Optimization and Simulation Modeling in Microeconomics

James R. Markusen
University of Colorado, Boulder

Pre-course notes and math review

Madison WI, January 2022

These notes follow the introductory notes and instructions for pre-course preparation that I sent along with this file. They relate primarily to my sections of the course, which consists of parts of Monday and Tuesday. You could read through these notes before downloading and installing GAMS and working though McCarl's slides, or do some of these things at the same time. (I believe that I read you were given McCarl's slides on usingGAMS Studio, if not, I can send these too.)

It is important that you read through these notes and perhaps refresh your knowledge of some of these points. I've written these notes in a fairly casual and, to a mathematician, fairly sloppy way. I am making no attempt at formal rigor here, just trying to remind us of some of the theory that underlies standard concepts we will use throughout the course.

There will be an emphasis on general-equilibrium problems. You can dust off your knowledge of Shepard's lemma and basic duality which you were probably taught at some point, but never really used. Follow this review, the next part of your preparation for my sections of the course is my paper "Global Comparative Statics in General Equilibrium: model building from theoretical foundations" (submitted to the Journal of Global Economic Analysis). This will introduce you to GAMS code as well.

Consider now a consumer maximizing utility U, which is a function of two goods X_1 and X_2, subject to a linear budget constraint in which M is income and p_1 and p_2 are the prices of the goods. The consumer can chose not to spend all of M, hence the budget constraint is a weak inequality, and the consumption of each good must be non-negative.

$$\max_{X_1,X_2} U(X_1,X_2) \quad \textit{subject to} \quad G(X_1,X_2) = M - p_1X_1 - p_2X_2 \geq 0, \quad X_1,X_2 \geq 0$$

How do we solve for the optimal choices of X_1 and X_2? The key is to convert this programming problem into a system of equations and unknowns. Formally, this is done in steps. First, the Karush-Kuhn-Tucker theorem turns the programming problem into a square system of equations. Then the implicit function theorem tells us that these can be converted into n explicit equations, each as a function of the exogenous parameters only, which in our case we call demand function, or later cost functions. These demand and cost functions will be the building blocks of a great deal of what we will do in this course.

Karush-Kuhn-Tucker (KKT) theorem: There exists non-negative constants λ and

$\mu = (\mu_1, \mu_2)$ such that the necessary condition for a maximum are:

$$\frac{\partial U}{\partial X_i} + \lambda \frac{\partial G}{\partial X_i} + \mu = 0 \qquad \text{first-order conditions (FOC) for } X_1 \text{ and } X_2 \text{ (i = 1,2)}$$

$G(X) \geq 0, \; X \geq 0$	primal feasibility
$\lambda \geq 0, \; \mu \geq 0$	dual feasibility
$\lambda G(X) = 0$	complementary slackness 1 (either $G = 0$ or $\lambda = 0$)
$\mu_i X_i = 0$	complementary slackness 2 (either $\mu = 0$ or $X = 0$)

Economists often construct a problem or simply make an implicit assumption that the optimization problem has an *interior solution*. In the present case, that means that the budget constraint holds with equality (all income is spent) and X_1 and X_2 are strictly positive. In that case, $\mu_1 = \mu_2 = 0$, and $\lambda > 0$. Later in this review, we will go into the KKT conditions in more detail and also show that the dual or "slack" variables μ and λ have interesting and important economic interpretations.

Important thing for our purposes: KKT *converts an optimization problem into a set of equations*. Here we have five equations in five unknowns:

Equations	Unknowns
FOC for X_1	X_1
FOC for X_2	X_2
complementary slackness 1	λ
complementary slackness 2	$\mu_1 \;\; \mu_2$

Implicit function theorem: Subject to some restrictions (the n-equations must be of full rank, etc.), we can solve the nxn (5x5) KKT conditions to express each of the five *endogenous* variables as a functions of the *exogenous* parameters only. In this case, we get Marshallian demand functions: demand as a function of prices and income, and these can be substituted into the utility function to get an indirect utility function.

$$X_1 = F_1(p_1, p_2, M) \qquad X_2 = F_2(p_1, p_2, M) \quad \Rightarrow \quad U(p_1, p_2, M) \text{ Indirect utility}$$

Alternatively, we can minimize the expenditure necessary to reach a target level of utility, denoted U^*:

$$\min_{X_1, X_2} p_1 X_1 + p_2 X_2 \quad \textit{subject to} \quad U(X_1, X_2) - U^* \geq 0$$

This allows us to derive Hicksian or Hicks-compensated demand functions: demands as a function of prices and the target utility level. These can then be substituted into $(p_1 X_1 + p_2 X_2)$

to get a *cost* or *expenditure function*: the *minimum* cost at prices p_1, p_2 needed to buy U^* units of utility.

$$X_1 = H_1(p_1,p_2,U^*) \qquad X_2 = H_2(p_1,p_2,U^*) \quad => \quad C(p_1,p_2,U^*) \quad \text{cost function}$$

Our next result is usually mentioned only briefly in a standard graduate course in microeconomic theory, but it turns out to be of immense practical value in formulating computable general-equilibrium models. Once cost functions for goods in consumption or factors of production used for producing output are derived, the demands for goods and factors needed for market-clearing equations are very easily derived.

Envelop theorem => Shepard's lemma: the partial derivative of a cost function with respect to a price is the optimal (cost minimizing) level of the good or factor of production needed to reach the target level of utility or output.

Because C(..) is the minimum cost of production, the X_i are chosen optimally. Below, we will refer to an optimized function as a *value function* and give a specific example. This has the important and very useful implication that a change in price has only a first-order (partial derivative holding input quantities fixed) effect equal to the quantity of the good purchased. The total derivative allowing optimal substitution among the X_i equal to the first-order effect.

$$C_i = p_1X_1(p_1,p_2,U^*) + p_2X_2(p_1,p_2,U^*)$$

$$\left[\frac{\partial C}{\partial p_i}\right]_{\bar{X}} = X_i = \frac{dC}{dp_i} = X_i + \left[p_1\frac{\partial X_1}{\partial p_i} + p_2\frac{\partial X_2}{\partial p_i}\right] \qquad \text{where [ ] = 0}$$

$$\text{since} \quad -\frac{\partial X_2}{\partial p_i} / \frac{\partial X_1}{\partial p_i} = MRS = \frac{p_1}{p_2} \quad \text{is a condition of optimality}$$

The right-hand term in brackets is zero as a condition of optimality. This then illustrates Shepard's lemma:

$$X_i = \frac{\partial C}{\partial p_i} = H_i(p_1,p_2,U^*) \qquad = \textit{optimal} \text{ level of demand for good } X_1$$

For production of good X, where w_1 and w_2 are the prices of factors of production V_1 and V_2, and where X* is the level of output level, optimization similarly yields a cost function

$$C(w_1,w_2,X^*) \qquad V_i = \frac{\partial C}{\partial w_i} \qquad = \textit{optimal} \text{ level of demand for factor } V_1$$

As I noted above, this is of great practical value because once the cost function for a good is derived, Shepard's lemma allows an easy short-cut to deriving the demand for goods (consumption) or factors of production (production).

Homogeneity of degree 1 (constant returns to scale): if a utility function or production function is homogeneous of degree 1, then the cost (or expenditure) function is separable

$$C(p_1, p_2, U^*) = c(p_1, p_2)U^* \qquad X_i = \frac{\partial c}{\partial p_i}U \qquad \text{and c(..) is hd1}$$

$$C(w_1, w_2, X^*) = c(w_1, w_2)X^* \qquad V_i = \frac{\partial c}{\partial w_i}X \qquad \text{and c(..) is hd1}$$

where X_i and V_i are now the total demands for goods/factors as a function of prices and total utility/output.

KKT conditions, complementarity, and interpretation of slack variables: a simple example

Suppose that a competitive firm produces an output X which it can sell at a fixed price p. It has a strictly convex (increasing marginal cost) function given by $C(X) = \alpha X + \beta X^2$. The optimization problem for the firm is to maximize profits, denoted by π, restricted only by a non-negativity condition on X.

$$\text{Maximize profits } \pi = pX - (\alpha X + \beta X^2) \qquad s.t. \qquad X \geq 0$$

Using the Karush-Kuhn-Tucker theorem, there exits a positive constant μ such that the necessary conditions for the maximum are:

$$\frac{d\pi}{dX} -= p - (\alpha + 2\beta X) + \mu = 0 \qquad \mu' X = 0 \qquad \mu \geq 0, \; X \geq 0$$

where $(\alpha + 2\beta X)$ is called marginal cost in economics: the cost of producing a little more output. α is the marginal cost of the first bit of output (marginal cost at $X = 0$). An "interior" solution with X > 0 and μ = 0 will occur if p is greater than α, otherwise X = 0 (the firm does not produce) and μ will be positive:

$$X = \frac{p - \alpha}{2\beta}, \; \mu = 0, \quad if \;\; p > \alpha, \qquad X = 0, \; \mu = \alpha - p, \quad if \;\; p \leq \alpha$$

Note here that the dual or slack variable μ, when it is positive, has a simple interpretation: it is the excess of the cost of producing the first bit of output over the price earned. It tells us how unprofitable this activity or technology is. GAMS will automatically report the value of μ as part of the solution it reports, with the value of μ named the "marginal" of the variable.

Optimality (profit maximization) in a production activity requires that marginal cost (*mc*) is greater than or equal to price, complementary with the level of production. An (optimal) positive output of X means price equals marginal cost. If price is strictly greater than marginal cost in equilibrium, then output is zero.

In GAMS following a tradition in mathematics, the optimality condition is written as an inequality and specifies the complementary variable. The multiplier μ is created "behind the scenes" by GAMS. Inequalities in GAMS are always written as greater than or equal to.

$mc \geq p \quad \perp \quad X$ which is read "the weak inequality is complementary to X"

$mc = p \;\; => \;\; X > 0 \qquad mc > p \;\; => \;\; X = 0$

If the inequality is strict in equilibrium, production is unprofitable and production is zero.

Though the modeler does not specify the slack variable μ, it will appear in the GAMS output file under the name "marginal". The value of marginal equals μ. It follows from the KKT conditions that if the variable X is strictly positive in the solution, then $\mu = 0$. If $X = 0$ in the solution, then μ is strictly positive. In other words, $\mu = mc - p$, and the slack variable μ measures the amount by which marginal cost exceeds p in equilibrium. Alternatively, μ is the difference between the left and right-hand sides of the weak inequality.

Profit maximization example to illustrate value functions and Shepard's lemma

Above, we had a profit equation that gives profits as a function of output and three parameters: $\pi(X, p, \alpha, \beta)$. Taking the first-order condition and solving for the optimal X, this value of X can be substituted back into the profit equation for X to get the optimal value of profits as a function of the exogenous parameters only. This is referred to as a *value function*, or specifically here a *profit function* (as opposed to profit equation).

$$\pi(X, p, \alpha, \beta) \;\; => \;\; \text{solve for } X \text{ from first-order condition } \frac{d\pi}{dX} = 0 \;\; => \;\; \pi(p, \alpha, \beta)$$

Assume that our profit-maximization example has an interior solutions so that

$X = \dfrac{p - \alpha}{2\beta}$ when substituted back into the profit equation gives

$$\pi = (p - \alpha)X - \beta X^2 = \frac{(p - \alpha)^2}{2\beta} - \beta\left[\frac{(p - \alpha)}{2\beta}\right]^2 = \frac{(p - \alpha)^2}{4\beta}$$

gives the *optimal* (maximum) value of π at (p, α, β)

The envelop theorem and Shepard's lemma tell us that the derivative of the value (profit) function must be the optimal supply of X at (p, α, β). This can be easily verified from these last two equations.

$$\frac{d\pi}{dp} = \frac{(p - \alpha)}{2\beta} = X \qquad \text{supply function } X(p, \alpha, \beta)$$

Market-clearing and complementarity

Complementarity between a weak inequality and a variable also arises in market clearing equations. Market clearing requires that the value (price times quantity) supplied in a market equal the value of demand. In the above example, the competitive firm is treating price p as exogenous, but market clearing conditions are added weak inequalities that we should think of as associated with the demand price, now an endogenous variable. Let D denote the demand for a good while X continues to denote supply: $pX = pD$. Here are three possibilities:

$X = D > 0 \quad mc = p > 0$ usual "interior" solution: both X and $p > 0$

$X > D > 0 \quad mc = p = 0$ excess supply in equilibrium, X is a "free good"

$X = D = 0 \quad mc > p > 0$ X is too expensive to produce (covered above)

One possibility that cannot occur is:

$X < D \quad mc = p \geq 0$ *violates* optimizing behavior: excess demand will cause the price to rise

Taken together, the first two possibilities imply that we should model a market-clearing condition as a weak inequality complementary with the price of the good.

$X \geq D \quad \perp \quad p$ which is read "the weak inequality is complementary to p"

When GAMS is given this weak inequality and complementary variable, it will create equation with a slack variable analogous to the procedure for the pricing equation.

$$X - D + \mu = 0 \qquad \mu' p = 0 \qquad X, D, p, \mu, \geq 0$$

$$X = D \;\Rightarrow\; p > 0 \qquad X > D \;\Rightarrow\; p = 0$$

If the inequality is strict in equilibrium, there is excess supply in equilibrium and the price is zero: X is a free good. In that case, the auxiliary variable μ, called the "marginal" in the GAMS listing file, has a positive value equal to the supply-demand imbalance in equilibrium: $\mu = X - D$.

Summing up this and the previous section, we should see that a properly specified supply-demand market-equilibrium model should have two equations, each specifically associated with one particular variable. Market equilibrium is the solution to two complementarity conditions:

$$mc \geq p \quad \perp \quad X$$

$$X \geq D \quad \perp \quad p$$

There are three possible (qualitative) solutions: (a) an interior solution in which both X and p are strictly positive; (b) a corner solution in which X is a free good; (c) a corner solution in which X is not produced. Early in the book, I will present exactly this market-equilibrium problem and show the three outcomes in GAMS.

KKT conditions derived from a Lagrangean function, complementarity and dual variables

In economics, the KKT conditions for a constrained optimization problem are typically derived by starting with what is referred to as a Lagranean equation or function (though Lagrange was a productive guy and this can also refer to quite different things). The equation to be optimized is the objective function (e.g., utility) followed by the sum of complementary slackness conditions: the dual variable multiplied by its complementary variable. The dual variables in economics are referred to as Lagranean multipliers. Let's take a specific quadratic utility function

$$U = (\alpha_1 X_1 - \beta_1 X_1^2) + X_2$$

Economists often (wrongly) fail to bother with the non-negativity conditions on the X_i, but the proper Lagrangean equation which goes with this utility function is:

$$U_L = (\alpha_1 X_1 - \beta_1 X_1^2) + X_2 + \lambda(M - p_1 X_1 - p_2 X_2) + \mu_1 X_1 + \mu_2 X_2$$

Note especially that, because of the KKT complementary slackness conditions, we will always have $U_L = U$ at the solution. Ignore the non-negativity constraints for the moment. The KKT conditions as written in economics for X_1, X_2, and λ are given by:

$$\frac{\partial U}{\partial X_1} = \alpha_1 - 2\beta_1 X_1 - \lambda p_1 + \mu_1 = 0$$

$$\frac{\partial U}{\partial X_2} = 1 - \lambda p_2 + \mu_2 = 0$$

$$\frac{\partial U}{\partial \lambda} = M - p_1 X_1 - p_2 X_2 \geq 0$$

If X_1 and X_2 are strictly positive at the solution, then $\mu_1 = \mu_2 = 0$, and if the budget constraint is binding (all income is spent), then the three equations can be solved for the three unknowns, though economists are often not interested in the solution value of λ (but we will see many cases where it is of great interest).

In this special case, all income will always be spent so that the budget constraint holds with equality and $\lambda > 0$. With each additional unit of X_2 consumed giving a marginal utility of one, there is no reason to throw away income. But it may not be the case that both X1 and X2 are consumed; let's *assume* that this is the case so that then $\mu_1 = \mu_2 = 0$. The second first-order condition then implies that $\lambda = 1/p_2$. Substituting this into the first equation given the demand for X1. Multiply this by p1 to get p1X1, and then substitute that into the budget constraint (third equation) to solve for X2. If my algebra is right, you should get

$$X_1 = \frac{\alpha - p_1/p_2}{2\beta} \qquad p_1X_1 = p_1\frac{\alpha - p_1/p_2}{2\beta} \qquad p_2X_2 = M - p_1\frac{\alpha - p_1/p_2}{2\beta}$$

$$X_2 = \frac{M}{p_2} - \frac{p_1}{p_2}\frac{\alpha - p_1/p_2}{2\beta}$$

$$p_1X_1 + p_2X_2 = p_1\frac{\alpha - p_1/p_2}{2\beta} + M - p_1\frac{\alpha - p_1/p_2}{2\beta} = M$$

where the last line verifies that the sum of the expenditures on the two goods equals income.

So we're finished right? Well, no. Suppose that these equations tell us that the expenditure on X1 exceeds income. Or suppose the opposite and suppose that expenditure in X1 is less than zero. We will then have a corner solution and this is when the slack variables then μ_1 and μ_2 come into play. In the first special case, X2 will be zero and μ2 will be positive, while in the second special case, X1 will be zero and μ1 will be positive. Specifically, the possible "corner solutions" are

If $\quad p_1\dfrac{\alpha - p_1/p_2}{2\beta} > M \quad \Rightarrow \quad X_1 = \dfrac{M}{p_1}, \quad X_2 = 0, \quad \mu_2 > 0$

If $\quad p_1\dfrac{\alpha - p_1/p_2}{2\beta} < 0 \quad \Rightarrow \quad X_1 = 0, \quad X_2 = \dfrac{M}{p_2}, \quad \mu_1 > 0$

Or to put it the other way around, the critical condition for an interior solution is that

$$\frac{2\beta M}{p_1} > \alpha - \frac{p_1}{p_2} > 0 \qquad \Rightarrow \qquad X_1, X_2 > 0$$

I think that this example shows that problems become a more complicated when we allow for corner solutions. I also think that it is fair to say that many economists, insisting on using only analytical methods avoid corner solutions and indeed often fail to check that their proposed interior solution is in fact valid. Corner solutions are viewed as an inconvenience to be avoided. The above case is in fact widely used in industrial organization and in international trade, and is referred to as "quasi-linear preferences", where one "outside good" has a constant marginal utility (X_2 in our case) and the other good is the one of primary interest (X_1 in our case). Invariably, and interior solution is *assumed*, and parameter restrictions needed to ensure that this is true are ignored.

But corner solutions are not just an inconvenient or pathological case of no importance. In many cases, they are vital to the economics of a problem as we shall see. Here are just a couple of examples before we move on. (A) changes in prices or taxes can lead to switches in which technologies are and are not used to produce a good like electricity. (B) changes in transportation costs, tariffs, or technologies can lead to switches in which trade links are active or inactive in international trade models (e.g., whether or not a particular good is imported by country i from country j, exported, or non-traded). (C) changes in costs or technologies can lead to switches in industrial structure, such as firms choosing between exporting to a foreign country or producing abroad in that country, or licensing to a local firm. (D) individuals or firms decide whether or not to buy insurance. (E) individuals decide on what assets to hold in a portfolio and which to avoid (hold zero). All of these examples will be examined in this book.

Exercise 1: Cobb-Douglas production and cost functions

A consumer has a Cobb-Douglas utility function and a linear budget constraint, with income *M*.

$$U = \left(\frac{X_1}{\alpha}\right)^{\alpha}\left(\frac{X_2}{1-\alpha}\right)^{1-\alpha} \qquad M = p_1X_1 + p_2X_2$$

(1) show that the optimal shares of income spent on X_1 and X_2 are given by α and (1 - α) respectively; i.e.,

$$\alpha = \frac{p_1X_1}{M}$$

(2) derive the Marshallian or uncompensated demand functions for X_1 and X_2 : $X_i = D_i(p_1, p_2, M)$

(3) derive the Hicksian or compensated demand functions for X_1 and X_2 : $X_i = H_i(p_1, p_2, U)$

(4) derive the expenditure (cost) function, which gives the minimum expenditure necessary at prices p_1 and p_2 to buy one unit of utility.

$$e = e(p_1, p_2) \qquad \text{And} \qquad E = e(p_1, p_2)U$$

(5) demonstrate Shepard's lemma: the derivative of the expenditure function with respect to the price of good i gives the Hicksian demand for good i:

$$X_i = \frac{\partial e(p_1, p_2)}{\partial p_i} U = H(p_1, p_2, U)$$

(6) show that in this Cobb-Douglas case that

$$X_1 = H(p_1, p_2, U) = [\alpha e(p_1, p_2)/p_1]\, U = [\alpha p_u/p_1]\, U \qquad p_u \equiv e(p_1, p_2)$$

where p_u is the "price" (cost) of buying one unit of utility.

(7) derive the indirect utility function, which gives the (maximum) level of utility reached at prices p_1 and p_2 and income M.

$$V = V(p_1, p_2, M)$$

(8) demonstrate Roy's identity:

$$X_i = D_i(p_1, p_2, M) = -V_{p_i}/V_m$$

Hint: start with the utility function and form the Lagrangean function.

$$\max \quad U = X_1^{\alpha} X_2^{1-\alpha} + \lambda(M - p_1 X_1 - p_2 X_2)$$

Take the three first-order conditions for maximization

$$\frac{\partial U}{\partial X_1} = \alpha X_1^{\alpha-1} X_2^{1-\alpha} - \lambda p_1 = 0$$

$$\frac{\partial U}{\partial X_2} = (1-\alpha) X_1^{\alpha} X_2^{-\alpha} - \lambda p_2 = 0$$

$$\frac{\partial U}{\partial \lambda} = M - p_1 X_1 - p_2 X_2 = 0$$

Divide the second equation by the first to eliminate λ, getting an expression for X_2 in terms of prices,parameters and X_1. Use this to eliminate either X_1 or X_2 in equation (3). Rearrange to get the Marshallian demand functions, the answer to part (b). Rearrange the Marshallian demand function to show (a).

Now you can do virtually everything else by making appropriate substitutions.

Alternative: minimize the expenditure necessary to reach a certain level of U.

$$\min \quad p_1X_1 + p_2X_2 \quad + \quad \mu(U - X_1^{\alpha}X_2^{1-\alpha})$$

Take the three first-order conditions for maximization

$$\frac{\partial E}{\partial X_1} = p_1 - \mu\alpha X_1^{\alpha-1}X_2^{1-\alpha} = 0$$

$$\frac{\partial E}{\partial X_2} = p_2 - \mu(1-\alpha)X_1^{\alpha}X_2^{-\alpha} = 0$$

$$\frac{\partial E}{\partial \mu} = U - X_1^{\alpha}X_2^{1-\alpha} = 0$$

Divide the second equation by the first to eliminate μ, getting an expression for X_2 in terms of prices,parameters and X_1. Use this to eliminate either X_1 or X_2 . Rearrange to get the Hicksian demand functions, the answer to part (c). Put these in the expenditure equation to derive the expenditure function (d).

Analogy and extension to cost functions for producers with constant-returns technologies

The expenditure function derived above is just a type of cost function: it gives the minimum expenditure necessary to buy one unit of *U*. Cost functions for firms are exactly the same, with different labels on the inputs and outputs. Suppose a firm produces *X* from inputs *L* and K (labor and capital), with the constant-returns technology

$$X = \left(\frac{L}{\alpha}\right)^{\alpha}\left(\frac{K}{1-\alpha}\right)^{1-\alpha} \qquad C = p_lL + p_kK$$

where p_l and p_k are the prices of labor and capital and *C* is cost. Following exactly the same minimization procedure used above, you should be able to show that the cost function is given by:

$$C(p_l, p_k, X) = c(p_l, p_k)X = p_l^{\alpha}p_k^{1-\alpha}X$$

The amount of labor (demand for labor) needed to produce X units of output is then given by Shepard's lemma:

$$L = \frac{\partial c}{\partial p_l} X = \alpha p_l^{\alpha - 1} p_k^{1 - \alpha} X = \alpha \left[p_l^{\alpha} p_k^{1 - \alpha} / p_l \right] X = \alpha \left[c(p_l, p_k) / p_l \right] X$$

and similarly for the demand for capital. As above, α gives the share of labor in total expenditure or cost.

$$\frac{p_l L}{c(p_l, p_k) X} = \alpha$$