OPTIMAL CONTROL AND DYNAMIC PROGRAMMING

PAUL SCHRIMPF
NOVEMBER 14, 2013
UNIVERSITY OF BRITISH COLUMBIA
ECONOMICS 526

1. INTRODUCTION

In the past few lectures we have focused on optimization problems of the form

$$\max_{x \in U} f(x) \text{ s.t. } h(x) = c$$

where $U \subseteq \mathbb{R}^n$. The variable that we are optimizing over, x, is a finite dimensional vector. There are interesting optimization problems in economics that involve an infinite dimensional vector.

Example 1.1. [Consumption-savings] An infinite horizon consumption-savings problem,

$$\max_{\{c_t\}_{t=0}^{\infty}, \{s_t\}_{t=1}^{\infty}} \sum_{t=0}^{\infty} \beta^t u(c_t) \text{ s.t. } s_{t+1} = (1 + r_t)(s_t - c_t),$$

involves maximizing over a countably infinite sequence of c_t and s_t. The interpretation of this problem is that $u(c)$ is the per-period utility from consumption. c_t is consumption at time t. s_t is the savings you have at time t. r_t is the return to savings at time t in period $t + 1$, and β is the discount factor.

Example 1.2. [Contracting with a continuum of types] On problem set 6, we studied a problem where a price discriminating monopolist was selling a good to two types of consumers. We could also imagine a similar situation where there is a continuous distribution of types. A consumer of type θ gets 0 utility from not buying the good, and $\theta v(q) - T$ from buying q units of the good at cost T. Let the types be indexed by $\theta \in [\theta_l, \theta_h]$ and suppose the density of θ is f_θ. The seller does not observe consumers' types. However, the seller can offer a menu of contracts $(q(\theta), T(\theta))$ such that type θ will choose contract $(q(\theta), T(\theta))$. The seller chooses the contracts to maximize profits subject to the requirement that each type chooses its contract.

$$\max_{q(\theta), T(\theta)} \int_{\theta_l}^{\theta_h} [T(\theta) - cq(\theta)] f_\theta(\theta) d\theta$$

$$\text{s.t.}$$

$$\theta v(q(\theta)) - T(\theta) \geq 0 \forall \theta \tag{1}$$

$$\theta v(q(\theta)) - T(\theta) \geq \max_{\tilde{\theta}} \theta v\left(q(\tilde{\theta})\right) - T(\tilde{\theta}) \forall \theta \tag{2}$$

The first constraint (1) is referred to as the participation (or individual rationality) constraint. It says that each type θ must prefer buying the good to not. The second constraint

(2) is referred to as the incentive compatibility constraint. It says that type θ must prefer buying the θ contract instead of pretending to be some other type $\tilde{\theta}$. This approach to contracting with asymmetric information—that you can setup contracts such that each type chooses a contract designed for it—is called the revelation principle because the choice of contract makes the consumers reveal their types. Note that this setup does not just apply to price discriminating monopolists. There are many other applications. For example, if you consider a government that needs to raise a certain amount of revenue by taxing workers with heterogeneous productivity, then you end up with essentially the same problem. θ would be worker productivity. $q(\theta)$ is the labor supplied by type θ, $\theta v(q(\theta))$ is the output of type θ, and $T(\theta)$ is the tax. The government maximizes total output subject to a revenue constraint and the constraints above.

Anyway, in these sorts of problems, we need to take the maximum with respect to a function instead of a finite dimensional vector.

Dynamic programming and optimal control are two approaches to solving problems like the two examples above. In economics, dynamic programming is slightly more often applied to discrete time problems like example 1.1 where we are maximizing over a sequence. Optimal control is more commonly applied to continuous time problems like 1.2 where we are maximizing over functions. However, dynamic programming can also be applied to continuous time problems and optimal control can be applied to discrete time problems. Optimal control and dynamic programming problems are often treated quite differently than finite dimensional optimization problems. Indeed there are specialized techniques for solving optimal control and dynamic programming problems that do not appear in finite dimensional optimization. However, the basic theory of infinite dimensional and finite dimensional optimization are the same. For infinite dimensional optimization problems, we can get the exact same first and second order conditions as we did in the finite dimensional case. To do this, we need to define derivatives in abstract (in particular infinite dimensional) vector spaces.

1.1. **References.** The last chapter of Chiang and Wainwright is a good practical introduction to optimal control and Pontryagin's maximum principle. A classic reference for optimization on vector spaces is *Optimization by vector space methods* by Luenberger (1969). *Applied dynamic programming* by Bellman and Dreyfus (1962) and *Dynamic programming and the calculus of variations* by Dreyfus (1965) provide a good introduction to the main idea of dynamic programming, and are especially useful for contrasting the dynamic programming and optimal control approaches. Stokey and Lucas *Recursive methods in economics dynamics* (1989) is the standard economics reference for dynamic programming. Bertsekas's *Dynamic programming and stochastic control* is the standard reference for dynamic programming with uncertainty. Acemoglu's *Introduction to Modern Economic Growth* includes two very nice chapters on optimal control and dynamic programming.

2. Differentiation in vector spaces

We covered differentiation in vector spaces already. We briefly review the main ideas here.

Recall from earlier that there are many sets of functions that are vector spaces. We talked a little bit about $\mathcal{L}^p$ spaces of functions. The set of all continuous functions and the sets of all k times continuously differentiable functions are also vector spaces. One of these vector spaces of functions will be appropriate for finding the solution to optimal control problems like example 1.2. Exactly which vector space is a slightly technical problem dependent question, so we will not worry about that for now (and we may not worry about it at all). Similarly, there are vector spaces of infinite sequences. Little ℓ^p is similar to big $\mathcal{L}^p$, but with sequences instead of functions

$$\ell^p = \{\{x_t\}_{t=1}^{\infty} : \left(\sum_{t=1}^{\infty} |x_t|^p\right)^{1/p} < \infty\}$$

There are others as well. Again, the right choice of vector space depends on the problem being considered, and we will not worry about it too much.

Recall our definition of the derivative for a function $f : \mathbb{R}^n \to \mathbb{R}^m$ was that Df_x satisfies

$$\lim_{\|h\| \to 0} \frac{\|f(x+h) - f(x) - Df_x h\|}{\|h\|} = 0$$

$\mathbb{R}^n$ and $\mathbb{R}^m$ are just two particular vector spaces. To write down this definition of the derivative, we needed to have norms, so to define derivatives in an abstract vector space, the space should have a norm. The definition of a derivative also involves taking limits. Recall that a complete normed vector space is a normed vector space where all Cauchy sequences converge. This ensures that limits are well behaved. We call complete normed vector spaces **Banach spaces**. Finally, Df_x is an m by n matrix. It is a linear transformation from $\mathbb{R}^n \to \mathbb{R}^m$.

Definition 2.1. Let V and W be Banach spaces and $f : V \to W$. The **Fréchet derivative** of f at $x \in V$ is a continuous linear transformation from V to W, denoted Df_x such that

$$\lim_{\|h\| \to 0} \frac{\|f(x+h) - f(x) - Df_x h\|}{\|h\|} = 0.$$

Notice that this definition of the derivative is exactly the same as our definition of the derivative for functions from $\mathbb{R}^n \to \mathbb{R}^m$. It is called the Fréchet derivative because there are other notions of the derivative on abstract vector spaces. The Gâteaux derivative is a directional derivative. It is a weaker concept in that Fréchet differentiability implies Gâteaux differentiability, but not vice versa. There is also an intermediate type of derivative called the Hadamard derivative, which is sometimes useful in econometrics and statistics. We will only use Fréchet derivatives in this class, so henceforth, whenever we talk about the derivative of a function on Banach spaces, we mean the Fréchet derivative.

Given two Banach spaces, V and W, let $BL(V, W)$ denote the set of all linear transformations from V to W. The derivative of $f : V \to W$ will be in $BL(V, W)$. We can show that $BL(V, W)$ is a vector space, and we can define a norm on $BL(V, W)$ by

$$\|D\|_{BL(V,W)} = \sup_{\|v\|_V = 1} \|Dv\|_W$$

where $\|\cdot\|_V$ is the norm on V and $\|\cdot\|_W$ is the norm on W. Moreover, we could show that $BL(V, W)$ is complete. Thus, $BL(V, W)$ is also a Banach space. Viewed as function of x, Df_x is a function from V to $BL(V, W)$. As we just said, there are both Banach spaces, so can differentiate Df_x with respect to x. In this way, we can define the second and higher derivatives of $f : V \to W$.

With this definition of the derivative, almost everything that we proved in lecture 8 for functions from $\mathbb{R}^n \to \mathbb{R}^m$ also holds functions on Banach spaces. In particular, Taylor's theorem, the implicit function theorem, and the inverse function theorem hold. The proofs of these theorems on Banach spaces are essential the same as in lecture 8, so we will not go over them. The mean value theorem is slightly more delicate. It still holds for functions $f : V \to \mathbb{R}^m$, but not when the target space of f is infinite dimensional. If you find this interesting, you may want to go through the proofs of all these claims, but it is not necessary to do so.

3. Optimization in vector spaces

In the previous section we saw that differentiation for functions on Banach spaces is the same as for functions on finite dimensional vector spaces. All of our proofs of first and second order conditions only relied on Taylor expansions and some properties of linear transformations. Taylor expansions and linear transformations are the same on Banach spaces as on finite dimensional vector spaces, so our results for optimization will still hold. Let's just state for the first order condition for equality constraints. The other results are similar, but stating them gets to be slightly cumbersome.

Theorem 3.1 (First order condition for maximization with equality constraints)**.** *Let $f : U \to \mathbb{R}$ and $h : U \to W$ be continuously differentiable on $U \subseteq V$, where V and W are Banach spaces. Suppose $x^* \in$ interior(U) is a local maximizer of f on U subject to $h(x) = 0$. Suppose that $Dh_{x^*} : V \to W$ is onto. . Then, there exists $\mu^* \in BL(W, \mathbb{R})$ such that for*

$$L(x, \mu) = f(x) - \mu h(x).$$

we have

$$\begin{aligned} D_x L(x^*, \mu^*) =& Df_{x*} - \mu Dh_{x^*} = 0_{BL(V,\mathbb{R})} \\ D_\mu L(x^*, \mu^*) =& h(x^*) = 0_W \end{aligned}$$

There are a few differences compared to the finite dimensional case that are worth commenting on. First, in the finite dimensional case, we had $h : U \to \mathbb{R}^m$, and the condition that rank$Dh_{x^*} = m$. This is the same as saying that the $Dh_{x^*} : \mathbb{R}^n \to \mathbb{R}^m$ is onto. Rank is not well-defined in infinite dimension, so we now state this condition as Dh_{x^*} being onto instead of being rank m.

Secondly, previously $\mu \in \mathbb{R}^m$, and the Lagrangian was

$$L(x, \mu) = f(x) - \mu^T h(x).$$

Viewed as a 1 by m matrix, μ^T is a linear transformation from $\mathbb{R}^m$ to $\mathbb{R}$. Thus, in the abstract case, we just say $\mu \in BL(W, \mathbb{R})$, which as when we defined transposes, is called the dual space of W and is denoted W^*.

Finally, we have subscripted the zeroes in the first order condition with $BL(V, \mathbb{R})$ and W to emphasize that the first equation is for linear transformations from V to $\mathbb{R}$, and the second equation is in W. Df_{x*} is a linear transformation from V to $\mathbb{R}$. Dh_{x*} goes from V to W. μ goes from W to $\mathbb{R}$, so μ composed with Dh_{x*}, which we just denoted by μDh_{x*} is a linear transformation from V to $\mathbb{R}$.

4. Optimal control

Optimal control is just one example of optimization in a vector space.

4.1. **Continous time optimal control.**

The classic continuous time optimal control problem has the following form:

$$\begin{aligned}
\max_{x(t),y(t)} & \int_0^T F(x(t), y(t), t)dt \\
& \text{s.t.} \\
& \frac{dy}{dt} = g(x(t), y(t), t) \forall t \in [0, T] \\
& y(0) = y_0
\end{aligned}$$

x is called a control variable, and y is called a state variable. The choice of x controls the evolution of y through the first constraint.

In terms of our notation for optimization in vector spaces, the space being maximized over, V is pairs of functions $(x(t), y(t))$ from $[0, T]$ to $\mathbb{R}$. The objective function is

$$f(x, y) = \int_0^T F(x(t), y(t), t)dt.$$

It is a mapping from V to $\mathbb{R}$. The constraint, $h(x, y)$ is a function from V to a space of functions from $[0, T]$ to $\mathbb{R}$. As above, we will call that space W, so $h : V \to W$ is

$$h(x, y)(t) = \frac{dy}{dt}(t) - g(x(t), y(t), t).$$

Let's apply theorem 3.1, and write out the first order condition. The multiplier, μ is a linear transformation from W to $\mathbb{R}$. There is also the constraint that $y_0 = y(0)$. Let μ_0 be the multiplier on that constraint. The Lagrangian, $L(x, y, \mu, \mu_0)$ is a map from $V \times BL(W, \mathbb{R}) \times \mathbb{R}$ to $\mathbb{R}$ given by

$$L(x, y, \mu, \mu_0) = \int_0^T F(x(t), y(t), t)dt - \mu\left(\frac{dy}{dt} - g(x, y, \cdot)\right) - \mu_0(y(0) - y_0).$$

I wrote $\frac{dy}{dt} - g(x, y, \cdot)$ to emphasize that this is some element of W, a function of t. μ takes this function of t and returns a real number. Many such transformations can be written as integrals. Let's assume that

$$\mu(w) = \int_0^T w(t)\lambda(t)dt$$

for some function λ. We can guarantee that this is true by appropriately defining V and W.[1] It may help to compare this to the finite dimensional case where if have constraints $h : \mathbb{R}^n \to \mathbb{R}^m$ we write $\mu^T h(x) = \sum_{j=1}^m \mu_j h_j(x)$. In this problem we have infinite constraints, so instead of a sum, we use an integral.

The first order conditions are

$$
\begin{aligned}
[x] : && D_x f_{x^*,y^*} - \mu D_x h_{x^*,y^*} =& 0 && (3)\\
[y] : && D_y f_{x^*,y^*} - \mu D_y h_{x^*,y^*} - \mu_0 =& 0 && (4)\\
[\mu] : && h(x^*, y*) =& 0 && (5)
\end{aligned}
$$

Let's work out exactly what each of these derivatives are. $D_x f$ is a linear map from $V \to \mathbb{R}$. Elements of V are pairs of functions from $[0, T]$ to $\mathbb{R}$. If we consider a single function from $[0, T]$ to $\mathbb{R}$, v, and look at

$$\frac{d}{da} \left(f(x + av, y) \right) |_{a=0}$$

, we get the directional derivative of f in direction v at x, y. By analogy with the finite dimensional case, when f is differentiable, all its directional derivatives exist, are equal, and

$$\frac{d}{da} \left(f(x + av, y) \right) |_{a=0} = D_x f v.$$

So to describe $D_x f$ it is enough to look at the directional derivatives of f.

$$
\begin{aligned}
\frac{d}{da} f(x + av, y) =& \frac{d}{da} \int_0^T F(x(t) + av(t), y(t), t) dt \\
=& \int_0^T \frac{\partial F}{\partial x}(x(t), y(t), t) v(t) dt
\end{aligned}
$$

where we need F to be continuously differentiable for all $t \in [0, 1]$ in a neighborhood of $(x(t), y(t), t)$ so that interchanging the derivative and integral is allowed. Thus,

$$D_x f_{x,y} v = \int_0^T \frac{\partial F}{\partial x}(x(t), y(t), t) v(t) dt.$$

Notice that this expression is linear in v. Similarly,

$$
\begin{aligned}
\frac{d}{da} f(x, y + av) =& \frac{d}{da} \int_0^T F(x(t), y(t) + av(t), t) dt \\
=& \int_0^T \frac{\partial F}{\partial y}(x(t), y(t), t) v(t) dt
\end{aligned}
$$

so

$$D_y f_{x,y} v = \int_0^T \frac{\partial F}{\partial y}(x(t), y(t), t) v(t) dt.$$

[1] One sufficient condition is to make W a Sobolev space that includes generalized functions. Sobolev spaces that include generalized functions are Hilbert spaces, so it follows from the Riesz representation theorem that $\mu(w) = \langle w, \lambda \rangle$. However, we only briefly mentioned Hilbert spaces, and we haven't talked at all about Sobolev spaces, generalized functions, or the Riesz representation theorem, so this footnote is likely very confusing and should be ignored.

Now we calculate the derivative of h. $h : V \to W$, so its derivative is a linear map from V to W. Here, W is a function from $[0, T]$ to $\mathbb{R}$, so $D_x h$ takes functions from $[0, T] \to \mathbb{R}$ and maps them to functions from $[0, T] \to \mathbb{R}$. Using the same reasoning as above,

$$\begin{aligned} \frac{d}{da} h(x + av, y)(t) =& \frac{d}{da} \left[\frac{dy}{dt}(t) - g(x(t) + av(t), y(t), t) \right] \\ =& - \frac{\partial g}{\partial x}(x(t), y(t), t) v(t) \end{aligned}$$

so

$$D_x h(v)(t) = -\frac{\partial g}{\partial x}(x(t), y(t), t) v(t).$$

Next,

$$\begin{aligned} \frac{d}{da} h(x, y + av)(t) =& \frac{d}{da} \left[\frac{d(y + av)}{dt}(t) - g(x(t), y(t) + av(t), t) \right] \\ =& \frac{dv}{dt}(t) - \frac{\partial g}{\partial y}(x(t), y(t), t) v(t) \end{aligned}$$

so

$$D_y h(v)(t) = \frac{dv}{dt}(t) - \frac{\partial g}{\partial y}(x(t), y(t), t) v(t).$$

Finally, as we assumed above, it can be shown that $\mu(w) = \int_0^T w(t)\lambda(t) dt$ for some function λ.

Combining all these facts, we can write the first order conditions for x and y as

$$\begin{aligned} [x]: \quad 0 =& \int_0^T \frac{\partial F}{\partial x}(x(t), y(t), t) v(t) dt - \int_0^T -\frac{\partial g}{\partial x}(x(t), y(t), t) v(t) \lambda(t) dt \\ =& \int_0^T v(t) \left(\frac{\partial F}{\partial x}(x(t), y(t), t) + \frac{\partial g}{\partial x}(x(t), y(t), t) \lambda(t) \right) dt \\ [y]: \quad 0 =& \int_0^T \frac{\partial F}{\partial y}(x(t), y(t), t) v(t) dt - \int_0^T \left(\frac{\partial v}{\partial t}(t) - \frac{\partial g}{\partial y}(x(t), y(t), t) v(t) \right) \lambda(t) dt - \mu_0 v(0) \\ =& \int_0^T v(t) \left(\frac{\partial F}{\partial y}(x(t), y(t), t) + \frac{\partial g}{\partial y}(x(t), y(t), t) \lambda(t) \right) dt - \int_0^T \frac{dv}{dt}(t) \lambda(t) dt - \mu_0 v(0) \\ =& \int_0^T v(t) \left(\frac{\partial F}{\partial y}(x(t), y(t), t) + \frac{\partial g}{\partial y}(x(t), y(t), t) \lambda(t) \right) dt + \int_0^T \frac{d\lambda}{dt}(t) v(t) dt - \\ & - \lambda(T) v(T) + \lambda(0) v(0) - \mu_0 v(0). \end{aligned}$$

The last line comes from integration by parts. These equations must be zero for all $v \in V$. This can be true for all v only if the integrands are zero everywhere. That is if

$$\begin{aligned} [x]: & \quad & 0 =& \frac{\partial F}{\partial x}(x(t), y(t), t) + \frac{\partial g}{\partial x}(x(t), y(t), t) \lambda(t) \\ [y]: & \quad & -\frac{d\lambda}{dt}(t) =& \frac{\partial F}{\partial y}(x(t), y(t), t) + \frac{\partial g}{\partial y}(x(t), y(t), t) v(t) \lambda(t) \\ [\mu]: & \quad & \frac{dy}{dt} =& g(x(t), y(t), t) \end{aligned}$$

and $\lambda(T) = 0$ and $\mu_0 = \lambda(0)$. These are the usual conditions that you get from using the Hamiltonian, which you may have seen before. The function $\lambda(t)$ is called the **costate** variable. The following theorem summarizes this result.

Theorem 4.1 (Pontryagin's maximum principle)**.** *Consider*

$$\max_{x,y \in U \subseteq X \times Y} \int_0^T F(x(t), y(t), t) dt$$

$$s.t. \tag{6}$$

$$\frac{dy}{dt} = g(x(t), y(t), t) \forall t \in [0, T] \tag{7}$$

$$y(0) = y_0.$$

where X and Y are some Banach spaces of differentiable functions from $[0, T]$ to $\mathbb{R}$, and $F, g : \mathbb{R}^2 \times [0, T] \rightarrow \mathbb{R}$ are continuously differentiable. Define the Hamiltonian as

$$H(x, y, \lambda, t) = F(x(t), y(t), t) + \lambda(t) g(x(t), y(t), t).$$

If x^* and y^* are a local constrained maximum of (7) in the interior of U, then there exists $\lambda^*(t)$ such that

$$\begin{aligned}
[x] : \quad & 0 = \frac{\partial H}{\partial x}(x^*, y^*, \lambda^*, t) \\
[y] : \quad & -\frac{d\lambda}{dt}(t) = \frac{\partial H}{\partial y}(x^*, y^*, \lambda^*, t) \\
[\lambda] : \quad & \frac{dy}{dt}(t) = \frac{\partial H}{\partial \lambda}(x^*, y^*, \lambda^*, t)
\end{aligned}$$

[2]

4.2. **Application: optimal contracting with a continuum of types.** Let's solve example 1.2.

$$\max_{q(\theta), T(\theta)} \int_{\theta_l}^{\theta_h} [T(\theta) - cq(\theta)] f_\theta(\theta) d\theta$$

s.t.

$$\theta v (q(\theta)) - T(\theta) \geq 0 \forall \theta \tag{8}$$

$$\theta v (q(\theta)) - T(\theta) \geq \max_{\tilde{\theta}} \theta v \left(q(\tilde{\theta})\right) - T(\tilde{\theta}) \forall \theta \tag{9}$$

First, notice that if the participation constraint (8) holds for type θ_l, and (9) holds for θ, then the participation constraint must also hold for θ.

[2]The condition for x is often stated somewhat more generally as

$$H(x^*, y^*, \lambda^*, t) = \max_x H(x, y^*, \lambda^*, t).$$

We can show that the incentive compatibility constraint (9) is equivalent to the following local incentive compatibility constraint and monotonicity constraint.

$$\theta v'(q(\theta))q'(\theta) - T'(\theta) = 0 \tag{10}$$

$$\frac{dq(\theta)}{d\theta} \geq 0 \tag{11}$$

Consider the incentive compatibility constraint (9). The first order condition for the maximization is

$$\theta v'(q(\tilde{\theta}))q'(\tilde{\theta}) = T'(\tilde{\theta}).$$

This is the same as the local incentive compatibility constraint with $\theta = \tilde{\theta}$.

The second order condition for (9) is

$$\theta v''(q(\tilde{\theta}))q'(\tilde{\theta})^2 + \tilde{\theta} v'(q(\tilde{\theta}))q''(\tilde{\theta}) - T''(\tilde{\theta}) \leq 0$$

On the other hand if we differentiate the local incentive compatibility constraint we get

$$\begin{aligned} v'(q(\theta))q'(\theta) + \theta v''(q(\theta))q'(\theta)^2 + \theta v'(q(\theta))q''(\theta) - T''(\theta) &= 0 \\ \theta v''(q(\theta))q'(\theta)^2 + \theta v'(q(\theta))q''(\theta) - T''(\theta) &= -v'(q(\theta))q'(\theta) \end{aligned}$$

We assume that $v' > 0$, and the monotonicity constraint says that $q' \geq 0$. Hence, this equation implies the second order condition. Therefore, we have shown that the local incentive compatibility constraint and monotonicity constraint are equivalent incentive compatibility constraint.

Now, we can write the seller's problem as

$$\begin{aligned} \max_{q(\theta),T(\theta)} \int_{\theta_l}^{\theta_h} & [T(\theta) - cq(\theta)] f_\theta(\theta) d\theta \\ \text{s.t.} & \\ & \theta_l v(q(\theta_l)) - T(\theta_l) \geq 0 \qquad (12) \\ & \theta v'(q(\theta))q'(\theta) - T'(\theta) = 0 \qquad (13) \\ & \frac{dq(\theta)}{d\theta} \geq 0 \qquad (14) \end{aligned}$$

The first order condition for T is for any $x : [\theta_l, \theta_h] \to \mathbb{R}$,

$$\begin{aligned} 0 &= \int_{\theta_l}^{\theta_h} x(\theta) f_\theta(\theta) d\theta - \int_{\theta_l}^{\theta_h} \mu(\theta) \frac{dx}{d\theta}(\theta) d\theta + \mu_0 x(\theta_l) \\ 0 &= \int_{\theta_l}^{\theta_h} x(\theta) \left(f_\theta(\theta) + \frac{d\mu}{d\theta}(\theta) \right) - \mu(\theta_h) x(\theta_h) + \mu(\theta_l) x(\theta_l) + \mu_0 x(\theta_l). \end{aligned}$$

From this we see that $\mu'(\theta) = -f_\theta(\theta)$, $\mu(\theta_l) = -\mu_0$, and $\mu(\theta_h) = 0$. Given μ' and $\mu(\theta_h)$, it must be that

$$\begin{aligned}\mu(\theta) =& \int_{\theta_h}^{\theta} -f_\theta(\hat{\theta})d\hat{\theta} \\ =& 1 - \int_{\theta_l}^{\theta} f_\theta(\hat{\theta})d\hat{\theta} \\ =& 1 - F_\theta(\theta)\end{aligned}$$

where F_θ is the cdf of f_θ. The first order condition for q is

$$\begin{aligned}0 =& \int_{\theta_l}^{\theta_h} cx(\theta)f_\theta(\theta)d\theta - \int_{\theta_l}^{\theta_h} \mu(\theta)\left(\theta v''(q(\theta))q'(\theta)x(\theta) + \theta v'(q(\theta))x'(\theta)\right)d\theta - \\ & - \mu_0\theta_l v'(q(\theta_l))x'(\theta_l) - \int_{\theta_l}^{\theta_h} \lambda(\theta)x'(\theta)d\theta \\ 0 =& \int_{\theta_l}^{\theta_h} x(\theta)\left(cf_\theta(\theta) - \mu(\theta)\left(\theta v''(q(\theta))q'(\theta)x(\theta) - v'(q(\theta)) - \theta v''(q(\theta))q'(\theta)\right)\right)d\theta + \\ & + \int_{\theta_l}^{\theta_h} \mu'(\theta)\theta v'(\theta)d\theta - \mu(\theta_h)\theta_h v'(q(\theta_h))x(\theta_h) + \mu(\theta_l)\theta_l v'(q(\theta_l))x(\theta_l) - \\ & - \mu_0\theta_l v'(q(\theta_l))x'(\theta_l) - \int_{\theta_l}^{\theta_h} \lambda(\theta)x'(\theta)d\theta \\ 0 =& \int_{\theta_l}^{\theta_h} x(\theta)\left(cf_\theta(\theta) + \mu(\theta)v'(q(\theta)) + \mu'(\theta)\theta v'(\theta)\right)d\theta - \\ & - \mu(\theta_h)\theta_h v'(q(\theta_h))x(\theta_h) + \mu(\theta_l)\theta_l v'(q(\theta_l))x(\theta_l) - \\ & - \mu_0\theta_l v'(q(\theta_l))x'(\theta_l) - \int_{\theta_l}^{\theta_h} \lambda(\theta)x'(\theta)d\theta\end{aligned}$$

If we assume that the monotonicy constraint does not bind, so $\lambda(\theta) = 0$, we see that

$$\begin{aligned}0 =& cf_\theta(\theta) + \mu(\theta)v'(q(\theta)) + \mu'(\theta)\theta v'(q(\theta)) \\ 0 =& cf_\theta(\theta) + (1 - F_\theta(\theta))v'(q(\theta)) - f_\theta(\theta)\theta v'(q(\theta)) \\ \theta v'(q(\theta)) =& c + \frac{1 - F_\theta(\theta)}{f_\theta(\theta)}v'(q(\theta)) \\ \left(\theta - \frac{1 - F_\theta(\theta)}{f_\theta(\theta)}\right)v'(q(\theta)) =& c\end{aligned}$$

You may recall from problem set 6 that with symmetric information, $\theta v'(q(\theta)) = c$. Since v' is decreasing in q, this implies that $q(\theta)$ is less than what it would be in the first best symmetric information case for all $\theta < \theta_h$. The highest type, θ_h gets the optimal level of consumption since $F_\theta(\theta_h) = 1$.

4.3. **Discrete time optimal control.** We can also consider discrete time optimal control problems of the form

$$\max_{x_t,y_y} \sum_{t=0}^{\infty} F(x_t,y_t,t)$$
$$\text{s.t.}$$
$$y_{t+1} - y_t = g(x_t,y_t,t)$$

with y_0 given. For simplicity, we will assume $x_t \in \mathbb{R}$ and $y_t \in \mathbb{R}$, but x_t and y_t could be elements of any normed vector spaces. Example 1.1 has this form with $x_t = c_t$, $y_t = s_t$, $F(x_t,y_t,t) = \beta^t u(c_t)$, and $g(x_t,y_t,t) = r_t s_t - (1+r_t)c_t$. We can apply theorem 3.1 to this problem. Here we are maximizing over a pair of infinite sequences, $\{x_t,y_t\}_{t=0}^{\infty} \in V$, where V is a Banach space of such sequences.[3] The constraint can be written as

$$h(x,y)_t = y_{t+1} - y_t - g(x_t,y_t,t).$$

h takes the pair of infinite sequences, $x,y \in V$ and returns another infinite sequence, which is in some Banach space W. The multiplier, μ is a map from W to $\mathbb{R}$. If W were finite dimension, we would have $\mu(w) = \mu^T w = \sum_{i=}^{m} \mu_m w_m$. It turns out that for infinite sequences, μ, we could show that μ must have the same form, so $\mu(w) = \sum_{t=1}^{\infty} \mu_t w_t$. Thus, we can write the Lagrangian as

$$L(x,y,\mu) = \sum_{t=1}^{\infty} \left(F(x_t,y_t,t) - \mu_t(y_{t+1} - y_t - g(x_t,y_t,t)\right).$$

Note that the derivative of L with respect to x is a linear transformation from V to $\mathbb{R}$. If we consider $v \in V$ and look at the directional derivative in direction v, like we did in the continuous case, we see that

$$D_x L(v) = \sum_{t=1}^{\infty} \left(\frac{\partial F}{\partial x}(x_t,y_t,t)v_t + \mu_t \frac{\partial g}{\partial x}(x_t,y_t,t)v_t \right) = 0$$

This must hold for all v_t, so the summand must be zero for each t. Applying similar reasoning to the first order conditions for y and μ, we obtain:

$$\begin{aligned}
[x]: && \frac{\partial F}{\partial x}(x_t,y_t,t) + \mu_t \frac{\partial g}{\partial x}(x_t,y_t,t) &= 0 \\
[y]: && \mu_{t-1} - \mu_t = \frac{\partial F}{\partial y}(x_t,y_t,t) + \mu_t \frac{\partial g}{\partial y}(x_t,y_t,t) &= 0 \\
[mu]: && y_{t+1} - y_t &= g(x_t,y_t,t)
\end{aligned}$$

and $y_0 = 0$ Notice the similarity to theorem 4.1. The only difference is that these equations involve $\mu_t - \mu_{t-1}$ instead of $\frac{d\lambda}{dt}$ and $y_{t+1} - y_t$ instead of $\frac{dy}{dt}$.

[3] As in the continuous time case, we will be somewhat vague about this space because its details depends on the problem at hand. For example, it would often be appropriate to use the space of all bounded sequences, $\ell^\infty = \{x_t, y_t : x_t \leq M, y_t \leq M \forall t \text{ for some } M < \infty\}$. Other times it would be appropriate to use some other space. For example, in 1.1 it might make sense to consider any sequence such that $\sum_{t=0}^{\infty} \beta^t u(c_t)$ is finite, or if $r_t = r$ is fixed, it might make sense to look at any sequence such that $\sum_{t=0}^{\infty} \frac{c_t}{(1+r)^t}$ is finite. We will not worry about these details.

5. Dynamic programming

For problems like examples 1.1 and 1.2, optimal control focuses on characterizing x^* through the first order conditions (given x^*, y^* is easily determined through the constraint). That is, optimal control focus on characterizing the maximizer. An alternative approach is to focus on the value of the maximized function. This value will depend on the entire problem, but in particular it depends on the initial condition y_0. Thus, we can think of the value as function of the initial state. Dynamic programming focuses on characterizing the value function.

The basic idea of dynamic programming can be illustrated in a familiar finite dimensional optimization problem. Consider a finite horizon discrete time consumption savings choice.

$$\max_{c_t, s_t} \sum_{t=0}^{T} \beta^t u(c_t) \text{ s.t. } s_{t+1} = (1 + r_t)(s_t - c_t)$$

with s_0 given, and the constraint that $s_{T+1} = 0$. We could just write down the first order conditions and try to solve them for c_t. However, if T is large, this might be very difficult. It can be especially difficult to calculate a solution numerically. The easiest maximization problems to solve numerically are ones where the objective function is linear or quadratic. In either of these cases, the amount work needed is proportional to the number of variables cubed. If T is large, T^3 can be so large that computing a solution takes prohibitively large.

We can divide this T dimensional problem to a series of smaller ones by first thinking about what happens at time T. At time T we have some savings s_T and want to choose c_T to solve

$$\max_{c_T} u(c_T) \text{ s.t. } s_{T+1} = (1 + r_T)(s_T - c_T) = 0$$

As long as u is increasing, it must be that $c_T^*(s_T) = s_T$. If we define the value of savings at time T as

$$V_T(s) = u(s),$$

then at time $T - 1$ given s_{T-1}, we can choose c_{T-1} to solve

$$\max_{c_{T-1}, s'} u(c_{T-1}) + \beta V_T(s') \text{ s.t.} s' = (1 + r_{T-1})(c_{T-1} - s_{T-1}).$$

This is relatively simple maximization problem with just two variables, so we can solve it without too much difficulty. Repeating in this way, for each t we can define the value of savings at time t as

$$V_t(s) = \max_{c_t, s'} u(c_t) + \beta V_{t+1}(s') \text{ s.t.} s' = (1 + r_t)(c_t - s). \tag{15}$$

This approach to sequential optimization was first proposed by Richard Bellman, so (15) is called a Bellman equation. Notice that if $(c_t^*(s_t), s_{t+1}^*(s_t)$ is a maximizer of (15) for each t, then the sequence of $c_0^*(s_0), s_1^*(s_0), c_1^*(s_1), ..., c_T^*$ must be a maximizer of the original problem. Bellman called this observation the **principle of optimality**. He described it as, "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 and Dreyfus 1962).

In finite horizon problems, it easy to see that the Bellman equations will exist. However, if we have an infinite horizon problem,

$$\max_{c_t, s_t} \sum_{t=0}^{\infty} \beta^t u(c_t) \text{ s.t. } s_{t+1} = (1 + r_t)(s_t - c_t)$$

then we cannot start from the last period to define the value function. However, if the problem is stationary, that is if the problem at time t and at time $t + 1$ are the same, then it seems reasonable to think that the value function would not depend on t and we could just write

$$V(s) = \max_{c, s'} u(c) + \beta V(s') \text{ s.t. } s' = (1 + r)(s - c).$$

Stokey and Lucas (1989) provide a fairly comprehensive analysis of various conditions when this is possible. We will just look at one case.

Consider a problem that is slightly more general than the consumption savings choice problem with fixed interest rate.

$$\max_{c_t, s_t} \sum_{t=0}^{\infty} \beta^t u(c_t, s_t)$$
$$s_{t+1} = g(c_t, s_t),$$

where $c \in \mathbb{R}$, $s \in \mathbb{R}$, $0 < \beta < 1$, and $u, g : \mathbb{R}^2 \to \mathbb{R}$. We want to show that the value function exists. Suppose we start with some guess at the value function $v_0 : \mathbb{R} \to \mathbb{R}$. Then we refine that guess by setting

$$v_1(s) = \max_{c, s'} u(c, s) + \beta v_0(s') \text{ s.t. } s' = g(c, s)$$

We could do this repeatedly. That is, if we let T be the operator defined by this equation,

$$T(v)(s) = \max_{c, s'} u(c, s) + \beta v_0(s') \text{ s.t. } s' = g(c, s),$$

we can construct a sequence with $v_{i+1} = T(v_i)$. Recall from lecture 9 that T is a contraction mapping if

$$\left\| T(v) - T(v') \right\| \leq c \left\| v - v' \right\|$$

for some $c < 1$ and all v, v'. We proved that contraction mappings have unique fixed points. Thus, if we can show that T is a contraction mapping, then $v_i \to V$, where V is the value function that satisfies the Bellman equation, like we want.

Suppose u is bounded, $u(c, s) \leq M$ for all c and s. Then

$$\sum_{t=0}^{\infty} \beta^t u(c_t, s_t) \leq \frac{M}{1 - \beta}.$$

Therefore, we should only look at possible value functions with $v(s) \leq \frac{M}{1-\beta}$ for all s. To show that T is a contraction, we must define the space that v lies in and its norm. Given the boundedness of v, it is natural to look at $\mathcal{L}^{\infty}(\mathbb{R})$, the space of all bounded real-valued function with norm $\|v\| = \sup_{x \in \mathbb{R}} |v(x)|$. Consider

$$T(v_0)(s) - T(v_1)(s) = \left(u(c_0, s) + \beta v_0(s_0')\right) - \left(u(c_1, s) + \beta v_1(s_1')\right)$$

where c_i, s_i' is the maximizer to

$$\max_{c_i, s_i'} u(c_i, s) + \beta v_i(s_i') \text{ s.t. } s_i' = g(c_i, s).$$

We should assume something that ensures c and s' exist. Assuming c', s lie in a compact set would be sufficient. Notice that

$$Tv_0(s) = u(c_0, s) + \beta v_0(s_0') \geq u(c_1, s) + \beta v_0(s_1').$$

Therefore,

$$T(v_0)(s) - T(v_1)(s) \geq \beta(v_0(s_1') - v_1(s_1').$$

Similarly,

$$T(v_0)(s) - T(v_1)(s) \leq \beta(v_0(s_0') - v_1(s_0').$$

It follows that

$$\begin{aligned} \|T(v_0) - T(v_1)\| &= \sup_s |T(v_0)(s) - T(v_1)(s)| \\ &\leq \sup_s |\beta(v_0(s) - v_1(s))| = \beta \|v_0 - v_1\|. \end{aligned}$$

Hence, T is a contraction and has a unique fixed point V.
Advantages of dynamic program. Dynamic programming and optimal control can both be used to solve the same sort of problems. Optimal control has the advantage that it uses very directly what we know about optimization in $\mathbb{R}^n$ and applies it to infinite dimensional spaces. Dynamic programming has the advantage that it lets us focus on one period at a time, which can often be easier to think about than the whole sequence. Because it only requires maximizing over a few variables at a time, dynamic programming can be a much more efficient way to calculate solutions. The computational advantage of dynamic programming is especially pronounced when some of the variables being maximized over are discrete. We will see some examples of this below.

5.1. **Solving dynamic programs.** There are three ways to solve a dynamic program. They are:

(1) Guess and verify the form of the value function
(2) Iterate the Bellman equation analytically
(3) Iterate the Bellman equation numerically

If you guess correctly, the first method is fairly straightforward. However, guessing correctly is difficult and often is not possible at all. The second method will always work, but may not lead to a closed form expression, and can be tedious. The third method is the main way dynamic programs are solved in practice, but we will not go into the details.

Example 5.1 (Optimal growth by guessing and verifying)**.** Consider an economy with a single infinitely lived representative consumer with per-period log utility from consumption and a discount factor of δ. The economy's production function is Cobb-Douglas with capital as the only input. Anything not consumed at time t becomes capital at time $t+1$. The optimal growth problem is

$$\max_{\{c_t\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \delta^t \log(c_t)$$
$$\text{s.t. } c_t + k_{t+1} = k_t^{\alpha}.$$

If we use the constraint to solve for c_t and substitute into the objective, then we have

$$\max_{\{k_t\}_{t=1}^{\infty}} \sum_{t=0}^{\infty} \delta^t \log(k_t^{\alpha} - k_{t+1})$$
$$\text{s.t.} 0 \leq k_{t+1} \leq k_t^{\alpha}$$

The Bellman equation for this problem is

$$v(k) = \max_{k' \in [0, k^{\alpha}]} \log(k^{\alpha} - k') + \delta v(k')$$

Now, we guess the functional form of v. Since the per-period utility function is logarithmic and production is Cobb-Douglas, it is sensible to guess that $v(k) = c_0 + c_1 \log(k^a)$ where c_0, c_1, and a are each constant for which we solve. Now, since $c_1 \log(k^a) = c_1 a \log(k)$, a and c_1 are redundant, so we can get rid of a, and just guess that $v(k) = c_0 + c_1 \log(k)$.

We now use the Bellman equation to solve for c_0 and c_1. First we solve for the optimal k' for a given c_0 and c_1. The Bellman equation is:

$$c_0 + c_1 \log k = \max_{k' \in [0, k^{\alpha}]} \log(k^{\alpha} - k') + \delta \left(c_0 + c_1 \log k'\right).$$

We could write the Lagrangean with the constraints that $k' \geq 0$ and $k' \leq k^{\alpha}$. If we were not sure whether these constraints would bind we would include them in the Lagrangean and check the complementary slackness conditions. However, it is slightly easier to just notice that these constraints cannot bind because utility approaches $-\infty$ as k approaches k^{α} and the next period's value approaches $-\infty$ as k' approaches 0, so neither constraint will bind. Without the constraints, the first order condition is:

$$-\frac{1}{k^{\alpha} - k'} + \delta c_1 \frac{1}{k'} = 0$$
$$-k' + \delta c_1 (k^{\alpha} - k') = 0$$
$$k' = \frac{\delta c_1}{1 + \delta c_1} k^{\alpha}$$

Now, we plug this back into the Bellman equation and solve for c_0 and c_1 by varying k.

$$\begin{aligned} c_0 + c_1 \log k &= \max_{k' \in [0,k^\alpha]} \log(k^\alpha - k') + \delta\left(c_0 + c_1 \log k'\right) \\ &= \log\left(k^\alpha - \frac{\delta c_1}{1+\delta c_1} k^\alpha\right) + \delta\left(c_0 + c_1 \log\left(\frac{\delta c_1}{1+\delta c_1} k^\alpha\right)\right) \\ &= \log\left(\frac{1}{1+\delta c_1}\right) + \alpha \log k + \delta\left(c_0 + c_1 \log\left(\frac{\delta c_1}{1+\delta c_1}\right) + \alpha \log k\right) \\ &= \underbrace{\log\left(\frac{1}{1+\delta c_1}\right) + \delta c_1 \log\left(\frac{\delta c_1}{1+\delta c_1}\right)}_{=c_0} + \underbrace{(\alpha + \delta c_1 \alpha)}_{=c_1} \log k \end{aligned}$$

Both the left and right sides of this equation are affince function of $\log k$. They can only be equal for all k if the coefficients are equal. Thus,

$$\begin{aligned} c_1 =& \alpha + \delta c_1 \alpha \\ c_1 =& \frac{\alpha}{1-\delta\alpha} \end{aligned}$$

and

$$\begin{aligned} c_0 =& \log\left(\frac{1}{1+\delta c_1}\right) + \delta c_1 \log\left(\frac{\delta c_1}{1+\delta c_1}\right) \\ =& -\left(1+\delta\frac{\alpha}{1-\delta\alpha}\right)\log\left(1+\delta\frac{\alpha}{1-\delta\alpha}\right) + \delta\frac{\alpha}{1-\delta\alpha}\log\left(\delta\frac{\alpha}{1-\delta\alpha}\right) \\ =& \log(1-\delta\alpha) + \frac{\delta\alpha}{1-\delta\alpha}\log(\delta\alpha). \end{aligned}$$

Finally, we should make sure that this solution doesn't violate the constraint. We have

$$k' = \frac{\delta c_1}{1+\delta c_1} k^\alpha = \delta\alpha k^\alpha,$$

so the constraints are satisfied as long as $\delta\alpha \in (0,1)$.

If we cannot guess the form of the value function, we can try to find it by repeatedly applying the Bellman operator. The Bellman operator is the T operator we defined above,

$$T(v)(s) = \max_{c,s'} u(c,s) + \beta v_0(s') \text{ s.t. } s' = g(c,s).$$

We already showed that T is a contraction (provided u is bounded and $|\beta| < 1$). Among other things, this means that if we start with an arbitrary guess of the value function, v_0, and then construct a sequence by repeatedly applying T, i.e.,

$$v_i = T(v_{i-1}),$$

then the sequence v_i will converge to a unique fixed point, v, that satisfies the Bellman equation.

Example 5.2 (Optimal growth by iterating)**.** The same optimal growth problem as in the previous example can also be solved by iterating the Bellman operator. We start with *any* guess of the value function for v_0. A common choice is the zero function, $v_0(k) = 0$ for all k. Then we find v_1 by solving

$$\begin{aligned} v_1(k) = & \max_{k' \in [0,k^\alpha]} \log(k^\alpha - k') + \delta v_0(k) \\ = & \max_{k' \in [0,k^\alpha]} \log(k^\alpha - k') \\ = & \alpha \log k. \end{aligned}$$

Then, we repeat to find v_2.

$$\begin{aligned} v_2(k) = & \max_{k' \in [0,k^\alpha]} \log(k^\alpha - k') + \delta v_1(k) \\ = & \max_{k' \in [0,k^\alpha]} \log(k^\alpha - k') + \delta \alpha \log(k') \\ = & c_2 + (\alpha + \delta \alpha^2) \log k, \end{aligned}$$

where c_2 is some constant that involves δ, α, and their logs. The third equality comes from writing the first order condition, solving for k', and subsituting back into the objective. We can explicitly solve for c_2, but it doesn't matter for the first order condition for v_3, so we don't need to know it exactly. We repeat again to get v_3

$$\begin{aligned} v_3(k) = & \max_{k' \in [0,k^\alpha]} \log(k^\alpha - k') + \delta v_1(k) \\ = & \max_{k' \in [0,k^\alpha]} \log(k^\alpha - k') + \delta \alpha \log(k') \\ = & c_3 + (\alpha + \delta \alpha^2 + \delta^2 \alpha^3) \log k. \end{aligned}$$

We could repeat again to get

$$\begin{aligned} v_4(k) = & c_4 + (\alpha + \delta \alpha^2 + \delta^2 \alpha^3 + \delta^3 \alpha^4) \log k \\ v_5(k) = & c_5 + (\alpha + \delta \alpha^2 + \delta^2 \alpha^3 + \delta^3 \alpha^4 + \delta^4 \alpha^5) \log k \\ & \vdots \end{aligned}$$

etc. Eventually, we hopefully notice a pattern. The more obvious pattern is that each v_i and will always be of the form $v_i(k) = c_i + m_i \log(k)$. Thus, we know that $v(k)$ will have that same form and we can go back to the guess and verify method. Better yet, we could notice that

$$v_i(k) = c_i + \alpha \sum_{j=0}^{i} (\alpha\delta)^j \log(k),$$

so

$$v(k) = C + \frac{\alpha}{1 - \alpha\delta} \log k.$$

If we care about C, we could find it by either explicitly writing c_i in terms of δ and α and taking the limit; or using the guess verify method just for C.

Solving for the value function, whether by guessing and verifying or iterating can be a bit tedious. Even worse, for most specifications of the per-period payoff u and constraints g, there will be no closed form solution for v. That makes it impossible to guess the form, and iterating the Bellman equation will not lead to a discernible pattern (although it will still give a convergent series). Using a computer to solve for the value function avoids both these problems. A computer does not care that Bellman operator iteration is tedious, and it can numerically compute $v(k)$ even if it has no closed form.

Another situation where dynamic programs can be solved analytically is when the control variable is discrete. For example, a person could be choosing to work or not each period, or a firm could be choosing to enter or exit a market. The section below and the last problem on problem set 6 are examples of dynamic programming with discrete control variables.

5.2. **Application: Diamond-Mortensen-Pissarides search model.** The standard neo-classical and neo-Keynesian macroeconomic models do not have any involuntary unemployment. If we include a labor-leisure choice, these models may not have everyone working (or at least not working full-time), but everyone who does not work does so voluntarily. We need to add something to the model if we want there to be involuntary unemployment. The standard way of modeling involuntary unemployment is through a search model. We will go through a simple version of a search model. It is often called a Diamond-Mortensen-Pissarides search model because those three people were the first to propose it.

There is a continuum of identical workers with total measure one. There is also a continuum of identical firms. Time is discrete. Workers can be either unemployed or employed. Firms can either employ one worker, post a vacancy, or do nothing. There is free entry. Posting a vacancy costs k. When a firm has a worker, they produce output y. Matches are dissolved with exogenous probability s. Unemployed workers produce a benefit of b each period. Workers' utility is just the discounted sum of their consumption.

Let u_t be the mass of unemployed workers at time t, and let v_t be the mass of vacancies. That is,

$$u_t = \int_{\text{all workers}} 1\{\text{worker } i \text{ unemployed}\} di$$

$$v_t = \int_{\text{all firms}} 1\{\text{firm } j \text{ unemployed}\} dj.$$

There is some matching technology $m(u_t, v_t)$ such that the probability that any given unemployment worker finds a vacancy at time t is $\frac{m(u_t,v_t)}{u_t}$. Similarly, the probability that any given vacant firm finds a worker is $\frac{m(u_t,v_t)}{v_t}$. We will assume that m has constant returns to scale, so in particular,

$$m(u_t, v_t) = u_t m(1, v_t/u_t) = v_t m(u_t/v_t, 1)$$

We define labor market tightness as $\theta_t = \frac{v_t}{u_t}$. Then we can write the unemployed worker's and vacant firm's matching probabilities as

$$\frac{m(u_t, v_t)}{u_t} = m(1, v_t/u_t) = m(1, \theta_t) \equiv \mu(\theta_t)$$

and

$$\frac{m(u_t, v_t)}{v_t} = \frac{m(1, v_t/u_t)u_t}{v_t} = \frac{\mu(\theta_t)}{\theta_t}.$$

To ensure that these are valid probabilities assume that $0 \leq \mu(\theta) \leq \min\{1, \theta\}$. Also, we will assume that μ is twice continuously differentiable with $\mu' > 0$ and $\mu'' < 0$.

5.2.1. *Social planner.* Suppose a social planner wants to maximize utility subject to the matching and production technologies.

$$\begin{aligned} \max_{c_t, v_t, u_t} & \sum_{t=0}^{\infty} \beta^t \left((1-u_t)c_t^e + u_t c_t^u\right) \\ & \text{s.t.} \\ & c_t^e + c_t^u = u_t b + (1-u_t)y - k v_t \\ & u_{t+1} = (1-\mu(v_t/u_t))u_t + s(1-u_t) \end{aligned}$$

Note that all that matters is total consumption, not c_t^e and c_t^u separately. Therefore we can eliminate the first constraint, leaving

$$\begin{aligned} \max_{v_t, u_t} & \sum_{t=0}^{\infty} \beta^t \left(u_t b + (1-u_t)y - k v_t\right) \\ & u_{t+1} = (1-\mu(v_t/u_t))u_t + s(1-u_t). \end{aligned}$$

Also, it is common to work with θ_t instead of v_t. Since $v_t = u_t\theta_t$, we can write the problem as

$$\begin{aligned} \max_{\theta_t, u_t} & \sum_{t=0}^{\infty} \beta^t \left(u_t b + (1-u_t)y - k\theta_t u_t\right) \\ & u_{t+1} = (1-\mu(\theta_t))u_t + s(1-u_t). \end{aligned}$$

The Bellman equation for this problem is

$$\begin{aligned} V(u) = & \max_{\theta, u'} ub + (1-u)y - k\theta u + \beta V(u') \\ & \text{s.t. } u' = (1-\mu(\theta))u + s(1-u) \end{aligned}$$

A common method for solving dynamic programming problems is to guess the form of the solution and then verify. It is often the case that $V(u)$ has the same form as the per-period payoff. Here, the per-period payoff is linear in u, so a good guess is that

$$V(u) = \alpha_0 + \alpha_1 u.$$

We plug this guess in for V, then use the first order condition for the Bellman equation to find out what α_0 and α_1 must be. If there are no α_0 and α_1 that make the first order condition hold, then our guess was incorrect. If our guess was correct, we will be able to uniquely solve for α_0 and α_1. Substituting the guess we have

$$\begin{aligned} \alpha_0 + \alpha_1 u = & \max_{\theta, u'} ub + (1-u)y - k\theta u + \beta\left(\alpha_0 + \alpha_1 u'\right) \\ & \text{s.t. } u' = (1-\mu(\theta))u + s(1-u) \end{aligned}$$

The first order conditions are

$$\begin{aligned} 0 = & -ku - \lambda\mu'(\theta)u \\ 0 = & \beta\alpha_1 - \lambda \end{aligned}$$

To solve for α_1 note that by the envelope theorem,

$$\alpha_1 = b - y - k\theta + \beta\alpha_1 (1 - \mu(\theta) - s)$$
$$\alpha_1 = \frac{b - y - k\theta}{\beta (1 - \mu(\theta) - s)}$$

From the first order condition,

$$\mu'(\theta^*) = \frac{-k}{\beta\alpha_1}$$
$$\mu'(\theta^*) = \frac{-k (1 - \mu(\theta^*) - s)}{(b - y - k\theta^*)} \tag{16}$$

In particular, θ^* does not depend on u. Therefore, $u^{*\prime}$ only depends on u through the constraint. Since the constraint is linear in u, $u^{*\prime}$ will also be linear in u. Finally, the guess $V(u) = \alpha_0 + \alpha_1 u$ is linear in u, so $V(u^{*\prime})$ will also be linear in u. Therefore, our guess is verified. By itself, solving the social planner problem is not very insightful. The one useful thing is that we know that the efficient level of labor market tightness satisfies (16). We can use this equation as a benchmark to compare what happens under other conditions.

5.2.2. *Competitive equilibrium.* Due to the matching friction, wages cannot be determined by supply and demand in the usual way. A worker cannot just go and find another job at the prevailing wage. Likewise for a firm. Together, a matched worker and firm are strictly better off than an unmatched worker and firm. Matched workers and firms earn a surplus, and we need some way of deciding how to divide this surplus. The typical way is through Nash bargaining. If the total output of a match is y, the bargaining power of the worker is η, the outside option of the worker is o_w and the outside option of the firm is o_f, then the Nash bargained wage solve

$$\max_w (w - o_w)^\eta (y - w - o_f)^{1-\eta}.$$

The first order condition is

$$0 = \eta(w - o_w)^{\eta-1}(y - w - o_f)^{1-\eta} - (1 - \eta)(w - o_w)^\eta (y - w - o_f)^{-\eta}.$$

Rearranging,

$$\frac{\eta}{1-\eta} = \frac{w - o_w}{y - w - o_f}.$$

The wage is such that benefit to the worker relative to the benefit to the firm is equal to the ratio of their bargaining powers.

To apply this result to our model, we must find the outside options of the worker and firm. That is, we must find the value of being unemployed and of being a vacant firm. Since there is free entry, the value of being a vacant firm must be zero. For simplicity, let's focus on the steady state when $\theta_t = \theta$ for all t. For workers, let $v(w)$ denote the value of value of being employed given wage w and $u(w)$ be the value of being unemployed. The

Bellman equations for these are

$$\begin{aligned} v(w) =& w + \beta\left((1-s)v(w) + su(w)\right) \\ u(w) =& b + \beta\left(\mu(\theta)v(w) + (1-\mu(\theta))u(w)\right) \end{aligned}$$

Similarly, if $j(w)$ is the value of a firm having a worker, and $i(w)$ is the value of posting a vacancy, then must satisfy

$$\begin{aligned} j(w) =& y - w + \beta(1-s)j(w) \\ i(w) =& 0 = -k + \beta\frac{\mu(\theta)}{\theta}j(w) \end{aligned}$$

$i(w)$ is zero because of free entry, but we still wrote down its Bellman equation since it will be useful later.

Now, we are ready to find the equilibrium wage. The benefit to a worker of getting wage w is not just w, it is $v(w)$. Similarly, the firm gets not $y-w$, but $j(w)$. The outside option of the firm is $i(w) = 0$. The outside option is $u(w^e)$, where w^e is the wage at other firms (the equilibrium wage). The bargaining problem is then

$$\max_w \left(v(w) - u(w^e)\right)^\eta j(w)^{1-\eta}$$

As above, the wage solution satisfies

$$\frac{\eta}{1-\eta} = \frac{v(w) - u(w)}{j(w)}$$

Combining this equation and the Bellman equations, we can solve for the wage and $\mu(\theta)$. The Bellman equation for $j(w)$ gives

$$j(w) = \frac{y-w}{1-\beta(1-s)}$$

Combining with the Bellman equation of $i(w)$ we get

$$\begin{aligned} k =& \beta\frac{\mu(\theta)}{\theta}\frac{y-w}{1-\beta(1-s)} \\ w =& y - k\frac{1-\beta(1-s)}{\beta\mu(\theta)} \end{aligned}$$

...

Eventually we get the Hosios-Mortensen condition, that the competitive equilibrium is efficient only if

$$1 - \eta = \frac{\theta\mu'(\theta)}{\mu(\theta)}.$$