Dynamic Programming

Econ 6105, Fall 2024

Prof. Josh Abel

(Hassler Lecture 5)

An Optimal Stopping Problem

At time $t = 0$, you draw $x_0 \sim U[0, 1]$.

- You can accept x_0, at which point the problem ends and your payoff is x_0
- or you can reject it, and you move to the next period and draw again from $U[0, 1]$
- The problem continues until you accept a draw
- You discount time geometrically (1, β, β^2,...)

The optimal rule is to pick some threshold $\hat{x}$ and accept a draw if and only if $x_t \geq \hat{x}$.

- Why would $\hat{x}$ not depend on t?

What is the optimal threshold, x^*, which maximizes the expected discounted value of your payoff?

A Brute Force Solution: Inputs

Take an arbitrary threshold, $\hat{x}$.
Probability that the problem ends in period:

- 0:
- 1:
- 2:
- ...
- t:
- ...

Average value of x, conditional on $x \geq \hat{x}$, i.e. $E[x|x \geq \hat{x}]$:

A Brute Force Solution: Inputs

Take an arbitrary threshold, $\hat{x}$.
Probability that the problem ends in period:

- 0: $1 - \hat{x}$
- 1:
- 2:
- ...
- t:
- ...

Average value of x, conditional on $x \geq \hat{x}$, i.e. $E[x|x \geq \hat{x}]$:

A Brute Force Solution: Inputs

Take an arbitrary threshold, $\hat{x}$.
Probability that the problem ends in period:

- 0: $1 - \hat{x}$
- 1: $\hat{x} \cdot (1 - \hat{x})$
- 2:
- ...
- t:
- ...

Average value of x, conditional on $x \geq \hat{x}$, i.e. $E[x|x \geq \hat{x}]$:

A Brute Force Solution: Inputs

Take an arbitrary threshold, $\hat{x}$.
Probability that the problem ends in period:

- 0: $1 - \hat{x}$
- 1: $\hat{x} \cdot (1 - \hat{x})$
- 2: $\hat{x}^2 \cdot (1 - \hat{x})$
- ...
- t:
- ...

Average value of x, conditional on $x \geq \hat{x}$, i.e. $E[x|x \geq \hat{x}]$:

A Brute Force Solution: Inputs

Take an arbitrary threshold, $\hat{x}$.
Probability that the problem ends in period:

- 0: $1 - \hat{x}$
- 1: $\hat{x} \cdot (1 - \hat{x})$
- 2: $\hat{x}^2 \cdot (1 - \hat{x})$
- ...
- t: $\hat{x}^t \cdot (1 - \hat{x})$
- ...

Average value of x, conditional on $x \geq \hat{x}$, i.e. $E[x|x \geq \hat{x}]$:

A Brute Force Solution: Inputs

Take an arbitrary threshold, $\hat{x}$.
Probability that the problem ends in period:

- 0: $1 - \hat{x}$
- 1: $\hat{x} \cdot (1 - \hat{x})$
- 2: $\hat{x}^2 \cdot (1 - \hat{x})$
- ...
- t: $\hat{x}^t \cdot (1 - \hat{x})$
- ...

Average value of x, conditional on $x \geq \hat{x}$, i.e. $E[x|x \geq \hat{x}]$:

$$E[x|x \geq \hat{x}] = \frac{\int_{x=\hat{x}}^{1} x \cdot dx}{1 - \hat{x}} = \frac{1}{2} \cdot \frac{1 - \hat{x}^2}{1 - \hat{x}} = \frac{1 + \hat{x}}{2} \tag{1}$$

A Brute Force Solution

$$\begin{aligned} \text{PDV}(\hat{x}) &= \sum_{t=0}^{\infty} \beta^t \cdot \left[\hat{x}^t \cdot (1-\hat{x})\right] \cdot \frac{1+\hat{x}}{2} \\ &= \frac{1-\hat{x}^2}{2} \cdot \sum_{t=0}^{\infty} (\beta \cdot \hat{x})^t \\ &= \frac{1-\hat{x}^2}{2} \cdot \frac{1}{1-\beta \cdot \hat{x}} \end{aligned} \tag{2}$$

$$\frac{d\text{PDV}(\hat{x})}{d\hat{x}} = \frac{1}{2} \cdot \frac{(1-\beta \cdot \hat{x}) \cdot -2 \cdot \hat{x} + \beta \cdot (1-\hat{x}^2)}{(1-\beta \cdot \hat{x})^2} \tag{3}$$

A Brute Force Solution

$$\begin{aligned} \text{PDV}(\hat{x}) &= \sum_{t=0}^{\infty} \beta^t \cdot \left[\hat{x}^t \cdot (1-\hat{x})\right] \cdot \frac{1+\hat{x}}{2} \\ &= \frac{1-\hat{x}^2}{2} \cdot \sum_{t=0}^{\infty} (\beta \cdot \hat{x})^t \\ &= \frac{1-\hat{x}^2}{2} \cdot \frac{1}{1-\beta \cdot \hat{x}} \end{aligned} \tag{2}$$

$$\frac{d\text{PDV}(\hat{x})}{d\hat{x}} = \frac{1}{2} \cdot \frac{(1-\beta \cdot \hat{x}) \cdot -2 \cdot \hat{x} + \beta \cdot (1-\hat{x}^2)}{(1-\beta \cdot \hat{x})^2} \tag{3}$$

Setting $\frac{d\text{PDV}(\hat{x})}{d\hat{x}} = 0$ yields

$$\beta \cdot \hat{x}^2 - 2 \cdot \hat{x} + \beta = 0 \tag{4}$$

Equation 4 is solved by $\hat{x} = \frac{1 \pm \sqrt{1-\beta^2}}{\beta}$; we take the minus option: $x^* = \frac{1-\sqrt{1-\beta^2}}{\beta}$

Solution Visualized

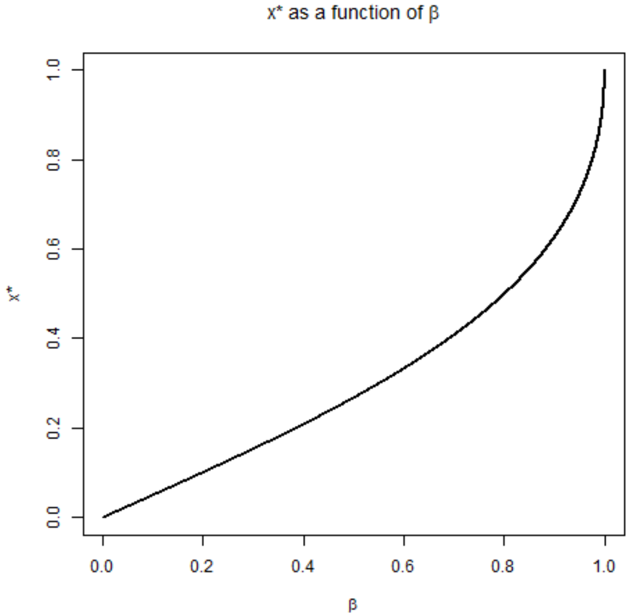

The Dynamic Programming Approach

Define $V(x_0)$ to be the value of drawing x_0, assuming you make the right decision in period 0 and all subsequent periods.

Then:

$$V(x_0) = \max\{x_0, \beta \cdot E[V(x)]\} \tag{5}$$

- If x_0 is larger than expected value of waiting, you accept it
- If the expected value of waiting is higher, your reject x_0 and continue

No longer looking at an infinite sum: all periods from $1 - \infty$ are baked into $E[V(x)]$

- Like how a stock's value can be broken down into 2 parts:
  1. Dividend
  2. Expected future price

Solution via Dynamic Programming

Define x^* to be the optimal threshold. Explain why we **cannot** have:

- $x^* < \beta \cdot E[V(x)]$
- $x^* > \beta \cdot E[V(x)]$

Solution via Dynamic Programming

Define x^* to be the optimal threshold. Explain why we **cannot** have:

- $x^* < \beta \cdot E[V(x)]$
  - Then should be rejecting more
- $x^* > \beta \cdot E[V(x)]$
  - Then should be rejecting less

So:

$$x^* = \beta \cdot E[V(x)]$$

Solution via Dynamic Programming

Define x^* to be the optimal threshold. Explain why we **cannot** have:

- $x^* < \beta \cdot E[V(x)]$
- $x^* > \beta \cdot E[V(x)]$

So:

$$x^* = \beta \cdot E[V(x)]$$

$$\begin{aligned} E[V(x)] &= {x^*}^2 + \int_{x^*}^{1} x \cdot dx \\ &= \frac{1}{2}(1 + {x^*}^2) \end{aligned}$$

These equations imply $0 = \beta \cdot {x^*}^2 - 2 \cdot x^* + \beta$.
This is the same quadratic we found previously, so we have the same solution!

Solution Visualized

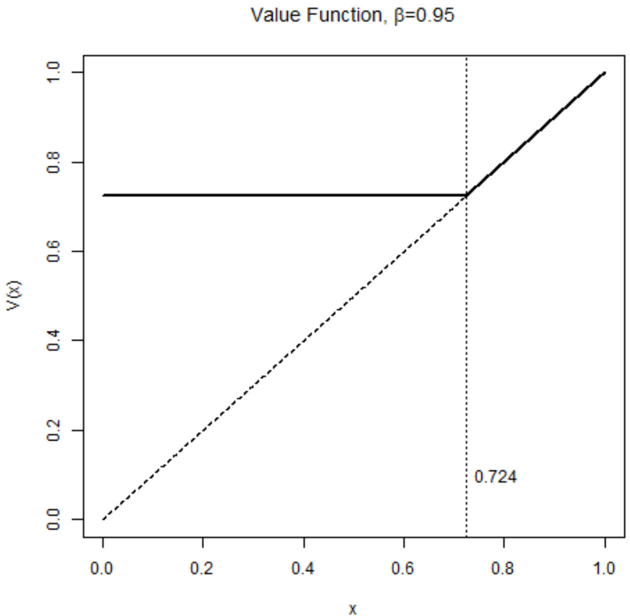

Difficult Problems

Simple as this example was, it had two features that can make problems very tricky:

1. Dynamic: decisions in one period affect future periods
2. Stochastic: there is uncertainty about how future periods will unfold

When combined, these two features can make problems intractable.

- This is particularly true when combined with a 3^{rd} feature: persistent shocks
  - I.e. if the draws did not automatically revert to the $U[0, 1]$ distribution every period

Advantages of Dynamic Programming

The brute force approach required us to find:

1. The full distribution of T, the time of accepting the offer
2. The conditional expectation of the offer, conditional on being accepted

In this simple example, we were able to solve for those things, but in more complex scenarios, they can be intractable

Dynamic programming takes a problem with many periods and it collapses it to a 2-period problem

- This often unlocks ways to solve it that are hidden when viewing the full complexity of the problem

Sequence Problem

Consider a "state variable," x, like draw from $U[0,1]$, and a "policy variable," θ, like whether to accept the draw. A common problem in economics is to maximize the following value function:

$$V(x_0) = E\left[\max_{\{\theta_t\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t \cdot u(x_t, \theta_t) \middle| x_0, \theta_0\right], \tag{6}$$

where the distribution of x_{t+1} depends on x_t and θ_t.
Given some current value of the state variable, x_0, the "value", $V(x_0)$, is:

- The expected (average) value of...
- the sum across all future periods of...
- each period's discounted flow of utility...
- assuming the agent implements the optimal policy in each period (in any state of the world).

Recursive Formulation

$$V(x_0) = \max_{\theta_0} \left\{ \underbrace{u(x_0, \theta_0)}_{\text{Flow utility}} + \underbrace{\beta \cdot E\left[V(x_1) | x_0, \theta_0 \right]}_{\text{Continuation value}} \right\}, \tag{7}$$

where the distribution of x_{t+1} depends on x_t and θ_t.

Given some current value of the state variable, x_0, the "value", $V(x_0)$, is:

- The flow of utility in period 0 plus...
- the expected (average) value of...
- the value in period 1...
- assuming the agent implements the optimal policy in period 0
- (and implicitly assuming the agent continues making optimal decisions in all future periods, in all states of the world).

This is often referred to as the "Bellman Equation."

Relationship Between Sequence Problem and Bellman Equation

- A function $V(x)$ that solves the sequence problem in Equation 6 also satisfies the Bellman Equation in 7.
  - V solves sequence $\Rightarrow$ V solves Bellman
- A function that solves the Bellman Equation in 7 also solves the sequence problem in Equation 6 if the following "Transversality Condition" holds:

$$\lim_{s \to \infty} \beta^s \cdot V(x_s) = 0 \tag{8}$$

  - V solves Bellman **and** Equation 8 holds $\Rightarrow$ V solves sequence

Transversality Condition

$$
\begin{aligned}
V(x_0) &= \max_{\theta_0} \left\{ u(x_0, \theta_0) + \beta \cdot E\Big[V(x_1) | x_0, \theta_0 \Big] \right\} \\
&= \max_{\theta_0} \left\{ u(x_0, \theta_0) + E\Big[\max_{\theta_1} \Big\{ \beta \cdot u(x_1, \theta_1) + \beta^2 \cdot E\Big[V(x_2) | x_1, \theta_1 \Big] \Big\} | x_0, \theta_0 \Big] \right\} \\
&= \max_{\theta_0} \Big\{ u(x_0, \theta_0) + E\Big[\max_{\theta_1} \Big\{ \beta \cdot u(x_1, \theta_1) + \ldots \\
&+ \max_{\theta_T} \Big\{ \beta^T \cdot u(x_T, \theta_T) + \beta^{T+1} \cdot E\Big[V(x_{T+1}) | x_T, \theta_t \Big] \Big\} | \ldots | x_0, \theta_0 \Big] \Big\}
\end{aligned} \tag{9}
$$

As $T \to \infty$, this replicates the sequence problem exactly, so long as that very last term ($\beta^s \cdot E[V(x_s)]$) converges to 0

- That's the Transversality Condition in Equation 8

A Harder Problem

Same as before, but now if you reject, there is a δ probability that you get no draws for the next n periods.

Bellman Equation:

$$V(x_t) = \max\Big\{x_t, \beta \cdot \big((1-\delta) \cdot E[V(x)] + \delta \cdot \beta^n \cdot E[V(x)]\big)\Big\} \tag{10}$$

Optimality condition:

$$x^* = \beta \cdot \big((1-\delta) \cdot E[V(x)] + \delta \cdot \beta^n \cdot E[V(x)]\big) \tag{11}$$

Just as before:

$$E[V(x)] = {x^*}^2 + \int_{x^*}^{1} x \cdot dx = \frac{1}{2} \cdot (1 + {x^*}^2) \tag{12}$$

A Harder Problem (2)

The previous 2 equations imply a quadratic:

$$0 = \beta \cdot {x^*}^2 - \frac{2}{1 - \delta + \delta \cdot \beta^n} \cdot x^* + \beta \qquad (13)$$

This is the same quadratic as before if $\delta = 0$ or $n = 0$.
For $\beta = 0.95, \delta = 0.3$, $n = 10$, we have $x^* = 0.539$.

A Harder Problem (3)

Using the brute force method would be very difficult, as the distribution of stopping times would be complex, as you need to keep track of the likelihood that you get a draw in each period.
If $n = 2$, the likelihood of the game ending in each period is:

- 0: $1 - \hat{x}$
- 1: $(1 - \hat{x}) \cdot (1 - \delta) \cdot \hat{x}$
- 2: $(1 - \hat{x}) \cdot \left((1 - \delta) \cdot \hat{x}\right)^2$
- 3: $(1 - \hat{x}) \cdot \left(\left((1 - \delta) \cdot \hat{x}\right)^3 + \delta \cdot \hat{x}\right)$
- 4: $(1 - \hat{x}) \cdot \left(\left((1 - \delta) \cdot \hat{x}\right)^4 + 2 \cdot \delta \cdot (1 - \delta) \cdot \hat{x}^2\right)$
- ...it's only going to get nastier

Becomes a challenging combinatorics problem, and the weighted sum is likely to be intractable.

Functional Operators

A "functional operator" is like a meta-function: a function for functions.

- It takes in function(s) and spits out – in our case – a (perhaps) changed function

Examples:

1. $T(f)(x) = 2 \cdot f(x)$
    - $T(x^2)(x) = 2 \cdot x^2$
    - $T(e^x)(x) = 2 \cdot e^x$
2. $T(f)(x) = \max\{0, f(x)\}$
    - $T(2 \cdot x)(x) = \begin{cases} 0 & \text{if } x \leq 0 \\ 2 \cdot x & \text{if } x > 0 \end{cases}$
    - $T(9 - x^2)(x) = \begin{cases} 0 & \text{if } |x| \geq 3 \\ 9 - x^2 & \text{if } |x| < 3 \end{cases}$

Fixed Points

A "fixed point" of a function, $f(x)$ is any point $x = a$ such that $f(a) = a$.

- The output of the function is equal to its input.

Examples:

- $f(x) = 1 + \frac{x}{2}$
  - Fixed point(s):
- $f(x) = x^3$
  - Fixed point(s):

Fixed Points

A "fixed point" of a function, $f(x)$ is any point $x = a$ such that $f(a) = a$.

- The output of the function is equal to its input.

Examples:

- $f(x) = 1 + \frac{x}{2}$
  - Fixed point(s): x=2
- $f(x) = x^3$
  - Fixed point(s): x=-1, 0, or 1

A fixed point of a functional operator is any function, f, such that $T(f)(x) = f(x)$

- This means they must be equal for every value of x in the domain!

The Bellman Operator

Define the function $W(x_0)$ to allow for maximization errors:

$$W(x_0) = u(x_0, \theta^W(x_0)) + \beta \cdot E\left[W(x_1)|x_0, \theta^W(x_0)\right] \tag{14}$$

This is the same as the Bellman Equation in 7 except with no assumption of optimal decisions.
Define B to be the following "Bellman Operator:"

$$B(W)(x) = \max_{\theta^W(x_0)} W(x_0) \tag{15}$$

The value function, $V(x_0)$, is a fixed point of the Bellman Operator:

$$V(x_0) = B(V)(x_0) \tag{16}$$

Contraction Mapping

A functional operator, T, is a “contraction mapping” on some space of functions, F, if:

$$\max_x |T(f)(x) - T(g)(x)| < \max_x |f(x) - g(x)| \tag{17}$$

for all $f, g \in F$.

- I've left out some technical stuff
- A contraction mapping is broader than what I have here, but this will work for us

Main idea: a contraction mapping moves two functions closer together, where closeness is defined by the maximum distance between them.

Contraction Mapping Examples

When x is bounded, $T(f)(x) = f(x)/2$ is a contraction mapping.

- $\max_x |T(f)(x) - T(g)(x)| = \frac{1}{2} \cdot \max_x |f(x) - g(x)| < \max_x |f(x) - g(x)|$

On $x \in [0, 2]$, $T(f)(x) = f(x)^2$ is not a contraction mapping.

- Counterexample: $f(x) = x$, $g(x) = 0$
  - $\max_x |T(f)(x) - T(g)(x)| = 4 > 2 = \max_x |f(x) - g(x)|$

Contraction Mapping Theorem

If $T : F \to F$ is a contraction mapping, then T has exactly one fixed point, $f^* \in F$. Furthermore, for any $f_0 \in F$:

$$\lim_{n \to \infty} T^n(f_0)(x) = f^*(x). \tag{18}$$

So if our Bellman Operator is a contraction mapping, then:

1. There is a function V that solves it.
2. There is only 1 function V that solves it.
3. We can find that function V by starting with any (!) function and "iterating the Bellman Operator."
    - $B^n(W_0)(x)$ converges to $V(x)$ no matter what W_0 you chose (!)

This provides a blueprint for numerically solving a Bellman Equation

- As long as the Transversality Condition holds, we can therefore solve the sequence problem (!)

Iterating the Bellman Operator: Another Approach

Go back to our simple example and define the Bellman Operator as:

$$B(W)(x_0) = \max\{x_0, \beta \cdot E[W(x)]\} \tag{19}$$

Let $W_0(x_0) = 1$ (a terrible guess)

1. $B(W_0)(x_0) =$
2. $B^2(W_0)(x_0) =$

Iterating the Bellman Operator: Another Approach

Go back to our simple example and define the Bellman Operator as:

$$B(W)(x_0) = \max\{x_0, \beta \cdot E[W(x)]\} \tag{19}$$

Let $W_0(x_0) = 1$ (a terrible guess)

1. $B(W_0)(x_0) = \begin{cases} \beta \text{ if } x < \beta \\ x \text{ if } x \geq \beta \end{cases}$
2. $B^2(W_0)(x_0) =$

Iterating the Bellman Operator: Another Approach

Go back to our simple example and define the Bellman Operator as:

$$B(W)(x_0) = \max\{x_0, \beta \cdot E[W(x)]\} \tag{19}$$

Let $W_0(x_0) = 1$ (a terrible guess)

1. $B(W_0)(x_0) = \begin{cases} \beta \text{ if } x < \beta \\ x \text{ if } x \geq \beta \end{cases}$
2. $B^2(W_0)(x_0) = \begin{cases} \beta^2 + \frac{\beta}{2} \cdot (1-\beta)^2 \text{ if } x < \beta^2 + \frac{\beta}{2} \cdot (1-\beta)^2 \\ x \text{ if } x \geq \beta^2 + \frac{\beta}{2} \cdot (1-\beta)^2 \end{cases}$
3. ...

Iterating the Bellman Operator: Visualized

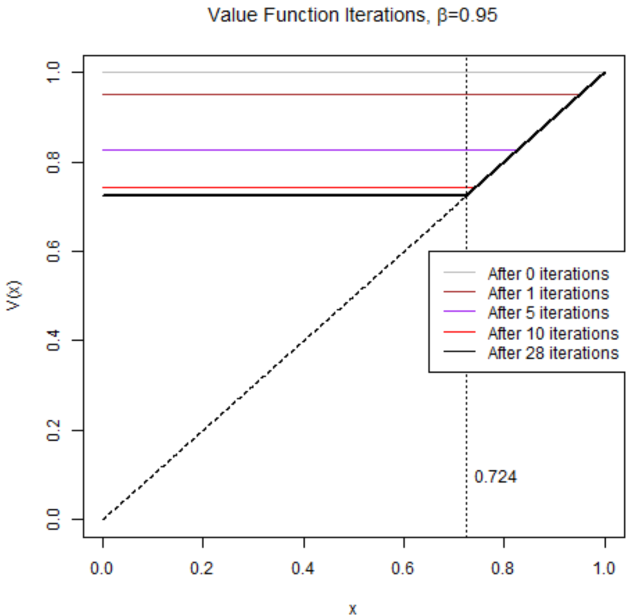

Blackwell's Theorem

Let F be the set of all bounded functions $f : S \subseteq R^n \rightarrow R^1$. A functional operator, $T : S \rightarrow S$ is a contraction mapping if it has the following 2 properties:

1. Monotonicity: $\forall f, g \in F$ such that $f(x) \leq g(x) \; \forall x$, $T(f)(x) \leq T(g)(x) \; \forall x$.
   - If f is "less" than g in the function-y sense, $T(f)$ is also less than $T(g)$.
2. Discounting: $\exists \delta \in [0, 1]$ such that, $\forall x \in S$, $a > 0$, $f \in F$:

$$T(f + a)(x) \leq T(f)(x) + \delta \cdot a \tag{20}$$

   - T does not "expand" additive constants.

Note that Blackwell's conditions are sufficient, but not necessary, for T to be a contraction mapping.

Blackwell's Conditions, Example

Consider the Bellman Operator from our simple example:

$$B(W)(x_0) = \max\{x_0, \beta \cdot E[W(x)]\}$$

1. Monotonicity
   - Assume $f(x) \leq g(x)\ \forall x \in [0,1]$
   - $B(f)(x_0) = \max\{x_0, \beta \cdot E[f(x)]\} \leq \max\{x_0, \beta \cdot E[g(x)]\} = B(g)(x_0) \checkmark$
2. Discounting
   - $B(f+a) = \max\{x_0, \beta \cdot (E[f(x)+a])\} \leq \max\{x_0, \beta \cdot E[f(x)]\} + \beta \cdot a = B(f)(x_0) + \beta \cdot a$
   - Since $\beta \in [0,1]$, discounting is satisfied by $\delta = \beta \checkmark$

So we can feel confident that our numerical approach was appropriate.

- After all, we've proven B is a contraction mapping, so our approach will find the unique fixed point

Mortgage Refinancing

A risk-neutral borrower with discount factor β has some interest rate on her mortgage, r_0.

- To further simplify, assume the mortgage has an infinite term and does not amortize.

At any time, she can pay some cost, C, to switch her interest rate for the market rate, r_t. The market interest rate follows a Random Walk, with shocks drawn from $N(0, \sigma^2)$.
Define $y_t \equiv r_t - r_0$. How far must y_t fall before the borrower should pay the cost to refinance?

Mortgage Refinancing (2)

"Real options" problems like this are very challenging:

- Need to think about every potential path of interest rates, which determines the key state variable, y...
- ...taking into account that the borrower's decision rule will impact how y evolves

Finding the optimal threshold y^* is totally infeasible with anything like the "brute force" approach:

- Good luck finding a distribution for all the times refinances will occur along with the savings they will generate!

But it's very solvable with dynamic programming!

- In continuous time, there is an analytical solution.
- In discrete time, we can numerically iterate the following Bellman Operator:

$$B(W)(y_0) = \min\left\{\frac{y_0}{1-\beta} + C + E[W(y_1)|\bar{y}_1 = 0], E[W(y_1)|\bar{y}_1 = y_0]\right\}$$

Iterating to the Mortgage Refinancing Solution

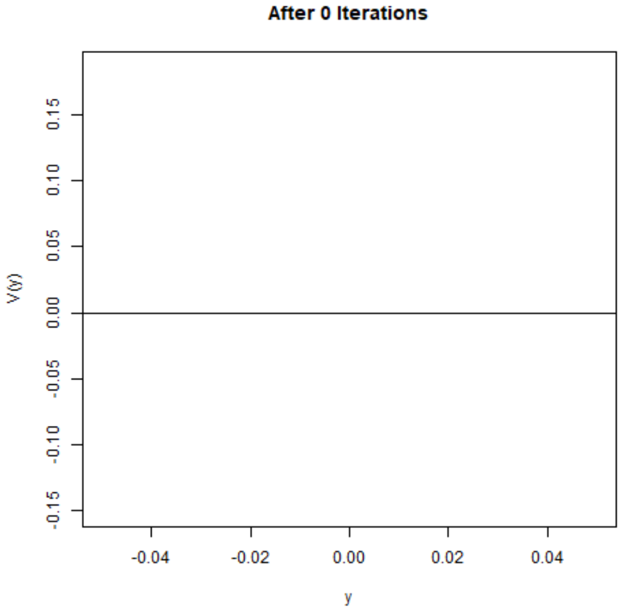

Iterating to the Mortgage Refinancing Solution

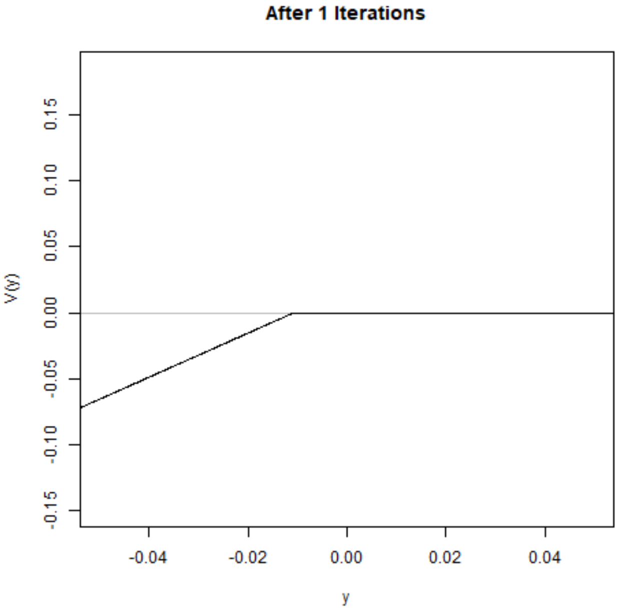

Iterating to the Mortgage Refinancing Solution

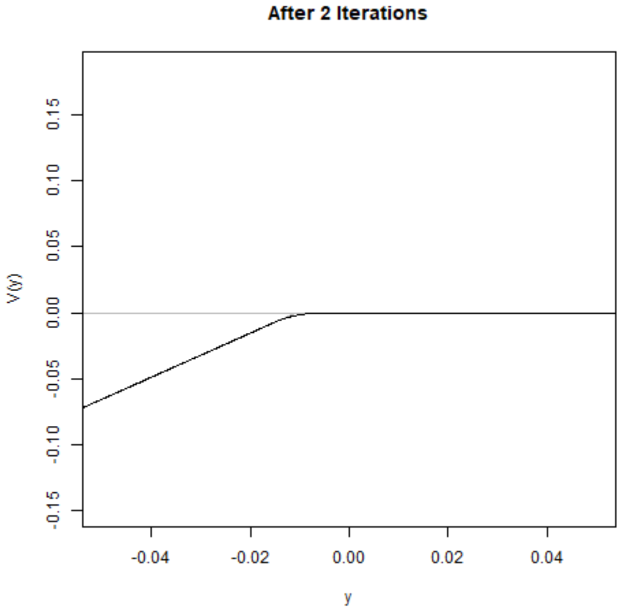

Iterating to the Mortgage Refinancing Solution

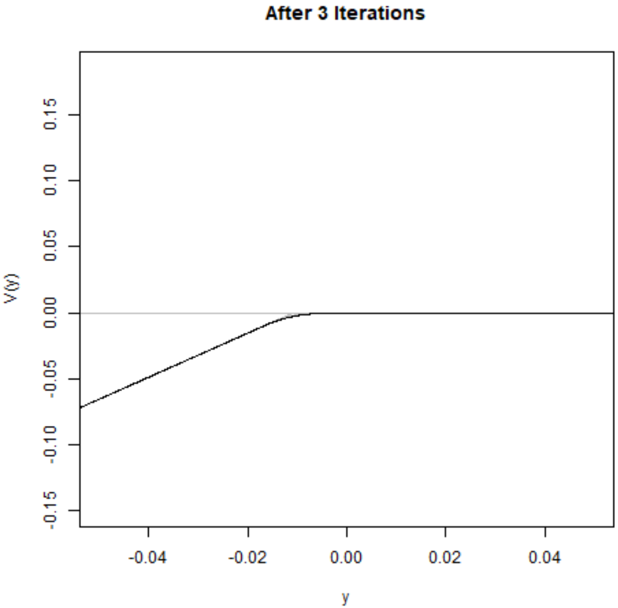

Iterating to the Mortgage Refinancing Solution

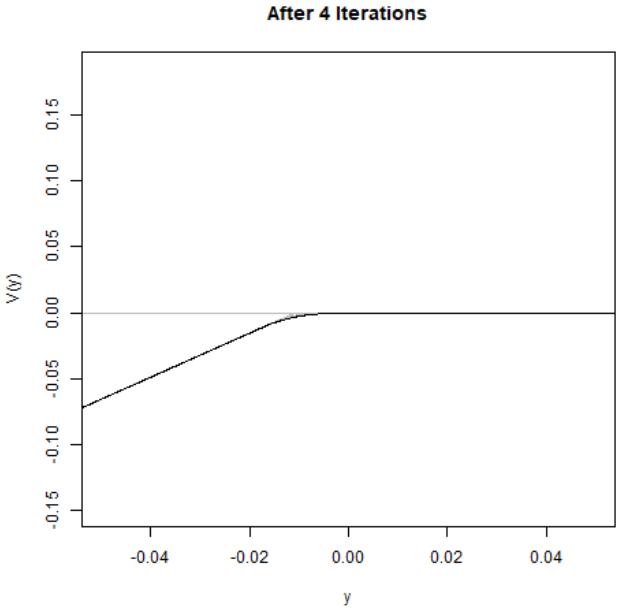

Iterating to the Mortgage Refinancing Solution

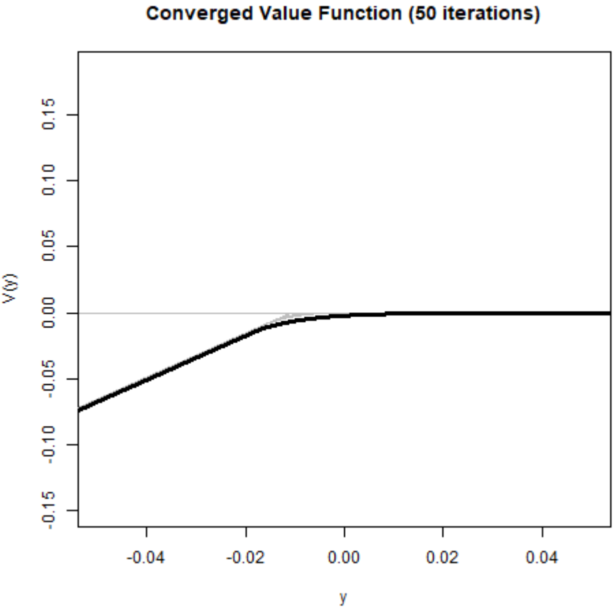

Mortgage Refinancing Solution, Visualized

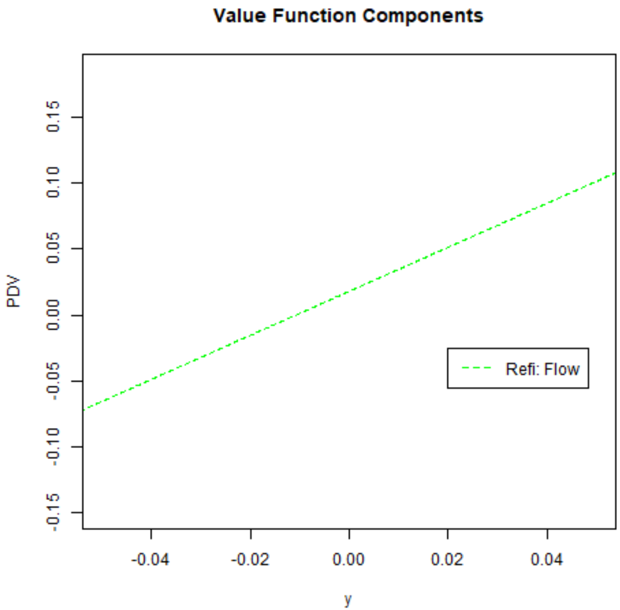

Mortgage Refinancing Solution, Visualized

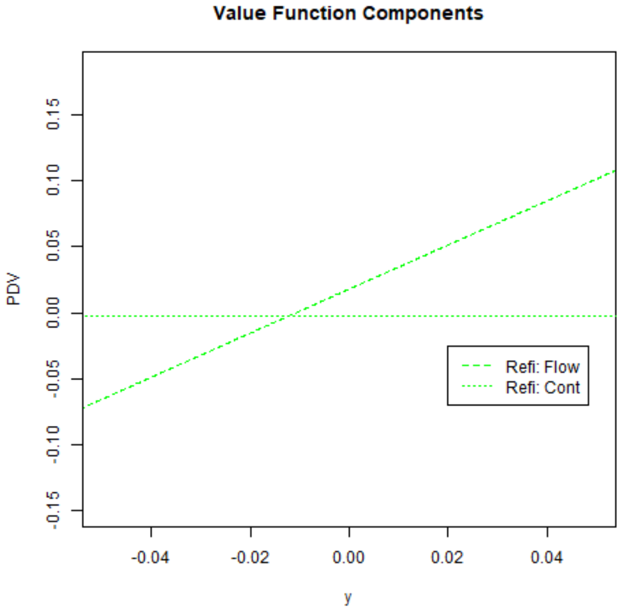

Mortgage Refinancing Solution, Visualized

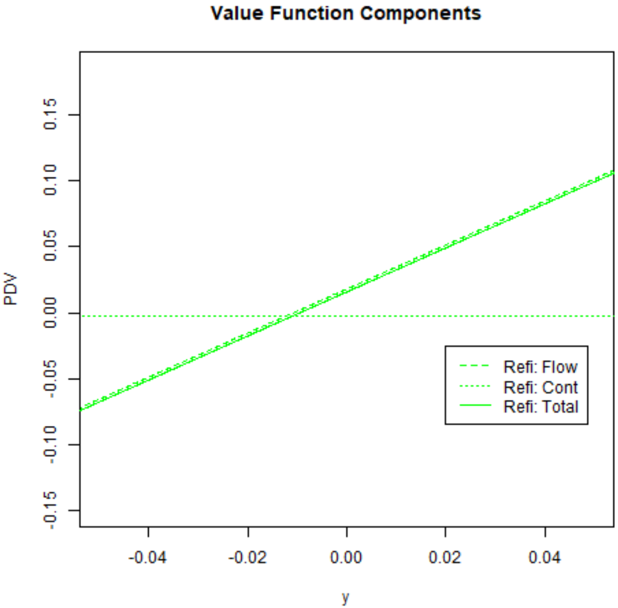

Mortgage Refinancing Solution, Visualized

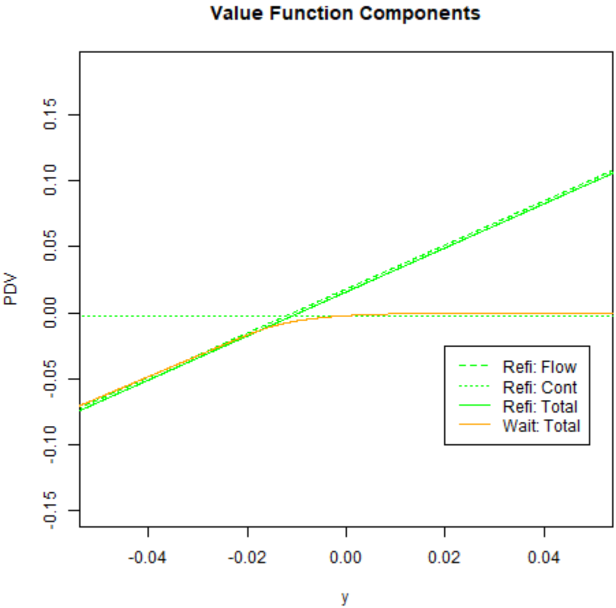

Mortgage Refinancing Solution, Visualized

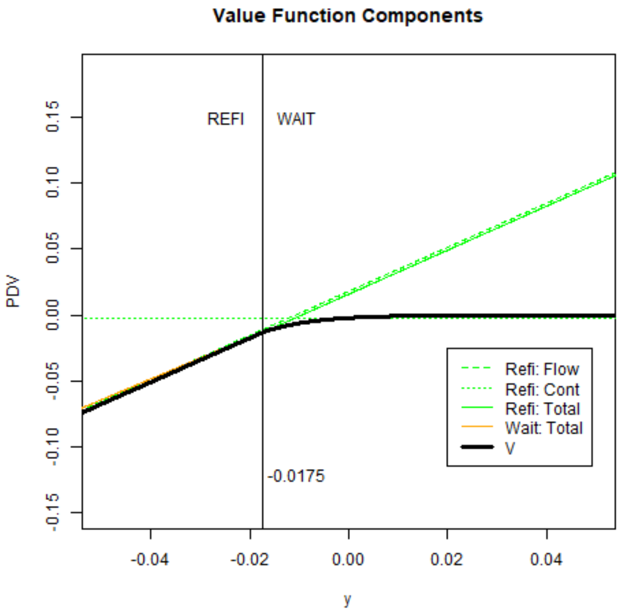

Reflective Barrier

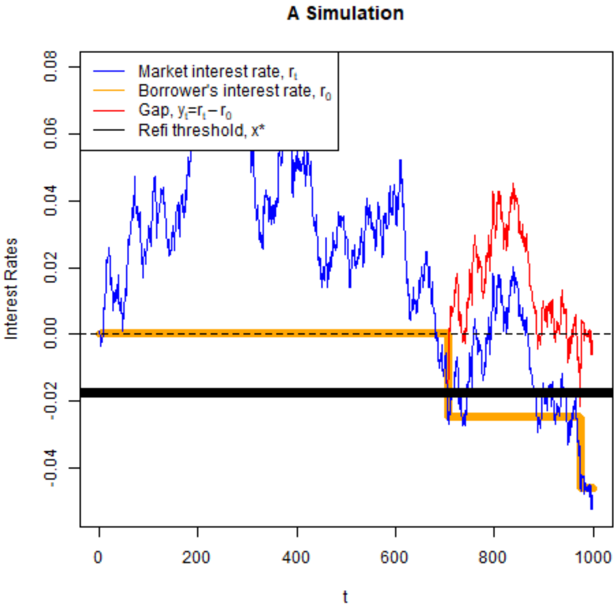

Reflective Barrier

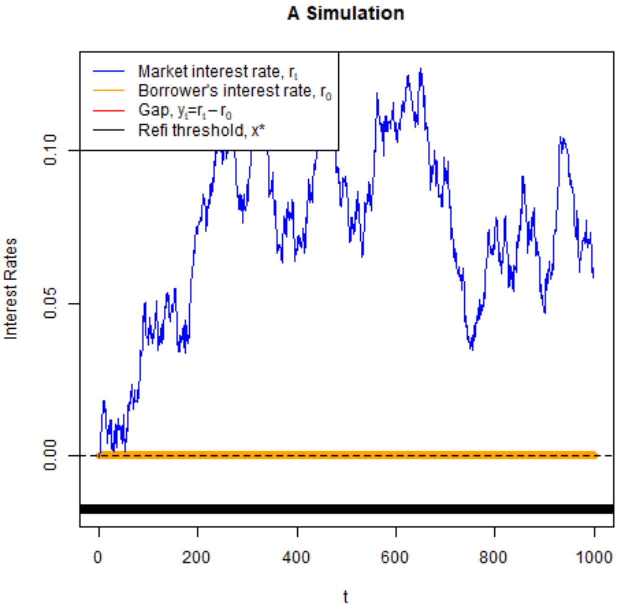

Reflective Barrier

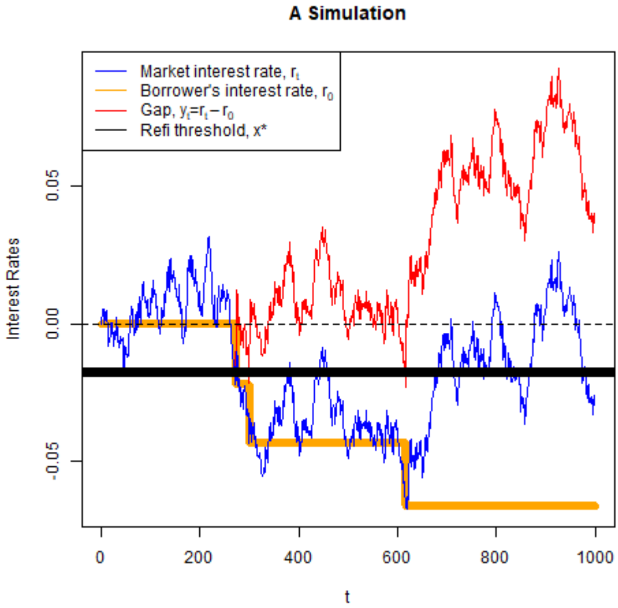

Reflective Barrier

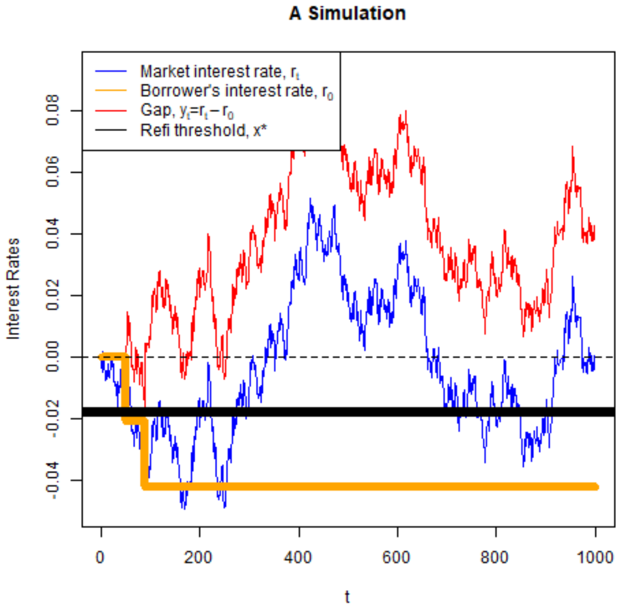

Finite-Horizon Problems

So far we have looked at "stationary," infinite horizon problems

- Every period is just like any other, except the state variable might be different
- There is no explicit dependence on time.

Many economic models violate this. Life-cycle models are common examples:

- For a given level of savings (x), a person might spend a different amount when they are old as opposed to young
- Age (time) explicitly matters here

So suppose your agent lives for a finite set of periods: $1, 2, ..., T$.

- You then need to solve for T distinct value functions.
- $V_1(x)$ gives the value in your first period for any x, $V_T(x)$ gives it in the last period

Solving a Finite-Horizon Problem: Backwards Induction

Conceptually, solving finite-horizon problems is quite easy.
In the final period, T, $V_T(x)$ should be easy to solve for, since there is only a flow utility: no continuation value to worry about:

$$V_T(x) = \max_{\theta_T} \left\{ u(x, \theta_T) \right\} \tag{21}$$

For the previous period, we have:

$$V_{T-1}(x) = \max_{\theta_{T-1}} \left\{ u(x, \theta_{T-1}) + \beta \cdot E[V_T(x_T)|x, \theta_{T-1}] \right\} \tag{22}$$

There are no clever tricks or iterating that need to be done: $V_{T-1}(x)$ is solved for in one try.
You can follow this process backwards for all $t \in \{0, 1, ..., T-1\}$:

$$V_t(x) = \max_{\theta_t} \left\{ u(x, \theta_t) + \beta \cdot E[V_{t+1}(x_{t+1})|x, \theta_t] \right\} \tag{23}$$

Finite-Horizon, Example

Take our simple example, but assume the problem only lasts 3 periods. So:

1. Period 3:
   - $V_3(x) =$
   - $x_3^* =$
2. Period 2:
   - $V_2(x) =$
   - $x_2^* =$
3. Period 1:
   - $V_1(x) =$
   - $x_1^* =$

Finite-Horizon, Example

Take our simple example, but assume the problem only lasts 3 periods. So:

1. Period 3:
   - $V_3(x) = x$
   - $x_3^* = 0$
2. Period 2:
   - $V_2(x) =$
   - $x_2^* =$
3. Period 1:
   - $V_1(x) =$
   - $x_1^* =$

Finite-Horizon, Example

Take our simple example, but assume the problem only lasts 3 periods. So:

1. Period 3:
   - $V_3(x) = x$
   - $x_3^* = 0$
2. Period 2:
   - $V_2(x) = \max\{x, \beta/2\}$
   - $x_2^* = \beta/2$
3. Period 1:
   - $V_1(x) =$
   - $x_1^* =$

Finite-Horizon, Example

Take our simple example, but assume the problem only lasts 3 periods. So:

1. Period 3:
   - $V_3(x) = x$
   - $x_3^* = 0$
2. Period 2:
   - $V_2(x) = \max\{x, \beta/2\}$
   - $x_2^* = \beta/2$
3. Period 1:
   - $V_1(x) = \max\{x, \beta^2/2 + \frac{\beta}{2}(1 - \beta/2)^2\}$
   - $x_1^* = \beta^2/2 + \frac{\beta}{2}(1 - \beta/2)^2$

If $\beta = 0.95$, $x_1^* \approx 0.582$, $x_2^* = 0.475$, $x_3^* = 0$,

Finite-Horizon Solution, Visualized