follows from the defining property $\sum_{i=1}^{n} \xi_{i} p_{i j}=$ $\xi_{j}$ of the steady-state probabilities.

- Using the lemma, the nonexpansiveness of $\Pi$, and the definition $T_{\mu} J=g+\alpha P J$, we have
$\left\|\Pi T_{\mu} J-\Pi T_{\mu} \bar{J}\right\|_{\xi} \leq\left\|T_{\mu} J-T_{\mu} \bar{J}\right\|_{\xi}=\alpha\|P(J-\bar{J})\|_{\xi} \leq \alpha\|J-\bar{J}\|_{\xi}$ for all $J, \bar{J} \in \Re^{n}$. Hence $\Pi T_{\mu}$ is a contraction of modulus $\alpha$.


## PROOF OF ERROR BOUND

- Let $\Phi r^{*}$ be the fixed point of $\Pi T$. We have

$$
\left\|J_{\mu}-\Phi r^{*}\right\|_{\xi} \leq \frac{1}{\sqrt{1-\alpha^{2}}}\left\|J_{\mu}-\Pi J_{\mu}\right\|_{\xi}
$$

Proof: We have

$$
\begin{aligned}
\left\|J_{\mu}-\Phi r^{*}\right\|_{\xi}^{2} & =\left\|J_{\mu}-\Pi J_{\mu}\right\|_{\xi}^{2}+\left\|\Pi J_{\mu}-\Phi r^{*}\right\|_{\xi}^{2} \\
& =\left\|J_{\mu}-\Pi J_{\mu}\right\|_{\xi}^{2}+\left\|\Pi T J_{\mu}-\Pi T\left(\Phi r^{*}\right)\right\|_{\xi}^{2} \\
& \leq\left\|J_{\mu}-\Pi J_{\mu}\right\|_{\xi}^{2}+\alpha^{2}\left\|J_{\mu}-\Phi r^{*}\right\|_{\xi}^{2}
\end{aligned}
$$

where

- The first equality uses the Pythagorean Theorem
- The second equality holds because $J_{\mu}$ is the fixed point of $T$ and $\Phi r^{*}$ is the fixed point of $П Т$
- The inequality uses the contraction property of $\Pi Т$.
Q.E.D.


## SIMULATION-BASED SOLUTION OF PROJECTED EQUATION

## MATRIX FORM OF PROJECTED EQUATION



- The solution $\Phi r^{*}$ satisfies the orthogonality condition: The error

$$
\Phi r^{*}-\left(g+\alpha P \Phi r^{*}\right)
$$

is "orthogonal" to the subspace spanned by the columns of $\Phi$.

- This is written as

$$
\Phi^{\prime} \Xi\left(\Phi r^{*}-\left(g+\alpha P \Phi r^{*}\right)\right)=0,
$$

where $\Xi$ is the diagonal matrix with the steadystate probabilities $\xi_{1}, \ldots, \xi_{n}$ along the diagonal.

- Equivalently, $C r^{*}=d$, where

$$
C=\Phi^{\prime} \Xi(I-\alpha P) \Phi, \quad d=\Phi^{\prime} \Xi g
$$

but computing $C$ and $d$ is HARD (high-dimensional inner products).

## SOLUTION OF PROJECTED EQUATION

- Solve $C r^{*}=d$ by matrix inversion: $r^{*}=C^{-1} d$
- Projected Value Iteration (PVI) method:

$$
\Phi r_{k+1}=\Pi T\left(\Phi r_{k}\right)=\Pi\left(g+\alpha P \Phi r_{k}\right)
$$

Converges to $r^{*}$ because $\Pi T$ is a contraction.


- PVI can be written as:

$$
r_{k+1}=\arg \min _{r \in \Re^{s}}\left\|\Phi r-\left(g+\alpha P \Phi r_{k}\right)\right\|_{\xi}^{2}
$$

By setting to 0 the gradient with respect to $r$,

$$
\Phi^{\prime} \Xi\left(\Phi r_{k+1}-\left(g+\alpha P \Phi r_{k}\right)\right)=0
$$

which yields

$$
r_{k+1}=r_{k}-(\underset{24}{\prime} \Xi \Phi)^{-1}\left(C r_{k}-d\right)
$$

## SIMULATION-BASED IMPLEMENTATIONS

- Key idea: Calculate simulation-based approximations based on $k$ samples

$$
C_{k} \approx C, \quad d_{k} \approx d
$$

- Matrix inversion $r^{*}=C^{-1} d$ is approximated by

$$
\hat{r}_{k}=C_{k}^{-1} d_{k}
$$

This is the LSTD (Least Squares Temporal Differences) Method.

- PVI method $r_{k+1}=r_{k}-\left(\Phi^{\prime} \Xi \Phi\right)^{-1}\left(C r_{k}-d\right)$ is approximated by

$$
r_{k+1}=r_{k}-G_{k}\left(C_{k} r_{k}-d_{k}\right)
$$

where

$$
G_{k} \approx\left(\Phi^{\prime} \Xi \Phi\right)^{-1}
$$

This is the LSPE (Least Squares Policy Evaluation) Method.

- Key fact: $C_{k}, d_{k}$, and $G_{k}$ can be computed with low-dimensional linear algebra (of order $s$; the number of basis functions).


## SIMULATION MECHANICS

- We generate an infinitely long trajectory $\left(i_{0}, i_{1}, \ldots\right)$ of the Markov chain, so states $i$ and transitions $(i, j)$ appear with long-term frequencies $\xi_{i}$ and $p_{i j}$.
- After generating each transition $\left(i_{t}, i_{t+1}\right)$, we compute the row $\phi\left(i_{t}\right)^{\prime}$ of $\Phi$ and the cost component $g\left(i_{t}, i_{t+1}\right)$.
- We form
$d_{k}=\frac{1}{k+1} \sum_{t=0}^{k} \phi\left(i_{t}\right) g\left(i_{t}, i_{t+1}\right) \approx \sum_{i, j} \xi_{i} p_{i j} \phi(i) g(i, j)=\Phi^{\prime} \Xi g=d$
$C_{k}=\frac{1}{k+1} \sum_{t=0}^{k} \phi\left(i_{t}\right)\left(\phi\left(i_{t}\right)-\alpha \phi\left(i_{t+1}\right)\right)^{\prime} \approx \Phi^{\prime} \Xi(I-\alpha P) \Phi=C$
Also in the case of LSPE

$$
G_{k}=\frac{1}{k+1} \sum_{t=0}^{k} \phi\left(i_{t}\right) \phi\left(i_{t}\right)^{\prime} \approx \Phi^{\prime} \Xi \Phi
$$

- Convergence based on law of large numbers.
- $C_{k}, d_{k}$, and $G_{k}$ can be formed incrementally. Also can be written using the formalism of temporal differences (this is just a matter of style)


## OPTIMISTIC VERSIONS

- Instead of calculating nearly exact approximations $C_{k} \approx C$ and $d_{k} \approx d$, we do a less accurate approximation, based on few simulation samples
- Evaluate (coarsely) current policy $\mu$, then do a policy improvement
- This often leads to faster computation (as optimistic methods often do)
- Very complex behavior (see the subsequent discussion on oscillations)
- The matrix inversion/LSTD method has serious problems due to large simulation noise (because of limited sampling) - particularly if the $C$ matrix is ill-conditioned
- LSPE tends to cope better because of its iterative nature (this is true of other iterative methods as well)
- A stepsize $\gamma \in(0,1]$ in LSPE may be useful to damp the effect of simulation noise

$$
r_{k+1}=r_{k}-\gamma G_{k}\left(C_{k} r_{k}-d_{k}\right)
$$

## MULTISTEP PROJECTED EQUATIONS

## MULTISTEP METHODS

- Introduce a multistep version of Bellman's equation $J=T^{(\lambda)} J$, where for $\lambda \in[0,1)$,

$$
T^{(\lambda)}=(1-\lambda) \sum_{\ell=0}^{\infty} \lambda^{\ell} T^{\ell+1}
$$

Geometrically weighted sum of powers of $T$.

- Note that $T^{\ell}$ is a contraction with modulus $\alpha^{\ell}$, with respect to the weighted Euclidean norm $\|\cdot\|_{\xi}$, where $\xi$ is the steady-state probability vector of the Markov chain.
- Hence $T^{(\lambda)}$ is a contraction with modulus

$$
\alpha_{\lambda}=(1-\lambda) \sum_{\ell=0}^{\infty} \alpha^{\ell+1} \lambda^{\ell}=\frac{\alpha(1-\lambda)}{1-\alpha \lambda}
$$

Note that $\alpha_{\lambda} \rightarrow 0$ as $\lambda \rightarrow 1$

- $T^{\ell}$ and $T^{(\lambda)}$ have the same fixed point $J_{\mu}$ and

$$
\left\|J_{\mu}-\Phi r_{\lambda}^{*}\right\|_{\xi} \leq \frac{1}{\sqrt{1-\alpha_{\lambda}^{2}}}\left\|J_{\mu}-\Pi J_{\mu}\right\|_{\xi}
$$

where $\Phi r_{\lambda}^{*}$ is the fixed point of $\Pi T^{(\lambda)}$.

- The fixed point $\Phi r_{\lambda}^{*}$ depends on $\lambda$.


## BIAS-VARIANCE TRADEOFF



- Error bound $\left\|J_{\mu}-\Phi r_{\lambda}^{*}\right\|_{\xi} \leq \frac{1}{\sqrt{1-\alpha_{\lambda}^{2}}}\left\|J_{\mu}-\Pi J_{\mu}\right\|_{\xi}$
- As $\lambda \uparrow 1$, we have $\alpha_{\lambda} \downarrow 0$, so error bound (and the quality of approximation) improves as $\lambda \uparrow 1$. In fact

$$
\lim _{\lambda \uparrow 1} \Phi r_{\lambda}^{*}=\Pi J_{\mu}
$$

- But the simulation noise in approximating

$$
T^{(\lambda)}=(1-\lambda) \sum_{\ell=0}^{\infty} \lambda^{\ell} T^{\ell+1}
$$

## increases

- Choice of $\lambda$ is usually based on trial and error


## MULTISTEP PROJECTED EQ. METHODS

- The projected Bellman equation is

$$
\Phi r=\Pi T^{(\lambda)}(\Phi r)
$$

- In matrix form: $C^{(\lambda)} r=d^{(\lambda)}$, where

$$
C^{(\lambda)}=\Phi^{\prime} \Xi\left(I-\alpha P^{(\lambda)}\right) \Phi, \quad d^{(\lambda)}=\Phi^{\prime} \Xi g^{(\lambda)},
$$

with

$$
P^{(\lambda)}=(1-\lambda) \sum_{\ell=0}^{\infty} \alpha^{\ell} \lambda^{\ell} P^{\ell+1}, \quad g^{(\lambda)}=\sum_{\ell=0}^{\infty} \alpha^{\ell} \lambda^{\ell} P^{\ell} g
$$

- The $\operatorname{LSTD}(\lambda)$ method is

$$
\left(C_{k}^{(\lambda)}\right)^{-1} d_{k}^{(\lambda)},
$$

where $C_{k}^{(\lambda)}$ and $d_{k}^{(\lambda)}$ are simulation-based approximations of $C^{(\lambda)}$ and $d^{(\lambda)}$.

- The $\operatorname{LSPE}(\lambda)$ method is

$$
r_{k+1}=r_{k}-\gamma G_{k}\left(C_{k}^{(\lambda)} r_{k}-d_{k}^{(\lambda)}\right)
$$

where $G_{k}$ is a simulation-based approx. to $\left(\Phi^{\prime} \Xi \Phi\right)^{-1}$

- $\mathrm{TD}(\lambda)$ : An important simpler/slower iteration [similar to $\operatorname{LSPE}(\lambda)$ with $G_{k}=I$ - see the text].


## MORE ON MULTISTEP METHODS

- The simulation process to obtain $C_{k}^{(\lambda)}$ and $d_{k}^{(\lambda)}$ is similar to the case $\lambda=0$ (single simulation trajectory $i_{0}, i_{1}, \ldots$, more complex formulas)

$$
\begin{gathered}
C_{k}^{(\lambda)}=\frac{1}{k+1} \sum_{t=0}^{k} \phi\left(i_{t}\right) \sum_{m=t}^{k} \alpha^{m-t} \lambda^{m-t}\left(\phi\left(i_{m}\right)-\alpha \phi\left(i_{m+1}\right)\right)^{\prime} \\
d_{k}^{(\lambda)}=\frac{1}{k+1} \sum_{t=0}^{k} \phi\left(i_{t}\right) \sum_{m=t}^{k} \alpha^{m-t} \lambda^{m-t} g_{i_{m}}
\end{gathered}
$$

- In the context of approximate policy iteration, we can use optimistic versions (few samples between policy updates).
- Many different versions (see the text).
- Note the $\lambda$-tradeoffs:
- As $\lambda \uparrow 1, C_{k}^{(\lambda)}$ and $d_{k}^{(\lambda)}$ contain more "simulation noise", so more samples are needed for a close approximation of $r_{\lambda}$ (the solution of the projected equation)
- The error bound $\left\|J_{\mu}-\Phi r_{\lambda}\right\|_{\xi}$ becomes smaller
- As $\lambda \uparrow 1, \Pi T^{(\lambda)}$ becomes a contraction for arbitrary projection norm

MIT OpenCourseWare
http://ocw.mit.edu

### 6.231 Dynamic Programming and Stochastic Control

Fall 2015

For information about citing these materials or our Terms of Use, visit: http://ocw.mit.edu/terms.

