Topic 2: Value function iteration in search problems

Jaroslav Borovička
Computational Dynamics (Spring 2023)

New York University

Content

Economic problem

- Worker receives repeated wage offer and decided to accept or continue searching.
- When should the worker accept the offer?
- What does the acceptance decision depend on?

Tools

- Recursification of infinite-horizon problems
- Discretization of continuous-state space problems to a finite grid (global approximation)

Literature

Textbook

- Ljungqvist and Sargent (2020), Chapter 7 (Sections 7.1–7.4)

Search models

- McCall (1970)

QuantEcon

- Quantitative Economics with Python: Topics 33–39 (the baseline search model and various extensions)

Problem setting

An infinite-horizon model of job search (McCall (1970))

- time is discrete and infinite, $t = 0, 1, 2, \dots$
- every period t, an iid wage offer w from distribution $F(w)$ is drawn, with $F(0) = 0$, $F(B) = 1$ for some $B > 0$

A worker decides to accept or reject the offer, $a_t \in \{\text{accept}, \text{reject}\}$

- when accepts, the worker receives income $y_t = w$ forever
- when rejects, the worker receives unemployment benefit $y_t = c$ and moves to next period where a new offer is drawn
- time is discounted at rate β

The worker solves the sequence problem

$$V_0^* = \max_{\{a_t\}_{t=0}^{\infty}} E_0 \left[\sum_{t=0}^{\infty} \beta^t y_t \right] \tag{2.1}$$

where $a_t \in \{\text{accept, reject}\}$ if the worker has not yet accepted any earlier offer, and $a_t \in \{\}$ otherwise.

- V_0^* is the value function, assume V_0^* conditions on the initial offer w_0
- every decision a_t is made conditional on the time-t information set, which contains the history of all offers up to time t, $w^t = (w_0, \ldots, w_t)$
- $E[\cdot]$ is the mathematical expectations operator

$$E[w] = \int_0^B w dF(w) = \int_0^B w f(w)\, dw.$$

Time-t information set based on perfect memory

- the current wage offer w_t,
- the history of previously rejected offers $(w_0, w_1, \ldots, w_{t-1}) \doteq w^{t-1}$,
- time t, and potentially other observed information.

Beliefs about future offers

- offers are iid with distribution $F(w)$
- correct beliefs coincide with the data-generating process $\implies$ rational expectations

Worker's strategy

- a complete description of worker's decisions in every contingency

Recursive formulation utilizes the principle of optimality

$$V_0^* = \max_{\{a_t\}_{t=0}^\infty} \left\{ y_0 + \beta E_0 \left[\sum_{t=1}^\infty \beta^{t-1} y_t \right] \right\} = \max_{a_0} \left\{ y_0 + \beta \max_{\{a_t\}_{t=1}^\infty} E_0 \left[\sum_{t=1}^\infty \beta^{t-1} y_t \right] \right\}$$
$$= \max_{a_0} \left\{ y_0 + \beta E_0 \left[\max_{\{a_t\}_{t=1}^\infty} \left\{ y_1 + \beta E_1 \left[\sum_{t=2}^\infty \beta^{t-2} y_t \right] \right\} \right] \right\} = \max_{a_0} \{ y_0 + \beta E_0 [V_1^*] \} . \quad (2.2)$$

In order to make the problem tractable, we need to find a representation in which V_0^* and V_1^* have the same structure.

- find the state that encodes all relevant information for worker's time-t decision problem
- finding the state is an art (Thomas Sargent)
- here, relevant information can be summarized by the current wage offer

Denote $V(w)$ the value associated with current offer w, before any decision is taken.

Accept a wage offer w at time t

- $y_{t+j} = w, \forall j \geq 0$, which implies value

$$V^a(w) = \sum_{t=0}^{\infty} \beta^t w = \frac{w}{1-\beta}.$$

Reject a wage offer w at time t

- unemployment benefit $y_t = c$, and then draw a next-period w'

$$c + \beta \int_0^B V(w')\, dF(w').$$

The function $V(w)$ is given by

$$V(w) = \max_{\{\text{accept, reject}\}} \left\{ V^a(w),\, c + \beta \int_0^B V(w')\, dF(w') \right\}. \tag{2.3}$$

The functional equation characterizing the decision problem is called the Bellman equation:

$$V(w) = \max_{\{\text{accept, reject}\}} \left\{ V^a(w), c + \beta \int_0^B V(w') \, dF(w') \right\} \tag{2.4}$$

- the solution consists of the function $V(w)$ and the decision rule $a(w) \in \{\text{accept, reject}\}$
- $V(w)$ is the fixed point of the Bellman equation
- this is in contrast to the sequence problem (2.2) where the solution $\{V_t^*\}_{t=0}^{\infty}$ and optimal decisions $\{a_t^*\}_{t=0}^{\infty}$ are stochastic processes that depend on the history of wage draws w^t

The recursive representation (2.3) is the foundation of the **dynamic programming** method

- the validity of this approach is based on the **principle of optimality**
- this principle, due to Richard Bellman (Bellman (1952, 1957)), breaks down the infinite-dimensional problem (2.1) into smaller subproblems:

  An optimal policy has the property that whatever the initial state and initial decision are, the remaining decisions must constitute an optimal policy with regard to the state resulting from the first decision. (Bellman (1957), Chapter III.3)

The equivalence between the value function V_0^* and the fixed point $V(w)$ of the Bellman equation needs to be shown and requires technical assumptions.

Value of accepting the offer is linear and increasing in w

$$V^a(w) = \frac{w}{1-\beta}$$

Value of rejecting the offer is constant

$$Q = c + \beta \int_0^B V(w')\, dF(w')\,. \tag{2.5}$$

Optimal decision must be in the form of a reservation wage $\bar{w}$ such that

- worker accepts if $w > \bar{w}$
- worker rejects if $w < \bar{w}$
- worker is indifferent between accepting and rejecting at $w = \bar{w}$

This implies that the function $V(w)$ is given by

$$V(w) = \begin{cases} c + \beta \int_0^B V(w')\, dF(w') = \dfrac{\bar{w}}{1-\beta} & \text{if } w \leq \bar{w} \\ \dfrac{w}{1-\beta} & \text{if } w \geq \bar{w} \end{cases} \tag{2.6}$$

The only unknown in the characterization is the reservation wage $\bar{w}$.

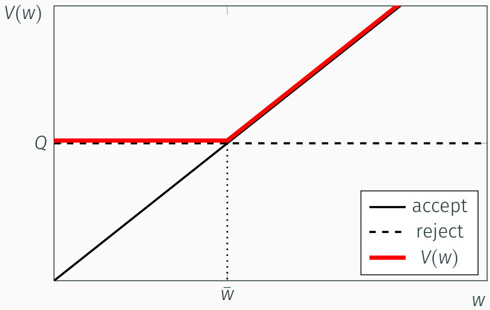

Solution $V(w)$ takes the form (2.6) but there could perhaps be two reservation wages $\bar{w}_1$ and $\bar{w}_2$.

The first line in (2.6) implies

$$\begin{aligned}\frac{\bar{w}}{1-\beta} &= c+\beta\int_0^{\bar{w}}\frac{\bar{w}}{1-\beta}dF(w')+\beta\int_{\bar{w}}^{B}\frac{w'}{1-\beta}dF(w') \\ &= c+\beta\int_0^{B}\frac{\bar{w}}{1-\beta}dF(w')+\frac{\beta}{1-\beta}\int_{\bar{w}}^{B}(w'-\bar{w})\,dF(w')\end{aligned}$$

and hence

$$\bar{w}-c=\frac{\beta}{1-\beta}\int_{\bar{w}}^{B}(w'-\bar{w})\,dF(w')\,. \tag{2.7}$$

- left-hand side is the cost of searching one more time when the current offer is $\bar{w}$
- right-hand side is the net benefit of searching one more time, denote it $h(\bar{w})$
- the only unknown is $\bar{w}$

LHS and RHS in (2.7) are equalized for a unique value of $\bar{w}$.

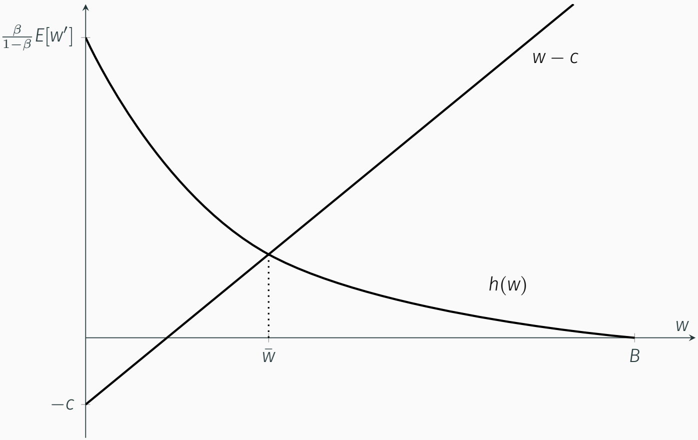

Introspection of equation (2.7) also yields the following associated results:

- $\bar{w} > c$ as long as $F(c) < 1$. Option value of waiting: worker rejects some offers higher than c in order to wait for better future offers.
- $\bar{w}$ does not depend on the shape of $F(w)$ on $[0, \bar{w})$. Reservation wage does not depend on the distribution of offers that get rejected anyway.
- $d\bar{w}/dc > 0$. An increase in c shifts the curve $w - c$ down, decreasing the cost of searching. This shifts $\bar{w}$ to the right. A higher unemployment benefit makes workers pickier.
- $d\bar{w}/d\beta > 0$. An increase in β makes $h(w)$ steeper. Consequently, $\bar{w}$ shifts to the right. Higher patience increases the option value of waiting.

We now modify the problem and assume that the economy has a finite horizon.

- time is discrete, $t = 0, 1, \ldots, T$
- if accepting offer w at time t, worker works at wage w until time T, yielding continuation value

$$V_t^a(w) = \frac{1-\beta^{T-t+1}}{1-\beta} w$$

- the state for the decision problem is (w, t), and value of optimal policy is denoted $V_t(w)$

$$\begin{aligned} V_t(w) &= \max_{\{\text{accept, reject}\}} \left\{ \frac{1-\beta^{T-t+1}}{1-\beta} w, c + \beta \int_0^B V_{t+1}(w')\, dF(w') \right\} \qquad t = 0, \ldots, T \quad (2.8) \\ V_{T+1}(w) &= 0 \end{aligned}$$

Characterizing the finite-horizon problem

Optimal policy is given by a reservation wage $\bar{w}_t$ which now depends explicitly on t.

- This can be shown using steps analogous to the infinite-horizon problem.
- $\bar{w}_T = c$. Any offer better than the unemployment benefit will be accepted in the last period, since there is no option value of waiting.
- As $T \to \infty$, we have $\bar{w}_t \to \bar{w}$ for every fixed t. When the terminal period is far in the future, the reservation wage approaches the time-invariant solution.
- The reservation wage $\bar{w}_t$ is decreasing in time t. Option value of waiting decreases as the worker approaches the terminal period T.

These results can be proven analytically, we will show them numerically.

Numerical implementation

Numerical solution

We implement a range of numerical methods that help us find the model solution.

1. Numerical evaluation of integrals
   - needed to evaluate the expectation operator
2. Solving the reservation wage equation

$$\bar{w} - c = \frac{\beta}{1-\beta} \int_{\bar{w}}^{B} \left(w' - \bar{w}\right) dF\left(w'\right)$$

   - root finding
   - using the contraction mapping property to implement an iteration scheme
3. Value function iteration

$$V_{n+1}\left(w\right) = \max_{\{\text{accept, reject}\}} \left\{ \frac{w}{1-\beta}, c + \beta \int_{0}^{B} V_n\left(w'\right) dF\left(w'\right) \right\}$$

   - backward induction

The model involves a continuous state space. Numerical implementation requires discretization.

- replace $[0, B]$ with a grid of nodes w^i, $i = 0, \ldots I$, such that $0 = w^0 < w^1 < \ldots < w^I = B$
- an equidistant grid that splits $[0, B]$ into I subintervals of length B/I, such that $w^i = iB/I$
- replace the distribution of wage offers $F(w)$ on $[0, B]$ with a discrete distribution $\hat{f}^i \doteq \hat{f}\left(w^i\right)$ on nodes w^i that approximates $F(w)$, for example,

$$\hat{f}^i = \begin{cases} F\left(\frac{1}{2}\left(w^1 + w^0\right)\right) & i = 0 \\ F\left(\frac{1}{2}\left(w^{i+1} + w^i\right)\right) - F\left(\frac{1}{2}\left(w^i + w^{i-1}\right)\right) & 0 < i < I \\ 1 - F\left(\frac{1}{2}\left(w^I + w^{I-1}\right)\right) & i = I \end{cases} \tag{2.9}$$

- concentrates continuous density $f(w)$ into nearest mass points on the grid

Using the approximation (2.9), the expectations operator is approximated as

$$E\left[g\left(w\right)\right] = \int_0^B g\left(w\right) dF\left(w\right) \approx \sum_{i=0}^{I} g\left(w^i\right) \hat{f}^i$$

- this method is called a quadrature rule

There exist efficient quadrature rules that achieve desirable properties with a sparse grid.

Gaussian quadrature designs the choice of nodes $\left\{\tilde{w}^j\right\}_{j=1}^J$ and associated weights $\left\{\tilde{f}^j\right\}_{j=1}^J$ in rule

$$\int_{\underline{w}}^{\overline{w}} g(w)\, dF(w) = \int_{\underline{w}}^{\overline{w}} g(w) f(w)\, dw \approx \sum_{j=0}^{J} g\left(\tilde{w}^j\right) \tilde{f}^j$$

to yield a good approximation for a particular class of functions on $(\underline{w}, \overline{w})$.

- a J-node approximation is constructed to provide an exact formula for the evaluation of the expectation of all polynomial functions up to degree $2J - 1$
- the choice of nodes and weights depends on the particular density $f(w)$
- $f(w)$ does not need to be a density, it can be a general weighting function

More detail in Tauchen and Hussey (1991) and on Wikipedia:

https://en.wikipedia.org/wiki/Gaussian_quadrature.

General approach to Gaussian quadrature

1. Pick a weighting function $f(w)$ and interval $(\underline{w}, \overline{w})$. These define an 'inner product' space: for any two functions $G(w)$, $H(w)$, their inner product is

$$\int_{\underline{w}}^{\overline{w}} G(w) H(w) f(w)\, dw.$$

2. Denote $P_n(w)$ a polynomial of degree n. Construct an orthogonal polynomial basis $P_n(w)$, $n = 0, \ldots, J$.

   Orthogonality: inner product of $P_m(w)$ and $P_n(w)$ is zero for $m \neq n$.

   Basis: Any polynomial of degree up to J is a linear combination of basis polynomials.

   The basis (its monic version) can be built recursively: $P_0(w) = 1$, $P_1(w) = w$, and

$$P_{n+1}(w) = (w - a_{n,n}) P_n(w) - a_{n,n-1} P_{n-1}(w) \tag{2.10}$$

   for appropriate coefficients $a_{n,n}$, $a_{n,n-1}$ (Gram–Schmidt orthogonalization).

Theorem 2.1

Let $P_n(w)$, $n = 0, 1, \ldots, J$ be the orthogonal basis of the space of polynomials of degree up to J on $[\underline{w}, \overline{w}]$ under a weighting function $f(w)$, and $\tilde{w}^j$, $j = 1, \ldots, J$ the roots of $P_J(w)$. Then there exist weights $\tilde{f}^j$, $j = 1, \ldots J$ such that the quadrature rule

$$\int_{\underline{w}}^{\overline{w}} h(w) f(w)\, dw = \sum_{j=1}^{J} h\left(\tilde{w}^j\right) \tilde{f}^j$$

is exact for all polynomials of degree up to $2J - 1$. Moreover, all the nodes $\tilde{w}^j$ lie in the open interval $(\underline{w}, \overline{w})$.

Proof. See notes. ■

3. It can be shown that the nodes $\left\{\tilde{w}^j\right\}_{j=1}^J$ are equal to the eigenvalues of the matrix

$$\hat{\Lambda} = \begin{pmatrix} a_{0,0} & \sqrt{a_{1,0}} & 0 & \dots & \dots & \dots \\ \sqrt{a_{1,0}} & a_{1,1} & \sqrt{a_{2,1}} & 0 & \dots & \dots \\ 0 & \sqrt{a_{2,1}} & a_{2,2} & \sqrt{a_{3,2}} & 0 & \dots \\ 0 & \dots & \dots & \dots & \dots & 0 \\ \dots & \dots & 0 & \sqrt{a_{J-2,J-3}} & a_{J-2,J-2} & \sqrt{a_{J-1,J-2}} \\ \dots & \dots & \dots & 0 & \sqrt{a_{J-1,J-2}} & a_{J-1,J-1} \end{pmatrix}$$

   where $a_{n,n}$ and $a_{n,n-1}$ are the coefficients in the monic rule (2.10).

4. Weights $\left\{\tilde{f}^j\right\}_{j=1}^J$ can be found from the associated eigenvectors ϕ^j:

$$\tilde{f}^j = \frac{\mu_0 \left(\phi_1^j\right)^2}{\|\phi^j\|^2}, \qquad \mu_0 = \int_{\underline{w}}^{\overline{w}} f(w)\, dw$$

   where ϕ_1^j denotes the first element of vector ϕ^j.

See notes for details.

Gauss–Legendre rule designed for weighting function and interval

$$f(w) = 1, \qquad (\underline{w}, \overline{w}) = (-1, 1) .$$

- nodes in the J-node approximation given by roots of Legendre polynomials
- (monic) Legendre polynomials defined recursively: $P_0(w) = 1$, $P_1(w) = w$

$$P_{n+1}(w) = wP_n(w) - \frac{n^2}{4n^2 - 1} P_{n-1}(w)$$

- hence coefficients to be used in matrix $\hat{\Lambda}$ are

$$a_{n,n} = 0 \qquad a_{n,n-1} = \frac{n^2}{4n^2 - 1}$$

The nodes and weights can also be directly translated to an approximation over an arbitrary interval $(\underline{w}, \overline{w})$ and weighting function $f(w) = c$ through a simple linear transformation.

In this case, the new nodes $\check{w}^j$ and weights $\check{f}^j$ are related to $\tilde{w}^j$ and $\tilde{f}^j$ through

$$\begin{aligned} \check{w}^j &= \underline{w} + \frac{\overline{w} - \underline{w}}{2}\left(\tilde{w}^j + 1\right) \\ \check{f}^j &= c\frac{\overline{w} - \underline{w}}{2}\tilde{f}^j \end{aligned}$$

For example, when $f(w)$ is a density on $(\underline{w}, \overline{w})$, then $c = (\overline{w} - \underline{w})^{-1}$

Gauss–Hermite rule

Gauss–Hermite rule designed for weighting function and interval

$$f(w) = \exp\left(-\frac{1}{2}w^2\right), \qquad (\underline{w}, \overline{w}) = (-\infty, \infty).$$

- the weighting function satisfies

$$\int_{-\infty}^{\infty} f(w)\, dw = \sqrt{\pi}.$$

- nodes in the J-node approximation given by roots of Hermite polynomial $P_J(w)$
- (monic) Hermite polynomials defined recursively: $P_0(w) = 1$, $P_1(w) = w$

$$P_{n+1}(w) = wP_n(w) - \frac{n}{2}P_{n-1}(w)$$

- hence coefficients to be used in matrix $\hat{\Lambda}$ are

$$a_{n,n} = 0 \qquad a_{n,n-1} = \frac{n}{2}$$

The Gauss-Hermite rule is frequently used to compute expectations of a normally distributed random variable with mean μ and variance σ^2.

The weighting function is then the density

$$f(w) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{1}{2}\frac{(w-\mu)^2}{\sigma^2}\right).$$

Since the density integrates to one, and the new variable is a linear transformation of the original one, the nodes and weights are transformed as

$$\check{w}^j = \mu + \sqrt{2}\sigma\tilde{w}^j \qquad \check{f}^j = \frac{1}{\sqrt{\pi}}\tilde{f}^j, \qquad j = 1, \ldots, J.$$

We now have two grids:

- $\left\{w^i\right\}$ for the discretization of the value function $V(w)$ (denoted $\hat{V}\left(w^i\right)$),
- $\left\{\tilde{w}^j\right\}$ for the quadrature rule.

These two grids may not align (unless specifically designed), so to evaluate

$$\int_0^B V(w')\, dF(w') \approx \sum_{j=1}^{J} V\left(\tilde{w}^j\right) \tilde{f}^j$$

we must interpolate the values $\hat{V}\left(w^i\right)$ onto $\left\{\tilde{w}^j\right\}$

The simplest form of interpolation is linear interpolation.

- imagine node $\tilde{w}^j$ lies in the interval $\left[w^i, w^{i+1}\right]$ for a particular value of $i = 0, \ldots, I = 1$. Then $V\left(\tilde{w}^j\right)$ can be approximated as

$$V\left(\tilde{w}^j\right) = \hat{V}\left(w^i\right) + \left(\tilde{w}^j - w^i\right) \frac{\hat{V}\left(w^{i+1}\right) - \hat{V}\left(w^i\right)}{w^{i+1} - w^i}.$$

Solving the reservation wage equation

We now return to the reservation wage equation (2.7)

$$\bar{w} - c = \frac{\beta}{1-\beta} \int_{\bar{w}}^{B} (w' - \bar{w}) \, dF(w') \, .$$

Since the LHS is increasing in $\bar{w}$ and RHS decreasing in $\bar{w}$, the root is unique.

- simple numerical problem (e.g., bisection method, Newton–Raphson method)
- requires the numerical evaluation of the integral (with lower bound changing)

Once we determine $\bar{w}$, we also have the whole function $V(w)$

$$V(w) = \begin{cases} \dfrac{\bar{w}}{1-\beta} & \text{if } w \leq \bar{w} \\[2ex] \dfrac{w}{1-\beta} & \text{if } w \geq \bar{w} \end{cases}$$

We can also approach the problem as finding a fixed point of a contraction mapping.

Let us return to equation (2.5) representing the value of rejecting an offer

$$Q = c + \beta \int_0^B V\left(w'\right) dF\left(w'\right) = c + \beta \int_0^B \max_{\{\text{accept, reject}\}} \left\{ \frac{w}{1-\beta}, Q \right\} dF\left(w'\right).$$

This is equivalent to the reservation wage equation since

$$Q = \frac{\bar{w}}{1-\beta}.$$

Define

$$Tq = c + \beta \int_0^B \max_{\{\text{accept, reject}\}} \left\{ \frac{w}{1-\beta}, q \right\} dF\left(w'\right)$$

Then Q is the solution (fixed point) to the equation $q = Tq$.

T is a contraction mapping operator if for two values $Q_1, Q_2 \in \left[0, (1-\beta)^{-1} B\right]$

$$|TQ_1 - TQ_2| \leq \rho \, |Q_1 - Q_2| \text{ for some } \rho \in (0,1)$$

The contraction mapping theorem then implies that

- solution Q that satisfies $Q = TQ$ exists and is unique,
- starting from any value $Q_0 \in \left[0, (1-\beta)^{-1} B\right]$, the fixed point Q can be found as the limit $Q = \lim_{n\to\infty} Q_n$ of successive approximations

$$Q_{n+1} = c + \beta \int_0^B \max_{\{\text{accept, reject}\}} \left\{ \frac{w}{1-\beta}, Q_n \right\} dF\left(w'\right).$$

- consequently, the reservation wage is $\bar{w} = \lim_{n\to\infty} (1-\beta) Q_n$

To show that T is a contraction mapping, notice that for Q_1, $Q_2 \in \left[0, (1-\beta)^{-1} B\right]$

$$
\begin{aligned}
|TQ_2 - TQ_1| &= \beta \left| \int_0^B \left(\max\left\{ \frac{w}{1-\beta}, Q_2 \right\} - \max\left\{ \frac{w}{1-\beta}, Q_1 \right\} \right) dF\left(w'\right) \right| \\
&\leq \beta \int_0^B \left| \max\left\{ \frac{w}{1-\beta}, Q_2 \right\} - \max\left\{ \frac{w}{1-\beta}, Q_1 \right\} \right| dF\left(w'\right) \\
&= \beta \int_0^{\bar{V}_1} |Q_2 - Q_1| \, dF\left(w'\right) + \beta \int_{\bar{V}_1}^{\bar{V}_2} \left| Q_2 - \frac{w}{1-\beta} \right| dF\left(w'\right) \\
&\quad + \beta \int_{\bar{V}_2}^{B} \left| \frac{w}{1-\beta} - \frac{w}{1-\beta} \right| dF\left(w'\right) \\
&\leq \beta \int_0^{\bar{V}_1} |Q_2 - Q_1| \, dF\left(w'\right) + \beta \int_{\bar{V}_1}^{\bar{V}_2} |Q_2 - Q_1| \, dF\left(w'\right) \\
&\leq \beta \, |Q_2 - Q_1|
\end{aligned}
$$

Alternatively, we can iterate on the whole value function using functional equation (2.4).

Define the Bellman operator

$$(Tv)(w) = \max_{\{\text{accept, reject}\}} \left\{ \frac{w}{1-\beta}, c + \beta \int_0^B v(w')\, dF(w') \right\}$$

- the operator maps a function $v : [0, B] \to \mathbb{R}_+$ into a new function Tv.
- the solution to the Bellman equation (2.4) is a function V that satisfies

$$V = TV.$$

If T is a contraction mapping (on a space of functions), then we can construct the recursive scheme

$$V_{n+1}(w) = \max_{\{\text{accept, reject}\}} \left\{ \frac{w}{1-\beta}, c + \beta \int_0^B V_n(w')\, dF(w') \right\} \tag{2.11}$$

- start from an initial guess, e.g., $V_0(w) = 0$.
- given $V_n(w)$, computing $V_{n+1}(w)$ is straightforward
- the contraction mapping property guarantees that $V_n(w) \to V_\infty(w) = V(w)$ as $n \to \infty$.

1. Approximate the distribution $F(w)$ on $[0, B]$ using a discrete distribution $\left\{\hat{f}^i\right\}_{i=0}^I$ on a grid of nodes $\left\{w^i\right\}_{i=0}^I$, with $0 = w^0 < w^1 < \ldots < w^I = B$.
2. Approximate the function $V(w)$ using an $(I+1) \times 1$ vector $\hat{V}$ with elements $\hat{V}^i$.
3. Start with an initial guess $\hat{V}_0$.
4. Replace the functional equation (2.11) with the algebraic system

$$\hat{V}^i_{n+1} = \max_{\{\text{accept, reject}\}} \left\{ \frac{w^i}{1-\beta}, c + \beta \sum_{j=0}^{I} \hat{V}^j_n \hat{f}^j \right\} \qquad i = 0, \ldots, I$$

5. If the numerical scheme is also a contraction mapping, $\hat{V}_n$ converges to a unique fixed point $\hat{V}$.
6. The degree to which $\hat{V}$ approximates well then true function V depends on the quality of the approximation scheme.

Compare backward induction on the finite horizon problem (2.8)

$$
\begin{aligned}
V_t(w) &= \max_{\{\text{accept, reject}\}} \left\{ \frac{1-\beta^{T-t+1}}{1-\beta} w, c + \beta \int_0^B V_{t+1}(w') \, dF(w') \right\} \qquad t = 0, \ldots, T \\
V_{T+1}(w) &= 0
\end{aligned}
$$

with successive approximations implementing the value function iteration

$$
\begin{aligned}
V_{n+1}(w) &= \max_{\{\text{accept, reject}\}} \left\{ \frac{w}{1-\beta}, c + \beta \int_0^B V_n(w') \, dF(w') \right\} \\
V_0(w) &= 0
\end{aligned}
$$

They would be equivalent if accepting in the finite-horizon problem implied employment forever.

Infinite-horizon solution (fixed point) is a limit of the finite-horizon problem.

Summary

Summary

Value function iteration is a widely used method to solve a range of decision problems in economics.

- it is often not the fastest but is robust and relies on very few assumptions
- for example, decision rules do not need to be continuous or differentiable

Discretization of the problem on a grid constitutes a global approximation.

- strives for accuracy of the solution across the whole state space
- curse of dimensionality: computational intensity grows exponentially with the dimension of the state space unless smart methods are used (e.g., endogenous grids), so the method is not suitable for high-dimensional problems
- in such cases, perturbation or projection methods will be more useful

Appendix

Bellman, Richard E. (1952) "On the Theory of Dynamic Programming," *Proceedings of the National Academy of Sciences*, 38 (8), 716–719.

——— (1957) *Dynamic Programming*: Princeton University Press, Princeton, NJ.

Ljungqvist, Lars and Thomas J. Sargent (2020) "Recursive Macroeconomic Theory," Unpublished manuscript, draft of 5th edition.

McCall, John (1970) "Economics of Information and Job Search," *Quarterly Journal of Economics*, 84 (1), 113–126.

Tauchen, George and Robert Hussey (1991) "Quadrature-Based Methods for Obtaining Approximate Solutions to Nonlinear Asset Pricing Models," *Econometrica*, 59 (2), 371–396.