Economics 2450A: Public Economics and Fiscal Policy I
Mathematics and Microeconomics Review

Michael Droste*

Fall 2022

This document is intended to serve as a brief refresher to the mathematical toolkit and microeconomic concepts that you will see in Harvard's public economics sequence. Ec2450 assumes some experience with graduate-level microeconomics and quasi-experimental econometric methods. This review document is not a substitute for this background; it is intended only to serve as a reference and refresh your memory on concepts you have seen in the past.

Section 1 reviews some mathematical tools that we will repeatedly use, including constrained optimization and the envelope theorem. Section 2 refreshes some core concepts from consumer theory. Section 3 provides some optional practice problems for your review. Section 4 provides some references on optimization, microeconomics, and empirical / econometric methods. This is not a comprehensive document; we will use and extend additional material as the semester goes on.

1 Mathematical Tools

1.1 Static Constrained Optimization

In this course, most often we will be considering the behavior of economic agents (i.e. consumers/households, firms, policymakers) whose behavior we characterize as the solution to (usually static) constrained optimization problems. For instance, a household may be choosing consumption and labor supply to maximize utility subject to a budget constraint; a firm may be choosing production inputs, like capital and labor, to maximize profits (perhaps subject to consumer demand); a government may be choosing a tax schedule to maximize social welfare subject to a desired level of tax revenue.

Let's work through a simple example involving a utility maximization problem. A household chooses a bundle $x \in \mathbb{R}^n$ in order to maximize a utility function $u(x) : \mathbb{R}^n \to \mathbb{R}$ subject to a budget or resource constraint $g(x, r) = 0$. This problem can be written as:

$$\max_x u(x) \text{ subject to: } g(x, r) = 0$$

You can think of r as any interesting parameter(s) of the problem that is exogenous to the optimizing agent (e.g. welfare benefits from the government, unearned income, etc...).

The standard approach to solve constrained optimization problems is the Lagrange multiplier method, likely very familiar to you already. The steps are always the same:

1. Construct the Lagrangian function corresponding to the problem:

$$\mathcal{L}(x, r) = u(x) + \lambda g(x, r) \tag{1}$$

   where λ is a constant that we will call the *Lagrangian multiplier*.

*These notes were created from materials shared by past 2450A teaching fellows. All errors are my own.

2. Take partial derivatives of the Lagrangian function $\mathcal{L}$ with respect to x and λ, and set them equal to 0 (thus obtaining the first-order conditions, or FOCs):

$$\forall k \in [1, ..., n] : \frac{\partial \mathcal{L}}{\partial x_k} = \frac{\partial u(x)}{\partial x_k} - \lambda \frac{\partial g(x,r)}{\partial x_k} = 0$$
$$\frac{\partial \mathcal{L}}{\partial \lambda} = g(x,r) = 0$$

Note that the last partial derivative simply gives you back the constraint.

3. Solve the system of $n+1$ equations and unknowns characterized by the first-order conditions and the variables $(x_1, ..., x_n, \lambda)$ for optimal choices of $x_1, ..., x_n$.

4. If the Lagrangian is concave (convex), then the candidate solutions obtained from the first-order conditions solve the maximization (minimization) problem.

1.1.1 Variations on this problem

- The Lagrange multiplier method generalizes easily to a setting with multiple constraints. Each constraint gets added to the Lagrangian function with its own multiplier.

- If you have an inequality constraint, $g(x,r) \geq 0$, you have to add the complementary slackness condition $\frac{\partial L}{\partial \lambda} \cdot \lambda = 0$. This tells us that either (i) the constraint binds, in which case the first-order condition is satisfied with equality ($g(x,r) = 0$); or (ii) the constraint does not bind ($g(x,r) > 0$), and therefore the Lagrange multiplier has to be $\lambda = 0$ (see the interpretation of the multiplier below to see why).

- If we have non-negativity constraints and they are binding, we need to consider the Kuhn-Tucker (KT) Conditions, which state that if x^* is a solution to the problem above, then x^* satisfies the KT conditions. That is, there exists a Lagrange multiplier $\lambda \geq 0$, such that for all $i \in [1, ..., n]$,

$$\frac{\partial u(x^*)}{\partial x_i} \leq \lambda \frac{\partial g(x^*, r)}{\partial x_i},$$

  with equality when $x_i^* > 0$ and $g(x^*, r) = 0$.

For the purposes of this course, we will generally not need to deal with inequality constraints or non-negativity constraints. For the models we consider, there is typically something about the economic environment that guarantees constraints bind with equality at an optimum or in equilibrium. For instance, in a consumer utility maximization problem, if the consumers' preferences exhibit local non-satiation, an optimal solution will always imply a binding budget constraint (Walras' law) - the consumer will always be on the budget frontier.

1.1.2 Interpreting the Lagrange Multiplier

Suppose $x^*(r)$ are the values of x that solve the optimization problem in the preceding section. The associated value of the objective function (e.g. utility) given the optimal x^* is therefore a function of r, with $V(r) = u(x^*(r))$ (the **value function** of the problem). Provided that certain mild regularity conditions are satisfied, we have the following result:

$$\frac{dV(r)}{dr} = \lambda(r)$$

Thus, the Lagrange multiplier $\lambda = \lambda(r)$ is the rate at which the optimal value of the objective function changes with respect to changes in the constraint parameter r. It is sometimes called the **shadow price** of resource r (because it indicates by how much the value function would increase in response to an increase in r).

Loosely, the Lagrangian multiplier λ answers the question, "If the constraint is relaxed by one unit, how much does the objective function change?" For example, in a static consumer utility maximization problem with a single budget constraint, the Lagrangian multiplier would tell us the marginal increase in utility induced by a marginal increase in income.

1.2 Comparative Statics and Implicit Differentiation

In the course, we will often be interested in how an agent's *behavior* changes as parameters of the economic environment (for instance, prices or taxes) change. If we have solved for the first order conditions that govern the agent's behavior, we can differentiate the first order condition(s) w.r.t. the parameter of interest. Suppose that we have a first order condition $\mathcal{L}_x(x,p) = 0$. Here, x is a variable that the agent optimizes over and p is an exogenous parameter of the model that the agent takes into account but does not choose herself. If we're interested in learning how the agent's choice of x changes as p changes, we can (implicitly) differentiate both sides of the equation $L_x(x,p) = 0$ w.r.t. p:

$$\frac{\partial \mathcal{L}_x(x,p)}{\partial x}\frac{dx}{dp} + \frac{\partial \mathcal{L}_x(x,p)}{\partial p} = 0$$

This allows us to express the effect of p on x *even if we cannot explicitly solve for x itself*:

$$\frac{dx}{dp} = -\frac{\frac{\partial \mathcal{L}_x(x,p)}{\partial p}}{\frac{\partial \mathcal{L}_x(x,p)}{\partial x}}$$

Implicit differentiation will work as long as the conditions for the implicit function theorem are met (you do not need to know the technical conditions for this course - except making sure that the denominator is different from 0). Remember the formula:

$$F(x,y) = 0 \Rightarrow \frac{dy}{dx} = \frac{F_x(x,y)}{-F_y(x,y)},\ (F_y(x,y) \neq 0) \tag{2}$$

This formula is known as the **implicit function theorem**.

1.3 Envelope Theorem

The envelope theorem is your best friend in the public economics sequence. It comes up all the time. The envelope theorem loosely says that if we are interested in computing the comparative static of a value function that depends on an optimized quantity with respect to some exogenous variable r, then we can ignore the indirect effect that dr has on the choice variables. It will be useful to illustrate the envelope theorem through a simple example. Nearly all of the instances in which we use the envelope theorem will have the same flavor.

Simple case with two choice variables

Suppose we wish to solve the following maximization problem:

$$V(\alpha) \equiv \left\{ \max_{x,y}\ f(x,y,\alpha) \text{ s.t. } g(x,y,\alpha) = 0 \right\} \tag{3}$$

where x, y are choice variables and α is some exogenous parameter(s). $V(\alpha)$ is the value function corresponding to this problem: it is the *value* that the objective function f takes, given the optimal choice of (x,y), as a function of the exogenous parameter α. For instance, if f is a utility function and g is a budget constraint in a simple static utility maximization problem, the value function V would tell you the level of utility realized by the consumer when they choose the optimal (x,y) as a function of the exogenous parameter(s) α.

The Lagrangian corresponding to this problem is:

$$L(x,y,\lambda,\alpha) = f(x,y,\alpha) + \lambda g(x,y,\alpha) \tag{4}$$

This Lagrangian yields the following first-order conditions:

$$\frac{\partial L}{\partial x} = f_x(x,y,\alpha) - \lambda g_x(x,y,\alpha) = 0$$
$$\frac{\partial L}{\partial y} = f_y(x,y,\alpha) - \lambda g_y(x,y,\alpha) = 0$$

The solution to this problem is given by $x^*(\alpha)$ and $y^*(\alpha)$. Substituting this into the value function, we get that $V(\alpha) = f(x^*(\alpha), y^*(\alpha), \alpha) = L(x^*(\alpha), y^*(\alpha), \lambda^*(\alpha), \alpha)$.

Taking the total derivative of V w.r.t. to exogenous parameter α:

$$\frac{dV}{d\alpha} = f_x \frac{dx^*}{d\alpha} + f_y \frac{dy^*}{d\alpha} + f_\alpha$$

Notice there is a direct effect of α on the value function from the last term. The first two terms are indirect effects in which α affects the control variables x and y.

The first-order conditions imply that at the optimum, $f_x = \lambda^* g_x$ and $f_y = \lambda^* g_y$. Plugging these in:

$$\frac{dV}{d\alpha} = \lambda^* \left(g_x \frac{dx^*}{d\alpha} + g_y \frac{dy^*}{d\alpha} \right) + f_\alpha$$

Totally differentiate the constraint with respect α at the optimum:

$$g_x \frac{dx^*}{d\alpha} + g_y \frac{dy^*}{d\alpha} - g_\alpha = 0$$

which implies that

$$\frac{dV}{d\alpha} = f_\alpha + \lambda^* g_\alpha = \frac{\partial L}{\partial \alpha} \tag{5}$$

This result is known as the **envelope theorem**: a change in the value function w.r.t a small in change in an exogenous parameter is equal to the partial derivative of the Lagrangian with respect to the parameter. The indirect effect **drops out** when we plug in the first-order conditions.

For a formal statement of the envelope theorem, see MWG Theorem M.L.1, p. 965. There are several useful intuitive ways to break down the envelope theorem, depending on how general you want to get. My preferred intuition is the following: at an optimum, changes in α have two impacts on the objective function: a first-order "direct effect" through the impact α has on the constraint or objective function, and a higher-order "indirect effect" where the change in α induces a change in the choice variables that impacts the objective / value function. If we only consider very small changes in α like $d\alpha$, the first-order effect completely dominates the higher-order indirect effect. This is all implicit when we apply the envelope theorem to a comparative static, since in this case $d\alpha$ is an infinitesimally small quantity.

It is important to note that while the envelope theorem only holds exactly for infinitesimally small changes in α, it does **not** rely on the objective function being differentiable with respect to the choice variables (as we assumed in our derivation above), or of continuous choice sets. Indeed, we will see a neat application of the envelope theorem to discrete choices in this course. The rather general conditions for which envelope theorem-like statements hold is covered by Milgrom and Segal (2002).

Why is this result so powerful?

The envelope theorem tells us that in some cases, we can ignore indirect effects of a change in an exogenous parameter on the value function. This is great, because it means that comparative statics involving value functions are often quite simple to compute. In addition, the envelope theorem often gets rid of parameters in our model from a comparative static, which is useful when those parameters might be difficult to estimate in practice.

Concretely, suppose we are interested in a comparative static like how consumer utility changes when the price of a single good changes (du/dp). The indirect effect that a change in p has on the consumer's choice of goods falls out: all that's left is how the price change impacts the consumer's budget. In an ideal case, we might not even need to place a functional form on the utility function to calculate the comparative static. That's great - it's less math to compute the comparative static, and our calculation does not depend on some utility function parameter(s) that we would rather not hang our hats on.

What are the limitations to this approach?

While the envelope theorem is super powerful and convenient, it has some limitations that you need to know. A key part of the course is distinguishing when the envelope theorem fails.

First, the envelope theorem is only exact for infinitesimal policy changes ($d\alpha$ very small). This is assumed in the calculus-based derivation we consider above, where the 'policy change' in our comparative static is $d\alpha$. For large policy changes (for instance, if we want to assess the real-world consequences of a big tax cut), we cannot simply ignore the impact that a change in some paraameter/policy has on value functions indirectly through a quantity being optimized. This is quite important in practice: many empirical papers use estimators or give interpretations to estimands that rely on the envelope theorem, which is a problem if the policy change in question is not small.

Next, the envelope theorem relies on a quantity being optimized, i.e. first-order conditions holding. In the presence of externalities, behavioral agents, incomplete markets, etc., the envelope theorem may not hold. We will encounter several concrete examples in 2450A that will allow us to think about these kinds of limitations in more detail.

2 Microeconomics review

2.1 Utility Maximization

To think about welfare, economists usually assume that consumers are rational actors who make economic decisions (about their consumption, labor supply, etc) according to their preferences. Under mild regularity conditions, a consumer's preferences can be represented by a utility function. Since consumers choose their most preferred element in their budget set, this corresponds to the **utility maximization problem** (UMP):

$$V(p,w) = \max_{x \in B(p,w)} u(x) \tag{6}$$

where $B(p,w) = \{x \in \mathbb{R}^n : p \cdot x \leq w\}$ denotes the budget set (which in this choice of notation reflects the constraint), x is a vector of goods, p is a vector of prices with each entry corresponding to a specific product, and w is exogenous wealth or income.

We say that x^* solves the UMP if x^* maximizes u on $B(p,w)$. We denote a solution to the UMP at prices p and wealth w as $x(p,w) = [x_1(p,w), ..., x_n(p,w)]$. The bundle is known as **Marshallian demand** (also known as uncompensated, gross, or Walrasian demand). Given solutions to the UMP, we plug them into the utility function to obtain the value function for this problem, the **indirect utility function** V:

$$V(p,w) \equiv u(x^*(p,w))$$

Properties of Marshallian Demand

- Walras' law: $p \cdot x(p,w) = w$ for all prices p and wealth w.
- Homogeneity of degree zero in (p,w). Suppose that Marshallian demand is unique at all p and w, then for any $\alpha > 0$, $x(p,w) = x(\alpha p, \alpha w)$.

Properties of Indirect Utility

- $V(p,w)$ is increasing in w and decreasing in p.
- $V(p,w)$ is homogeneous of degree zero in (p,w).
- Roy's identity: For all p and w:

$$x_k(p,w) = -\frac{\partial V(p,w)/\partial p_k}{\partial V(p,w)/\partial w} \tag{7}$$

  for $k = [1, ..., n]$.

2.1.1 Solution to UMP

To formally characterize the solutions to the UMP, we assume that utility is smooth and quasiconcave. For now, it is useful to derive the first order conditions to the UMP assuming that the solution is in the interior.

The Lagrangian corresponding to the UMP is:

$$\mathcal{L}(x, p, w, \lambda) = u(x) + \lambda(w - p \cdot x)$$

where λ is the Lagrange multiplier. The corresponding FOCs are:

$$\forall i \in [1, ..., n]: \quad \frac{\partial u(x)}{\partial x_i} - \lambda p_i = 0$$

Combining the first-order conditions for any two goods i and j yields the following characterizations of the household's optimality conditions:

1. The marginal rate of substitution between any two goods i and j must equal the price ratio:

$$\frac{\partial u(x)/\partial x_i}{\partial u(x)/\partial x_j} = \frac{p_i}{p_j}$$

2. The utility increments between any two goods i and j must be equal at the optimum:

$$\frac{\partial u(x)/\partial x_i}{p_i} = \frac{\partial u(x)/\partial x_j}{p_j}$$

Note that these two statements are equivalent (one is redundant), and both characterize optimal allocation between any two goods i, j.

In the truly unfortunate event when the solution to the UMP is not in the interior, we need to draw on the Kuhn-Tucker (KT) conditions. In this course, we will typically consider models where we assume regularity conditions (i.e. quasi-concave utility) that guarantee interior solutions.

The solution to this utility-maximization problem are the Marshallian demand functions $x(p, w)$ and $y(p, w)$. They tell us the consumer's optimal quantities of goods x and y as a function of the price vector p and income w. We cannot solve explicitly for these functions unless we assume a functional form for utility.

2.2 Expenditure Minimization

In addition to asking what the highest level of utility the household could attain subject to their budget set (the utility-maximization problem), we could also ask: what is the smallest amount of wealth (expenditure) needed to achieve a target utility level $\bar{u}$? This question defines the consumer's expenditure minimization problem (EMP):

$$e(p, \bar{u}) = \min_{x \in \mathcal{U}} \; p \cdot x \tag{8}$$

where $\mathcal{U} = \{x : u(x) \geq \bar{u}\}$.

We say that x^c solves the EMP if x^c minimizes $p \cdot x$ on $\mathcal{U}$. We denote a solution to the EMP at prices p and target utility $\bar{u}$ as $x^c(p, \bar{u}) = [x_1^c(p, \bar{u}), ..., x_n^c(p, \bar{u})]$. The bundle is known as **Hicksian demand** (also known as compensated demand). Here, utility is held constant when prices vary. To keep utility constant, we are changing w in the background or compensating the agent. The solution to the EMP can be derived analogously to how the UMP was solved.

Properties of Hicksian Demand

- Homogeneity of degree zero in prices. For all p, $\bar{u}$, and $\alpha > 0$, $x^c(\alpha p, \bar{u}) = x^c(p, \bar{u})$.
- Symmetry of cross-price effects. For all p and $\bar{u}$,

$$\frac{\partial x^c_\ell(p,\bar{u})}{\partial p_k} = \frac{\partial x^c_k(p,\bar{u})}{\partial p_\ell}$$

  for all goods ℓ, k.
- Compensated law of demand. For all p, p'', and $\bar{u}$,

$$(p'' - p') \cdot (x^c(p'', \bar{u}) - x^c(p', \bar{u})) \leq 0$$

  Hicksian demand is actually downward slopping! Recall that Marshallian demand need not slope downward.

Properties of Expenditure Function

- Increasing in prices and target utility.
- Homogeneity of degree one in prices.
- Concavity in prices: For all p, p' and $\gamma \in (0, 1)$,

$$e(\gamma p + (1-\gamma)p', \bar{u}) \geq \gamma e(p, \bar{u}) + (1-\gamma)e(p', \bar{u})$$

- Shephard's Lemma: For all p and $\bar{u}$,

$$x^c_k(p, \bar{u}) = \frac{\partial e(p,\bar{u})}{\partial p_k} \tag{9}$$

  for all $k = [1, ..., n]$.

2.3 Connection between UMP and EMP

The connection (aka duality) between the consumer's utility-maximization problem and their expenditure-minimization problem is really important! The following fun facts establish that in a certain sense, the solution for the UMP yields a quick solution for the EMP and vice versa:

- If x^* solves the UMP at wealth w, then x^* solves the EMP for target utility $u(x^*)$ and $e(p, u(x^*)) = w$.
- If x^c solves the EMP for target utility $\bar{u}$, then x^c solves the UMP at wealth $p \cdot x^c$ and $V(p, p \cdot x^c) = \bar{u}$.

2.4 Comparative Statics & Slutsky

Comparative statics answer questions of the form, 'How does an endogenous variable change as we vary an exogenous variable?' In the case of the UMP, we are typically concerned about how demand changes as prices vary or $\partial x_l(p, w)/\partial p_k$.

The impact of an uncompensated price change on (Marshallian) demand can be decomposed into two effects:

$$\underbrace{\frac{\partial x_l(p,w)}{\partial p_k}}_{\text{Uncompensated effect}} = \underbrace{\frac{\partial x^c_l(p,\bar{u})}{\partial p_k}}_{\text{Substitution effect}} \underbrace{-\frac{\partial x_l(p,w)}{\partial w} x_k(p,w)}_{\text{Income effect}} \tag{10}$$

This decomposition is known as the **Slutsky equation**. It is derived from the deep connection (duality) between the household's UMP and EMP:

$$\begin{aligned}\frac{\partial x_l^c(p,\bar{u})}{\partial p_k} &= \frac{\partial x_l(p,e(p,\bar{u}))}{\partial p_k} \\ &= \frac{\partial x_l(p,w)}{\partial p_k} + \frac{\partial x_l(p,w)}{\partial e(p,u)} x_k(p,w)\end{aligned}$$

We can write the uncompensated effect in terms of **elasticities** (the percent change in x_l due to a percent change in variable v). In general, for two variables x, y the elasticity of x w.r.t y is $\epsilon_{xy} = \frac{y}{x}\frac{\partial x}{\partial y} = \frac{\partial \log(x)}{\partial \log(y)}$.

Now we've developed all the microeconomics machinery we need to coherently speak about three types of elasticities that may be of interest to us in the real world:

1. **Uncompensated price elasticity** of demand of good l with respect to k:

$$\epsilon_{lk}^u \equiv \frac{p_k}{x_l}\frac{\partial x_l(p,w)}{\partial p_k} \tag{11}$$

   Note that when $l = k$, it is known as the own-price elasticity and when $l \neq k$, the cross-price elasticity.

2. **Compensated price elasticity** of demand of good l with respect to k:

$$\epsilon_{lk}^c \equiv \frac{p_k}{x_l}\frac{\partial x_l^c(p,\bar{u})}{\partial p_k} \tag{12}$$

3. **Income elasticity** of good l:

$$\eta_l = \frac{w}{x_l}\frac{\partial x_l(p,w)}{\partial w} \tag{13}$$

With these two definitions, we re-write the Slutsky equation as:

$$\epsilon_{lk}^u = \epsilon_{lk}^c - b_k \eta_l \tag{14}$$

where $b_k = \frac{p_k x_k}{w}$ is the budget share of good k.

3 Practice Problems

3.1 Envelope theorem and utility maximization

Consider the following utility maximization problem:

$$\max_{(x_1,...,x_n)} u\left(x_1, ..., x_n\right) \qquad \text{s.t. } \sum_i p_i x_i = y$$

where y is exogenous income and p_i denotes the price of x_i.

1. Use the envelope theorem to show that, in an interior solution, the value of the Lagrangian multiplier is equal to the marginal utility of income.
   Solution:
   The Lagrangian is $L = p \cdot x + \lambda(y - p \cdot x)$. The FOCs imply that $\lambda^* p_i = \frac{\partial u}{\partial x_i}$ for all $i = 1, ..., n$, which yield Marshallian demands $x(p, y)$. This defines indirect utility as

   $$V(p, y) = L(p, y) = u(x(p, y)) + \lambda(y - p \cdot x(p, y))$$

   By the envelope theorem, $\frac{dV(p,y)}{dy} = \frac{\partial L}{\partial y} = \lambda$.

2. Show Roy's identity, that is:

   $$x_i^* = -\frac{\frac{\partial V(p_1,...,p_n,y)}{\partial p_i}}{\frac{\partial V(p_1,...,p_n,y)}{\partial Y}}$$

   Solution:
   Starting with the numerator, we have

   $$-\lambda x_i(p, y) + \sum_{k=1}^{n} \left(\underbrace{\frac{\partial u(x)}{\partial x_k} - \lambda p_k}_{=0} \right) \frac{\partial x_k}{\partial p_i} = -\lambda x_i(p, y)$$

   Using the solution from (1), we get the desired result.

   $$x_i^* = -\frac{-\lambda x_i}{\lambda} = x_i(p, y)$$

3.2 Constrained optimization, comparative statics, and the envelope theorem

Suppose that a consumer is solving the following utility maximization problem:

$$\max_{s,c} u(s, c) = \frac{1}{1 + \frac{1}{\beta}} s^{1+\frac{1}{\beta}} + c \qquad \text{subject to: } (p + t) \cdot s + c = w$$

where s denotes consumption of a specific good and c consumption of goods other than s. The consumer faces a price of p and an excise tax of t on s; w denotes the consumer's wealth. The parameter β is strictly less than 0 and strictly larger than -1.

1. Solve for the levels of s and c that maximize utility in this problem.
   Solution:
   We write: $\max_s \frac{s^{1+\beta^{-1}}}{1+\beta^{-1}} + w - (p + t)s$. The first order condition w.r.t. to s is:

   $$s^{1/\beta} - (p + t) = 0$$

   and as a result, we derive demands for s and c:

   $$c = w - (p + t)^{1+\beta}$$
   $$s = (p + t)^{\beta}$$

2. What are the comparative statics of s and c w.r.t. the gross price $q = p + t$ and wealth w?
Solution:
Differentiating explicitly:

$$\frac{\partial c}{\partial q} = -(1+\beta)(p+t)^{\beta}$$
$$\frac{\partial s}{\partial q} = \beta(p+t)^{\beta-1}$$
$$\frac{\partial c}{\partial w} = 1$$
$$\frac{\partial s}{\partial w} = 0$$

3. What is the elasticity of demand for the specific good s w.r.t. the consumer price $q = p + t$?
Solution:
The price elasticity of s w.r.t. to q is

$$\epsilon_{s,q} = \frac{q}{s}\beta q^{\beta-1} = \frac{\beta}{(p+t)^{\beta}}(p+t)^{\beta} = \beta$$

4. Suppose the government increases the tax t by a "small" amount. How does this affect the level of utility achieved by the consumer? You can solve this "by hand" or appeal to the logic of the envelope theorem. For this problem, you should assume that the price p is not affected by the level of the tax t (for instance, because s is traded on the world market).
Solution:
The envelope theorem says we only worry about the direct effects of a policy change (we can ignore behavioral responses). It follows that

$$\frac{dv(p,t,w)}{dt} = \frac{\partial v(p,t,w)}{\partial t} = -s^* = -(p+t)^{\beta} < 0$$

5. Provide some intuition for the result in 4.
Solution:
Small changes in prices might change s^*, but because the consumer is optimizing and has already chosen s^* to maximize utility, a small change dt does not change the value function. Hence, the effect is only through the "direct" effect of taxation on the consumer's budget constraint.

3.3 Recovering Compensated Demand Elasticities via the Slutsky Equation

1. Prove Shephard's lemma for a case with n goods.
Solution:
The expenditure function at the optimum is

$$e(p,\bar{u}) = p \cdot x^c(p,\bar{u}) + \mu(\bar{u} - u(x^c(p,\bar{u}))$$

By the envelope theorem:

$$\frac{de(p,\bar{u})}{dp_j} = \frac{\partial e(p,\bar{u})}{\partial p_j} = x_j^c(p,\bar{u})$$

2. Derive the Slutsky equation for change in p_j for good i.
Solution:
We note that at the optimum, Marshalian and compensated demands coincide: $x(p, e(p,\bar{u})) = x^c(p,\bar{u})$. Differentiating both sides w.r.t. to p_i

$$\frac{\partial x_j^c(p,\bar{u})}{\partial p_i} = \frac{\partial x_j(p,w)}{\partial p_i} + \frac{\partial x_j(p,w)}{\partial e(p,\bar{u})}\frac{\partial e(p,\bar{u})}{\partial p_i}$$
$$= \frac{\partial x_j(p,w)}{\partial p_i} + \frac{\partial x_j(p,w)}{\partial w}x_i(p,w)$$

The second line follows from Shephards Lemma and $x_i = x_i^c$ at the optimum. Rearranging, we get the Slutsky equation:

$$\frac{\partial x_j(p,w)}{\partial p_i} = \frac{\partial x_j^c(p,\bar{u})}{\partial p_i} - \frac{\partial x_j(p,w)}{\partial w} x_i(p,w)$$

3. Write the Slutsky equation in terms of elasticities and the share of income spent on good j.
Solution:
Multiplying the equation on both sides by p_i/x_j, we have

$$\frac{p_i}{x_j}\frac{\partial x_j(p,w)}{\partial p_i} = \underbrace{\frac{p_i}{x_j}}_{=x_i^c}\frac{\partial x_j^c(p,\bar{u})}{\partial p_i} - \frac{p_i}{x_j}\frac{\partial x_j(p,w)}{\partial w} x_i(p,w)$$

$$\epsilon_{j,p_i} = \epsilon_{j,p_i}^c - \frac{p_i x_i(p,w)}{w}\eta_j$$

where $\eta_j = \frac{w}{x_j} \cdot \frac{\partial x_j}{\partial w}$, the income elasticity of good j.

4. Suppose you know the uncompensated elasticities. What further information do you need to compute the compensated elasticity?
Solution: You need the budget share of good j and the income elasticity of good j.

5. Consider a partial equilibrium setup with two goods: good x_1 and the numeraire x_2. The consumer maximizes utility subject to a budget constraint with exogenous income y. Suggest a research design and explain how you would use it to estimate the compensated elasticity of demand for x_1 with respect to its price p.
Solution: There can be many answers for this problem. One way to think about this problem is through an RCT where we randomize locations to get price changes and people to get vouchers. Some groups get both changes. This should allow us to identify this compensated elasticity.

4 Useful References

4.1 Optimization and Modeling

The following chapters of *Microeconomic Theory* by Mas-Collell, Whinston and Green (1995) are a great resource that have a formal treatment of the mathematical methods used in this course:

- Classical Demand Theory (chapter 3)
- Chapters M.E, M.J, M.K, and M.L from the Mathematical Appendix

Note, however, that in this course we will typically deal with "well-behaved" problems satisfying the appropriate regularity conditions that MWG wrestles with in all their gory detail.

4.2 Empirical Tools

This course will cover research making use of a wide variety of empirical tools. These include research designs like event studies, difference-in-differences, instrumental variables, regression discontinuity designs, randomized experiments, matching methods, and panel data methods (often used in combination!). We will assume you have seen these at least in the context of a first-year graduate econometrics course. *Mostly Harmless Econometrics: An Empiricist's Companion* by Angrist and Pischke (2009) is a great high-level resource for these methods, although its coverage is somewhat dated.

We would also be remiss not to recommend a textbook reference for econometrics. *Econometrics* by Hayashi (2001) is a fantastic textbook for the foundations of statistical (frequentist) inference and common estimators in economics. Alternatively, *Econometric Analysis of Cross Section and Panel Data* by

Wooldridge (2010) is a good (somewhat dated) reference on cross-section and panel econometrics.

Be advised that empirical best practices for quasi-experimental methods are always changing. As a rule, reference textbooks cannot describe the frontier of methods, and recently published papers will be your best resource.

References

Angrist, Josh, and Jörn-Steffen Pischke. 2009. *Mostly Harmless Econometrics: An Empiricist's Companion.* Princeton University Press.

Hayashi, Fumio. 2001. *Econometrics.* Princeton University Press.

Mas-Collell, Andreu, Michael D. Whinston, and Jerry R. Green. 1995. *Microeconomic Theory.* Oxford University Press.

Milgrom, Paul, and Ilya Segal. 2002. "Envelope Theorems for Arbitrary Choice Sets." *Econometrica*, 70(2): 583–601.

Wooldridge, Jeffrey. 2010. *Econometric Analysis of Cross Section and Panel Data.* MIT Press.