John Nachbar
Washington University
February 7, 2022

Competitive Consumer Demand[1]

1 Introduction.

These notes sketch out the basic elements of competitive demand theory. The main result is the Slutsky Decomposition theorem, Theorem 9, which gives an exact statement of the fact that price changes affect demand through two channels, a substitution effect and an income effect. The Slutsky Decomposition builds off results from two closely related but distinct problems: the utility maximization problem and the expenditure minimization problem. I discuss these in turn.

2 Utility Maximization.

Given a price vector $p \in \mathbb{R}^N_{++}$ and wealth $m \in \mathbb{R}_{++}$, the *budget set* is

$$B(p, m) = \{x \in \mathbb{R}^N_+ : p \cdot x \leq m\}.$$

By the *utility maximization problem*, I mean,

$$\begin{array}{cc} \max & u(x) \\ s.t. & x \in B(p, m). \end{array}$$

where $u : \mathbb{R}^N_+ \to \mathbb{R}$.

To avoid complicating the discussion, I do not state results for the weakest possible assumptions on u. Instead, I assume that u satisfies what I call the *standard assumptions*:

1. On $\mathbb{R}^N_+$, u is continuous, locally non-satiated, and strictly quasi-concave;
2. On $\mathbb{R}^N_{++}$, u is $\mathcal{C}^{r+1}$, where $r \geq 1$ is an integer, differentiably locally non-satiated ($Du(x) \neq 0$), and differentiably strictly quasi-concave.

2.1 Demand.

Solving the utility maximization problems yields the demand function ϕ, where $\phi(p, m)$ gives the demanded vector at prices p and wealth m. The following result establishes some basic properties of ϕ, beginning with its existence.

Theorem 1. *Under the standard assumptions, the following hold.*

1. *For every* $(p, m) \in \mathbb{R}^{N+1}_{++}$*, the utility maximization problem has a unique solution,* $\phi(p, m)$*.*

2. $\phi : \mathbb{R}^{N+1}_{++} \to \mathbb{R}^{N}_{+}$ *is continuous.*

3. *If* u *is* $\mathcal{C}^{r+1}$ *and* $\phi(p^*, m^*) \gg 0$ *then* ϕ *is* $\mathcal{C}^r$ *on an open set containing* (p^*, m^*)*.*

4. *For any* $(p, m) \in \mathbb{R}^{N+1}_{++}$*,* $p \cdot \phi(p, m) = m$*.* (Walras's Law.)

5. *For any* $(p, m) \in \mathbb{R}^{N+1}_{++}$ *and any* $\delta > 0$*,* $\phi(\delta p, \delta m) = \phi(p, m)$*.* (Homogeneity of degree 0 in p and m.)

6. *Consider any* $(p^*, m^*), (\hat{p}, \hat{m}) \in \mathbb{R}^{N+1}_{++}$*. Let*

$$\begin{aligned} B^* &= B(p^*, m^*), & x^* &= \phi(p^*, m^*), \\ \hat{B} &= B(\hat{p}, \hat{m}), & \hat{x} &= \phi(\hat{p}, \hat{m}). \end{aligned}$$

   (a) $x^* \in \hat{B}$ *and* $\hat{x} \in B^*$ *iff* $x^* = \hat{x}$*.* (The Weak Axiom.)

   (b) *Moreover, if* $x^* \gg 0$ *then* $x^* = \hat{x}$ *iff* $B^* = \hat{B}$*.*

Proof.

1. *Existence and uniqueness.* Since $p \gg 0$, $m > 0$, $B(p, m)$ is compact and non-empty. Since u is continuous, a solution exists. To see that the solution is unique, consider any $x, \hat{x} \in B(p, m)$. If both were solutions then $u(x) = u(\hat{x})$. I argue by contraposition. Suppose that $u(x) = u(\hat{x})$ and $x \neq \hat{x}$. For any $\theta \in (0, 1)$, let $x_\theta = \theta\hat{x} + (1 - \theta)x$. Then $x_\theta \in B(p, m)$ and by strict quasi-concavity, $u(x_\theta) > u(x)$. Hence neither x nor $\hat{x}$ is a solution for $B(p, m)$.

2. *Continuity.* Continuity of ϕ is a corollary of a general result, the Theorem of the Maximum, that establishes continuity of the solution for many parameterized optimization problems. Here is a proof tailored to this particular situation.

   Consider sequences $\{p_t\}$, $\{m_t\}$ with $p_t \gg 0$ and $m_t > 0$. Let $x_t = \phi(p_t, m_t)$. I need to show that if $p_t \to p^* \gg 0$ and $m_t \to m^* > 0$ then $x_t \to x^* = \phi(p^*, m^*)$.

   Suppose that $\{x_t\}$ converges, say to $\hat{x}$. I need to show that $\hat{x} = x^*$. By continuity of $p \cdot x$, it follows that $p^* \cdot \hat{x} \leq m^*$, hence $\hat{x} \in B(p^*, m^*)$. Since x^* is optimal for (p^*, m^*), it must be that $u(\hat{x}) \leq u(x^*)$. I claim that $u(\hat{x}) = u(x^*)$. If this is true then $\hat{x}$ is a solution at (p^*, m^*). Since the solution is unique, $\hat{x} = x^*$, as was to be shown.

   If $x^* = 0$ it follows that $x^* \in B(p_t, m_t)$ for every t, hence $u(x_t) \geq u(x^*)$ for every t. Since u is continuous, this implies $u(\hat{x}) \geq u(x^*)$. Hence $u(\hat{x}) = u(x^*)$ as was to be shown.

On the other hand, if $x^* > 0$, then for each t let

$$x_t^* = \left[\frac{p^* \cdot x^*}{m^*} \frac{m_t}{p_t \cdot x^*}\right] x^*.$$

By construction, since $p^* \cdot x^* \leq m^*$, $p_t \cdot x_t^* \leq m_t$, hence $x_t^* \in B(p_t, m_t)$, hence $u(x_t) \geq u(x_t^*)$. Taking the limit, we again have $u(\hat{x}) \geq u(x^*)$, hence $u(\hat{x}) = u(x^*)$.

This argument assumes that x_t is convergent. To see that it is, I first argue that $\{x_t\}$ must be bounded. If it were not bounded then there would be a subsequence $\{x_{t_k}\}$ with $p_{t_k} \cdot x_{t_k} \to \infty$, contradicting the fact that $p_t \cdot x_t \leq m_t \to m^* < \infty$.

Since $\{x_t\}$ is bounded, it lies in a compact set. By the above argument, every convergent subsequence of $\{x_t\}$ converges to x^*. By compactness, the sequence itself converges to x^*.

3. *Differentiability.* Let $x^* = \phi(p^*, m^*)$. From the Karush-Kuhn-Tucker theorem (KKT),

$$\nabla u(x^*) = \gamma^* p^*,$$

where $\gamma^* \geq 0$ is the KKT multiplier on the budget constraint. Since $\nabla u(x^*) \neq 0$ it follows that $\gamma^* > 0$. Therefore, again by KKT, $p^* \cdot x^* = m^*$. Define

$$f(p, m, x, \gamma) = \begin{bmatrix} \nabla u(x) - \gamma p \\ p \cdot x - m \end{bmatrix}.$$

Then $f(p^*, m^*, x^*, \gamma^*) = 0$. More generally, by KKT, if $x = \phi(p, m) \gg 0$ then there is a $\gamma > 0$ such that $f(p, m, x, \gamma) = 0$. Conversely, if $x \gg 0$, $\gamma > 0$, and $f(p, m, x, \gamma) = 0$ then the KKT conditions are satisfied. Moreover, under the standard assumptions, the KKT conditions are sufficient as well as necessary. Thus, $x = \phi(p, m) \gg 0$ iff there is a $\gamma > 0$ such that $f(p, m, x, \gamma) = 0$. This means that the equation $f(p, m, x, \gamma) = 0$ implicitly defines $\phi(p, m)$.

Since u is $\mathcal{C}^{r+1}$, f is $\mathcal{C}^r$. If $D_{x,\gamma} f(p^*, m^*, x^*, \gamma^*)$ has full rank then, by the Implicit Function Theorem, there is a $\mathcal{C}^r$ function ψ defined on an open set containing (p^*, m^*) such that for any (p, m) in this set, $f(p, m, \psi(p, m)) = 0$. Thus, $\psi(p, m)$ gives both $x = \phi(p, m)$ *and* the associated λ. The function ϕ is simply the first n coordinates of ψ. Since ψ is $\mathcal{C}^r$, so is ϕ.

It remains to show that $D_{x,\gamma} f(p^*, m^*, x^*, \gamma^*)$ has full rank. To simplify notation, let

$$A = D_{x,\lambda} f(p^*, m^*, x^*, \gamma^*) = \begin{bmatrix} D^2 u(x^*) & -p^* \\ p^{*\prime} & 0 \end{bmatrix}.$$

Because u is assumed to be differentiably strictly quasiconcave, the matrix

$$B = \begin{bmatrix} D^2 u(x^*) & \nabla u(x^*) \\ Du(x^*) & 0 \end{bmatrix}$$

has full rank, hence $|B| \neq 0$. From the KKT conditions, moreover, $\nabla u(x^*) = \gamma^* p^*$. Since the determinant is multilinear, this implies,

$$\gamma^2 |A| = -|B|.$$

Hence $|A| \neq 0$, as was to be shown.

4. *Walras's Law.* I argue by contraposition. Consider any $x \in B(p, m)$ with $p \cdot x < m$. By local nonsatiation and the continuity of $p \cdot x$ there is an $\hat{x} \in B(p, m)$ such that $u(\hat{x}) > u(x)$, which implies $x \neq \phi(p, m)$.

5. *Homogeneity.* This follows immediately from the fact that, for any $\delta > 0$, $B(\delta p, \delta m) = B(p, m)$, since $\delta p \cdot x \leq \delta m$ iff $p \cdot x \leq m$. If the constraint set does not change then the solution cannot change.

6. *Weak Axiom.*

   (a) If $x^* = \hat{x}$ then, trivially, $x^* \in \hat{B}$ and $\hat{x} \in B^*$. For the converse direction, since ϕ is derived from utility maximization it must satisfy the Weak Axiom (WA) as defined for choice functions. Explicitly, if $x^* \in \hat{B}$, then $\hat{x}$ is revealed preferred to x^*. WA then requires that x^* not be revealed strictly preferred to $\hat{x}$. Since the choice here is unique, this means that *either* $\hat{x} \notin B^*$ *or* $\hat{x} = x^*$. Therefore, if, in fact, $\hat{x} \in B^*$, then $\hat{x} = x^*$, as was to be shown. Thus, part (a) is logically equivalent to WA.

   (b) If $B^* = \hat{B}$ then, trivially, $x^* = \hat{x}$. For the converse direction, suppose $x^* \gg 0$ and consider any $(p, m) \in \mathbb{R}^{N+1}_{++}$ such that $x^* = \phi(p, m)$. By KKT, there is a $\gamma \geq 0$ such that

   $$\nabla u(x^*) = \gamma p.$$

   Since $\nabla u(x^*) \neq 0$, this implies $\gamma > 0$. Since $x^* = \phi(p^*, m^*)$, KKT implies that there is a γ^* such that

   $$\nabla u(x^*) = \gamma^* p^*,$$

   again with $\gamma^* > 0$. Combining, $p = \tilde{\gamma} p^*$, where $\tilde{\gamma} = \gamma^*/\gamma$. By Property (4) (Walras's Law), it follows that $m = \tilde{\gamma} m^*$ and hence that $B(p, m) = B(\tilde{\gamma} p^*, \tilde{\gamma} m^*) = B(p^*, m^*)$, where the last equality follows from Property 5.

■

The demand function ϕ is sometimes called *Marshallian* demand in honor of the economist Alfred Marshall. Let me make a few remarks.

On differentiability, there are two basic ways that differentiability could fail, even if u is $\mathcal{C}^r$. First, differentiability could fail at the boundary. For example, if $u(x) = \ln(x_1) + x_2$ then demand for good 2 is given by

$$\phi_2(p, m) = \begin{cases} \frac{m}{p_2} - 1 & \text{if } p_2 \leq m, \\ 0 & \text{otherwise.} \end{cases}$$

This is not differentiable with respect to either p_2 or m at any $(p^*, m^*) \in \mathbb{R}^{N+1}_{++}$ at which $p_2^* = m^*$. For example, looking at $D_m\phi_2(p^*, m^*)$ with $p_2^* = m^*$, the left-hand derivative is $1/p^* > 0$ while the right-hand derivative is 0.

The statement of Property 3 rules out this problem by the crude expedient of restricting attention to (p^*, m^*) for which $\phi(p^*, m^*) \gg 0$. Although I do not do so explicitly, it is straightforward to strengthen the assumptions on u slightly (on the boundary of $\mathbb{R}^N_+$) to get differentiability of ϕ even when $\phi_n(p^*, m^*) = 0$ for some n, provided there is an open set containing (p^*, m^*) on which the set of commodities with zero demand does not change.

The second way differentiability could fail is more subtle. Even if $\phi(p^*, m^*) \gg 0$ and even if strict quasiconcavity holds, if differentiable strict quasiconcavity fails then it is possible for the derivative of ϕ_n with respect to one of the parameters to be infinite. This problem is usually viewed as a pathology: demand with an infinite slope is the limiting case of demand that is very steep but still differentiable.

Walras's Law says all wealth is spent on consumption. This may seem odd: what about saving? The budget constraint here is most easily interpreted as being for a single-period world: there is no tomorrow and hence there is no reason to save for tomorrow. In a multi-period model, there is an analog of Walras's Law that says that all wealth in a period is spent either on consumption in that period *or* on saving for consumption in future periods (including, possibly, bequests); the alternative would be to burn the wealth, benefitting nobody. A related, somewhat subtle, point is that in many (but not all) multi-period models, the consumer's multi-period optimization problem can be formulated using a single budget constraint, reflecting the fact that in deciding about how much to save today, one is effectively deciding about how much to consume tomorrow. The single budget constraint formulation focuses directly on ultimate consumption and abstracts away from savings/borrowing. For this formulation of the budget constraint in multi-period settings, Walras's Law then says that all wealth is spent on consumption at the various dates.

On homogeneity, this is a formalization of the idea of no *money illusion*: if wealth increases (or falls) by a multiple δ but prices also change by the same multiple then there is no change in the budget set and hence there should be no change in consumption. In a more elaborate model, wealth might depend on the sale of endowment, including wages. In this case, homogeneity refers to situations in which all prices, *including* wages, increase (or fall) by the same multiple. Obviously, if wages remain constant while the price of consumption goods goes up then the demanded bundle will have to fall and the consumer will be worse off.

Finally, the Weak Axiom (WA) is equivalent to a weak version of the generalized Law of Demand. I discuss this point in more detail later in the course.

2.2 Indirect utility.

The *indirect utility function* is the value function for the utility maximization problem: when prices and income are (p^*, m^*), the indirect utility, written $v(p^*, m^*)$, is defined by

$$v(p^*, m^*) = u(\phi(p^*, m^*)) = u(x^*).$$

The special terminology and notation for this value function originated in the early history of demand theory.

Here are some basic facts about v.

Theorem 2. *Under the standard assumptions, the following hold.*

1. *$v(p, m)$ is defined for every $(p, m) \in \mathbb{R}^{N+1}_{++}$.*

2. *$v : \mathbb{R}^{N+1}_{++} \to \mathbb{R}$ is continuous.*

3. *If $\phi(p^*, m^*) \gg 0$ then v is $\mathcal{C}^r$ on an open set containing (p^*, m^*).*

4. *For any $(p, m) \in \mathbb{R}^{N+1}_{++}$ and any $\delta > 0$, $v(\delta p, \delta m) = v(p, m)$.* (Homogeneity of degree 0 in (p, m).)

5. *v is weakly decreasing in p_n for every n, strictly decreasing if $\phi_n(p, m) > 0$. v is strictly increasing in m.*

6. *v is quasiconvex.*

Proof.

1. *Existence.* This is immediate since $\phi : \mathbb{R}^{N+1}_{++} \to \mathbb{R}^N$ and $v(p, m) = u(\phi(p, m))$.

2. *Continuity.* This is immediate since both u and ϕ are continuous.

3. *Differentiability.* This is immediate since u is $\mathcal{C}^{r+1}$ and ϕ is $\mathcal{C}^r$ on an open set containing (p^*, m^*).

4. *Homogeneity.* This is immediate since ϕ is homogeneous of degree zero in (p, m).

5. *Monotonicity.* Suppose that p^* is the same as $\hat{p}$ in every coordinate except n, with $p^*_n > \hat{p}_n$. Let $x^* = \phi(p^*, m)$ and $\hat{x} = \phi(\hat{p}, m)$. Since $m = p^* \cdot x^* \geq \hat{p} \cdot x^*$, it follows that $x^* \in B(\hat{p}, m)$, which implies $u(\hat{x}) \geq u(x^*)$.

   Moreover, if $\hat{x}_n > 0$ then $p^* \cdot \hat{x} > \hat{p} \cdot \hat{x} = m$. Since $p^* \cdot x^* = m$, it follows that $\hat{x} \neq x^*$, hence $u(\hat{x}) > u(x^*)$ (since the solution to the $B(\hat{p}, m)$ utility maximization problem is unique).

Let $m^* > \hat{m}$ and let $x^* = \phi(p, m^*)$ and $\hat{x} = \phi(p, \hat{m})$. Then $\hat{x} \in B(p, m^*)$, hence $u(x^*) \geq u(\hat{x})$. Moreover, $p \cdot \hat{x} = \hat{m} < m^*$. Local nonsatiation then implies that $u(x^*) > u(\hat{x})$ (since local nonsatiation implies that, since $p \cdot \hat{x} < m^*$, there is an $x \in B(p, m^*)$ that has strictly higher utility than $\hat{x}$).

6. *Quasiconvexity.* Consider any $c \in \mathbb{R}$ and suppose that $v(p^*, m^*) \leq c$ and $v(\hat{p}, \hat{x}) \leq c$. Consider any $\theta \in (0, 1)$, let $p_\theta = \theta\hat{p} + (1-\theta)p^*$ and let $m_\theta = \theta\hat{m} + (1-\theta)m^*$. I must show that $v(p_\theta, m_\theta) \leq c$.

   Let $x_\theta = \phi(p_\theta, m_\theta)$. I claim that either $\hat{p} \cdot x_\theta \leq \hat{m}$ or $p^* \cdot x_\theta \leq m^*$. I argue by contraposition. Consider any x such that $\hat{p} \cdot x > \hat{m}$ and $p^* \cdot x > m^*$. Multiplying the first by θ and the second by $1 - \theta$ and adding yields,

   $$p_\theta \cdot x > m_\theta.$$

   Since $p_\theta \cdot x_\theta \leq m_\theta$, it follows that $x \neq x_\theta$, and the claim follows.

   Suppose then that $\hat{p} \cdot x_\theta \leq \hat{m}$. Then x_θ is feasible for the $(\hat{p}, \hat{m})$ problem, and hence $v(p_\theta, m_\theta) = u(x_\theta) \leq u(\phi(\hat{p}, \hat{m})) = v(\hat{p}, \hat{m}) \leq c$. The argument is similar if $p^* \cdot x_\theta \leq m^*$.

■

2.3 Roy's Identity: the Envelope theorem, v and ϕ.

Theorem 3 (Roy's Identity)**.** *Under the standard assumptions, if $x^* = \phi(p^*, m^*) \gg 0$,*

$$x^* = -\frac{1}{D_m v(p^*, m^*)} \nabla_p v(p^*, m^*).$$

Proof. By the Envelope theorem, letting γ^* be the KKT multiplier on the budget constraint,

$$\nabla_p v(p^*, m^*) = -\gamma^* x^*,$$
$$D_m v(p^*, m^*) = \gamma^*.$$

Dividing the first by the second yields the result. ■

The intuition is that if the price of good n goes up by Δp_n, then nominal wealth would have to go up by $\Delta p_n x_n^*$ in order for x^* to remain affordable. Since nominal wealth does not go up, an increase in p_n causes a decrease in real wealth of roughly $\Delta p_n x_n^*$. This is only an approximation, and in particular ignores the fact that since p changes, the change in relative prices will by itself cause a change in x. But the Envelope Theorem says, among other things, that the change in x from the relative price effect has negligible impact on v. The main impact is from the effective decrease in real wealth. So the change in utility from an increase p_n is approximately $-\Delta p_n x_n^*$ times the marginal utility of income. Dividing both sides by Δp_n and taking the limit as Δp_n goes to zero yields

$$D_{p_n} v(p^*, m^*) = -x_n^* \, D_m v(p^*, m^*).$$

Rearranging terms yields the result.

The significance of Roy's Identity is that it allows us to compute ϕ quickly, via simple differentiation, once we know v. In applied work, it is often convenient to work with v and skip an explicit specification of u.

In competitive production theory, Roy's Identity has an analog, not exact but similar, called *Hotelling's Lemma*: the derivative of the profit function with respect to prices, both input and output prices, gives optimal input and output levels.

3 Expenditure Minimization.

Closely related to the utility maximization problem is the expenditure minimization problem: given prices $p \in \mathbb{R}^N_{++}$ and a utility level $c \in \mathbb{R}$, find the least costly $x \in \mathbb{R}^N_+$ for which $u(x) \geq c$:

$$\begin{array}{rc} \min_x & p \cdot x \\ & u(x) \geq c \\ & x \geq 0 \end{array}$$

Note that the expenditure minimization will not have a solution, because the constraint set will be empty, if c is greater than the maximum possible value of u. On the other hand, the minimization problem is trivial if $c < u(0)$ (the solution in such cases is $x^* = 0$). Therefore, to avoid triviality, always assume that c belongs to the set A defined by

$$A = \{c \in u(\mathbb{R}^N_+) | c \geq u(0)\}$$

Since u is continuous and $\mathbb{R}^N_+$ is connected, A will be an interval, possibly unbounded. (Note that if u is monotone then $u(x) \geq u(0)$ for every $x \in \mathbb{R}^N_+$ and hence $A = u(\mathbb{R}^N_+)$. I have not assumed monotonicity, however.)

3.1 Hicksian demand.

Solving the expenditure minimization problem yields the *Hicksian demand function* h, where $h(p, c)$ is the consumption that minimizes expenditure given prices p and utility of at least c.

By the standard assumptions, I continue to mean the assumptions on u discussed when I introduced the utility maximization problem. The following theorem establishes basic properties about h.

Theorem 4. *Under the standard assumptions, the following hold.*

1. *For every* $(p, c) \in \mathbb{R}^N_{++} \times A$*, the expenditure minimization problem has a unique solution, written* $h(p, c)$*.*

2. $h : \mathbb{R}^N_{++} \times A \to \mathbb{R}^N_+$ *is continuous.*

3. *If* $h(p^*, c^*) \gg 0$ *then* h *is* $\mathcal{C}^r$ *on an open set containing* (p^*, c^*)*.*

4. *For any* $(p, c) \in \mathbb{R}^N_{++} \times A$*,* $u(h(p, c)) = c$*.*

5. *For any* $(p, c) \in \mathbb{R}^N_{++} \times A$ *and any* $\delta > 0$*,* $h(\delta p, c) = h(p, c)$*.* (Homogeneity of degree 0 in p.)

Proof. These results are close analogs of the results for ϕ and so I will be brief.

1. *Existence.* The argument is similar to that for ϕ except for one point: the constraint set $\{x \in \mathbb{R}^N_+ : u(x) \geq c,\ x \geq 0\}$ is not compact. Take any $\hat{x}$ in this set (which is not empty since $c \in A$). Let $\hat{m} = p \cdot \hat{x}$. Then, since $\hat{x}$ is feasible, any solution must lie in the intersection of the constraint set and $B(p, \hat{m})$, and this intersection is compact.

2. *Continuity.* Suppose that $p_t \to p^*$ and $c_t \to c^*$, let $x_t = h(p_t, c_t)$ and $x^* = h(p^*, c^*)$. I must show that $x_t \to x^*$.

   Suppose that $x_t \to \hat{x}$. I must show that $\hat{x} = x^*$. By continuity of u, since $u(x_t) \geq c_t$ for all t, and $c_t \to c^*$, $u(\hat{x}) \geq c^*$. So $\hat{x}$ is feasible for the (p^*, c^*) minimization problem. Since x^* is the solution to this problem, $p^* \cdot \hat{x} \geq p^* \cdot x^*$. Suppose that, in addition, $p^* \cdot x^* \geq p^* \cdot \hat{x}$, so that $p^* \cdot x^* = p^* \cdot \hat{x}$. Then $\hat{x}$ is a solution to the (p^*, c^*) problem. Since this solution is unique, $\hat{x} = x^*$, as was to be shown.

   To see that $p^* \cdot x^* \geq p^* \cdot \hat{x}$, note that, by local nonsatiation, there is a point $x_1^* \in N_1(x^*)$ such that $u(x_1^*) > u(x^*) \geq c^*$. Since $c_t \to c^*$, there is an index t_1 such that $u(x_1^*) > c_{t_1}$. Hence x_1^* is feasible for the (p_{t_1}, c_{t_1}) problem, hence $p_{t_1} \cdot x_1^* \geq p_{t_1} \cdot x_{t_1}$. Similarly, there is a point $x_2^* \in N_{1/2}(x^*)$ such that $u(x_2^*) > u(x^*)$ and an index $t_2 > t_1$ such that $u(x_2^*) > c_{t_2}$. Hence x_2^* is feasible for the (p_{t_2}, c_{t_2}) problem, hence $p_{t_2} \cdot x_2^* \geq p_{t_2} \cdot x_{t_2}$. And so on. Proceeding in this way, there is a sequence $\{x_k^*\}$ such that $x_k^* \to x^*$ and, for each k, $p_{t_k} \cdot x_k^* \geq p_{t_k} \cdot x_{t_k}$. Taking limits, $p^* \cdot x^* \geq p^* \cdot \hat{x}$, as was to be shown.

   This argument assumes that $\{x_t\}$ converges. To see that it does, first note that $\{x_t\}$ must be bounded. If not, then there is a subsequence $\{x_{t_k}\}$ such that $p_{t_k} \cdot x_{t_k} \to \infty$. To see that this is impossible, choose any $x \in \mathbb{R}^N_+$ such that $u(x) > c^*$. (If $u(x) \leq c^*$ for all x, so $c^* = \sup A$, then use the following argument with $x = x^*$.) Since $c_t \to c^*$, x is feasible for the (p_t, c_t) problem for all t sufficiently large. Hence $p_t \cdot x_t \leq p_t \cdot x$ for all t sufficiently large. Since $p_t \cdot x \to p^* \cdot x < \infty$, there is no subsequence for which $p_{t_k} \cdot x_{t_k} \to \infty$ and hence, by contraposition, $\{x_t\}$ must be bounded.

Since $\{x_t\}$ is bounded, it lies in a compact set. By the previous argument, every convergent subsequence converges to x^*. By compactness, it follows that the sequence itself converges to x^*. As discussed in the proof of continuity for ϕ, compactness is important for this last step.

3. *Differentiability.* This is very similar to the proof of differentiability for ϕ, so I will be brief. The KKT first order conditions are now

$$\begin{aligned} p^* &= \gamma^* \nabla u(x^*) \\ u(x^*) &= c^* \end{aligned}$$

Define f by

$$f(p, c, x, \gamma) = \begin{bmatrix} p - \gamma \nabla u(x) \\ c - u(x) \end{bmatrix}$$

By KKT, since the sufficiency conditions hold (the objective function is linear, hence convex, while the binding constraint is quasi-concave), if $x \gg 0$ then $x = h(p, c)$ iff there is a $\gamma > 0$ such that $f(p, c, x, \gamma) = 0$. Then the result follows from the Implicit Function theorem provided

$$D_{x,\gamma} f(p^*, c^*, x^*, \gamma^*) = \begin{bmatrix} -\gamma^* D^2 u(x^*) & -\nabla u(x^*) \\ -Du(x^*) & 0 \end{bmatrix}$$

has full rank. But this is almost immediate from differentiable strict quasi-concavity and the fact that the determinant is multilinear.

4. *Utility is exactly c.* From the constraint, $u(h(p, c)) \geq c$. If $c = u(0)$, then $h(p, 0) = 0$ and we are done. Otherwise, I argue by contraposition. Suppose $u(x) > c \geq u(0)$. Then $x > 0$. By continuity of u, there is a point $\tilde{x}$ such that $\tilde{x} < x$, hence $p \cdot \tilde{x} < p \cdot x$, but $u(\tilde{x}) > c$, hence $\tilde{x}$ is feasible for the (p, c) problem. Since $\tilde{x}$ is feasible but costs less than x, x cannot be optimal: $x \neq h(p, c)$. The result then follows by contraposition.

5. *Homogeneity.* Almost immediate since the effect of switching from p to δp is to multiply the objective function by δ, and any strictly increasing transformation of the objective function yields the same solution.

■

3.2 The Expenditure function.

The *expenditure function* is the value function of the expenditure minimization problem: if x^* solves the expenditure minimization when prices and utility are (p^*, c^*) then

$$e(p^*, c^*) = p^* \cdot x^*.$$

The following theorem establishes basic properties of e.

Theorem 5. *Under the standard assumptions, the following hold.*

1. *$e(p, m)$ is defined for every $(p, c) \in \mathbb{R}^N_{++} \times A$. Thus, there is an expenditure function $e : \mathbb{R}^N_{++} \times A \to \mathbb{R}$.*

2. *e is continuous.*

3. *If $h(p^*, c^*) \gg 0$ then e is $\mathcal{C}^r$ on an open set containing (p^*, c^*).*

4. *For any $(p, c) \in \mathbb{R}^N_{++} \times A$ and any $\delta > 0$, $e(\delta p, c) = \delta e(p, c)$.* (Homogeneity of degree 1 in p.)

5. *e is weakly increasing in each p_n, strictly so if $h_n(p^*, c^*) > 0$. e is strictly increasing in c.*

6. *e is concave in p.*

Proof.

1. *Existence.* Immediate from the existence of h.

2. *Continuity.* Immediate from the continuity of h and the fact that $e(p, c) = p \cdot h(p, c)$.

3. *Differentiability.* Immediate from the differentiability of h and the fact that $e(p, c) = p \cdot h(p, c)$.

4. *Homogeneity.* Immediate from the fact $e(\delta p, c) = \delta(p \cdot h(\delta p, c)) = \delta(p \cdot h(p, c)) = \delta e(p, c)$.

5. *Monotonicity.* Let p^* be identical to $\hat{p}$ in every coordinate except n, with $p^*_n > \hat{p}_n$. Let $x^* = h(p^*, c)$ and $\hat{x} = h(\hat{p}, c)$. Then x^* is feasible for the $(\hat{p}, c)$ problem, hence $e(\hat{p}, c) \leq \hat{p} \cdot x^* \leq p^* \cdot x^* = e(p^*, c)$.

   It remains to show that if $\hat{x}_n > 0$ then $e(p^*, c) > e(\hat{p}, c)$. I argue by contraposition. Suppose $e(p^*, c) = e(\hat{p}, c)$. Then $p^* \cdot x^* = \hat{p} \cdot \hat{x}$. From above, $\hat{p} \cdot x^* \leq p^* \cdot x^*$. Combining $\hat{p} \cdot x^* \leq \hat{p} \cdot \hat{x}$. Since x^* is feasible for the $(\hat{p}, c)$ problem, this implies that x^* is a solution to the $(\hat{p}, c)$ problem. Since the solution is unique, $x^* = \hat{x}$, which implies that $\hat{x}_n = x^*_n = 0$.

   Suppose $c^* > \hat{c}$. Let $x^* = h(p, c^*)$ and let $\hat{x} = h(p, \hat{c})$. If $u(x^*) > \hat{c}$ and $\hat{c} \geq u(0)$ then $x^* > 0$. By continuity of u, there is a point $\tilde{x}$ such that $\tilde{x} < x^*$, hence $p \cdot \tilde{x} < p \cdot x^*$, while $u(\tilde{x}) > \hat{c}$. Thus $p \cdot x^* > p \cdot \tilde{x} \geq p \cdot \hat{x}$, where the last inequality comes from the fact that $\tilde{x}$ is feasible for the $(p, \hat{c})$ problem. Hence $e(p, c^*) > e(p, \hat{c})$.

6. *Concavity.* Fix c and consider any p^* and $\hat{p}$ in $\mathbb{R}^N_{++}$. Take any $\theta \in (0,1)$ and let $p_\theta = \theta\hat{p} + (1-\theta)p^*$. Let $x_\theta = h(p_\theta, c)$. I must show that

$$e(p_\theta, c) \geq \theta e(\hat{p}, c) + (1-\theta)e(p^*, c).$$

Note that

$$e(p_\theta, c) = p_\theta \cdot x_\theta = \theta\hat{p} \cdot x_\theta + (1-\theta)p^* \cdot x_\theta.$$

The result then follows from the fact that $\hat{p} \cdot x_\theta \geq e(\hat{p}, c)$, since x_θ is feasible at utility level c but not necessarily optimal, and similarly $p^* \cdot x_\theta \geq e(p^*, c)$.

■

3.3 The Envelope theorem, e, and h.

Theorem 6. *Under the standard assumptions on u, if $x^* = h(p^*, c^*) \gg 0$ then*

$$x^* = \nabla_p e(p^*, c^*).$$

Proof. This is almost immediate from the Envelope theorem, since p does not show up in the constraints. ■

The intuition is that if the price of good n goes up by Δp_n then expenditure goes up by Δp_n times consumption of good n, namely $\Delta p_n x^*_n$, assuming no shift in consumption. Of course, expenditure minimization *will* call for a shift in consumption in response to the change in relative prices, but the Envelope theorem implies, among other things, that the consumption shift has, in and of itself, negligible impact on expenditure. So the expenditure change is approximately $\Delta p_n x^*_n$. Dividing both sides by Δp_n and taking the limit as Δp_n goes to zero gives

$$D_{p_n} e(p^*, c^*) = x^*_n.$$

In competitive production theory, this theorem has an exact analog called *Shephard's Lemma*: the gradient of the cost function equals the vector of conditional factor demands. By analogy, Theorem 6 is sometimes also called Shephard's Lemma, although there is no consensus on this.

3.4 The properties of $D_p h$.

Theorem 7 says that $D_p h$ is negative semi-definite. It cannot be negative definite because (as stated in Theorem 7, Property 1), the homogeneity of h implies that $D_p h(p^*, c^*)p^* = 0$ so that, in particular, $p^{*\prime} D_p h(p^*, c^*)p^* = 0$. This is, however, the only way that $D_p h$ fails to be negative definite. As discussed after the proof, the fact that $D_p h$ is (almost) negative definite can be interpreted as saying that h satisfies a multivariate version of the Law of Demand. In contrast, ordinary demand ϕ need not satisfy the Law of Demand.

Theorem 7. *Under the standard assumptions, if u is $\mathcal{C}^3$ and if $h(p^*, c^*) \gg 0$ then $D_p h(p^*, c^*)$ has the following properties.*

1. *$D_p h(p^*, c^*) p^* = 0$.*
2. *$D_p h(p^*, c^*)$ is symmetric.*
3. *$D_p h(p^*, c^*)$ is negative semi-definite and "almost negative definite" in the sense that, for any vector $w \in \mathbb{R}^N$, $w \neq 0$, if w is not collinear with p^* then*

$$w' D_p h(p^*, c^*) w < 0,$$

while if w is collinear with p^* then

$$w' D_p h(p^*, c^*) w = 0.$$

Proof. Since $h(p, c) = \nabla_p e(p, c)$, $D_p h(p, c) = D_p^2 e(p, c)$.

1. This is a differential version of the homogeneity of h (Property 4 of Theorem 4). $D_p h(p^* c^*) p^*$ equals the directional derivative in the direction p^*. Thus

$$D_p h_n(p^*, c^*) p^* = \lim_{t \to 0} \frac{h_n(p^* + tp^*, c^*) - h_n(p^*, c^*)}{t} = 0$$

since, by homogeneity, $h(p^* + tp^*, c^*) = h((1 + t)p^*, c^*) = h(p^*, c^*)$.

2. Symmetry follows from Young's theorem and the fact that e is $\mathcal{C}^2$ if u is $\mathcal{C}^3$.

3. Since e is concave, it follows that $D^2 e(p^*, c^*)$ is negative semidefinite, hence $D_p h(p^*, c^*)$ is negative semidefinite. From Property 1, we already know that

$$D_p h(p^*, c^*) p^* = 0,$$

which implies that if w is collinear with p^* then

$$w' D_p h(p^*, c^*) w = 0.$$

Thus, it remains to show that if w is not collinear with p^* then

$$w' D_p h(p^*, c^*) w < 0.$$

From the proof of differentiability of h in Theorem 4, applying the Chain Rule to

$$f(p, c, \psi(p, c)) = 0$$

and manipulating (see the notes on the Implicit Function Theorem),

$$D\psi(p^*, c^*) = -\left[D_{x,\gamma} f(p^*, c^*, x^*, \gamma^*)\right]^{-1} \left[D_{p,c} f(p^*, c^*, x^*, \gamma^*)\right]$$

or

$$\begin{bmatrix} D_p h(p^*,c^*) & D_c h(p^*,c^*) \\ D_p \psi_{N+1}(p^*,c^*) & D_c \psi_{N+1}(p^*,c^*) \end{bmatrix}$$
$$= - \begin{bmatrix} -\gamma D_p^2 u(p^*,c^*) & -\nabla u(x^*) \\ -Du(x^*) & 0 \end{bmatrix}^{-1} \begin{bmatrix} I_{N\times N} & 0_{N\times 1} \\ 0_{1\times N} & 1 \end{bmatrix}$$

or

$$\begin{bmatrix} D_p h(p^*,c^*) & D_c h(p^*,c^*) \\ D_p \psi_{N+1}(p^*,c^*) & D_c \psi_{N+1}(p^*,c^*) \end{bmatrix} = \begin{bmatrix} \gamma D_p^2 u(p^*,c^*) & \nabla u(x^*) \\ Du(x^*) & 0 \end{bmatrix}^{-1}.$$

To simplify notation, let

$$A = \begin{bmatrix} D_p h(p^*,c^*) & D_c h(p^*,c^*) \\ D_p \psi_{N+1}(p^*,c^*) & D_c \psi_{N+1}(p^*,c^*) \end{bmatrix}$$

and let

$$B = \begin{bmatrix} \gamma D_p^2 u(p^*,c^*) & \nabla u(x^*) \\ Du(x^*) & 0 \end{bmatrix}.$$

For any $w \in \mathbb{R}^N$,

$$(w,0)'A(w,0) = w'D_p h(p^*,c^*)w.$$

Therefore, $w'D_p h(p^*,c^*)w < 0$ iff

$$(w,0)'B^{-1}(w,0) < 0.$$

Let T_{p^*} denote the set of points b in $\mathbb{R}^N$ such that $p^* \cdot b = 0$. T_{p^*} is an $N-1$ dimensional linear subspace (hyperplane through the origin). Note that by KKT, $p^* = \gamma \nabla u(x^*)$, hence p^* and $\nabla u(x^*)$ are collinear. Hence $p^* \cdot b = 0$ iff $\nabla u(x^*) \cdot b = 0$. This implies that for any $b \in T_{p^*}$, $B(b,0) = (w,0)$ for some $w \in \mathbb{R}^N$. Since B has full rank and T_{p^*} is $N-1$ dimensional, the set $S = \{w \in \mathbb{R}^N : \exists b \in T_{p^*} \text{ such that } w = Bb\}$ is likewise $N-1$ dimensional.

For any $w \in S$, let $b = B^{-1}w$. Then, since B^{-1} is symmetric,

$$\begin{aligned} (w,0)'B^{-1}BB^{-1}(w,0) &= \left[B^{-1}(w,0)\right]' B \left[B^{-1}(w,0)\right] \\ &= (b,0)'B(b,0) \\ &= \gamma b' D_p^2 u(p^*,c^*) b \\ &< 0, \end{aligned}$$

where the inequality comes from the fact that b is orthogonal to $\nabla u(x^*)$ and u is differentiably strictly quasi-concave. Hence

$$\begin{aligned} (w,0)'B^{-1}(w,0) &= (w,0)'B^{-1}\left[BB^{-1}\right](w,0) \\ &= (w,0)'B^{-1}BB^{-1}(w,0) \\ &< 0, \end{aligned}$$

where the inequality is form the previous argument. Thus, for any $w \in S$,

$$w' D_p h(p^*, c^*) w < 0.$$

It follows that $p^* \notin S$.

To complete the proof, I need to show that $w' D_p h(p^*, c^*) w < 0$ for all w not collinear with p^*, and not just $w \in S$, $w \neq 0$. This is almost immediate. Since S is an $N-1$ dimensional subspace and since $p^* \notin S$, any $w \in \mathbb{R}^N$ can be written as a linear combination of p^* and some $\hat{w} \in S$. Suppose $w = a_1 p^* + a_2 \hat{w}$. Then if w is not collinear with p^*, $w \neq 0$, hence $a_2 \neq 0$,

$$\begin{aligned} w' D_p h(p^*, c^*) w &= (a_1 p^* + a_2 \hat{w})' D_p h(p^*, c^*)(a_1 p^* + a_2 \hat{w}) \\ &= a_2^2 \hat{w}' D_p h(p^*, c^*) \hat{w} \\ &< 0. \end{aligned}$$

This proves the result.

■

As noted above, Property 3 can be viewed as stating that h must satisfy a multivariate version of the Law of Demand. In somewhat more detail, if the vector w gives a price change then $D_p h(p^*, c^*) w$ gives a linear approximation to the resulting vector change in quantities. The statement

$$w \cdot [D_p h(p^*, c^*) w] = w' D_p h(p^*, c^*) w \leq 0$$

then says that price changes are negatively related to quantity changes, Geometrically, the two change vectors must be at least 90 degrees apart. If w is zero except in component n (meaning that only the price of good n changes) then w is not collinear with p^* (since $p^* \gg 0$) and so

$$0 > w' D_p h(p^*, c^*) w = w_n^2 \frac{\partial h}{\partial p_n}(p^*, c^*)$$

which is to say that

$$\frac{\partial h}{\partial p_n}(p^*, c^*) < 0.$$

That is, if the price of good n goes up then the Hicksian demand for good n must fall. Thus the multivariate Law of Demand implies the more familiar own-price Law of Demand as a special case.

4 Duality.

It should be intuitive that the utility maximization problems and the expenditure minimization problems are closely related. Formally, the problems are "dual." The connection between solutions to the two problems is recorded in the following theorem.

Theorem 8. *Under the standard assumptions, the following hold.*

1. *For any* $(p, m) \in \mathbb{R}^{N+1}_{++}$
   (a) $\phi(p, m) = h(p, v(p, m))$.
   (b) $e(p, v(p, m)) = m$.
2. *For any* $(p, c) \in \mathbb{R}^{N}_{++} \times A$,
   (a) $h(p, c) = \phi(p, e(p, c))$.
   (b) $v(p, e(p, c)) = c$.

Proof.

1. Let $c = v(p, m)$ and let $x = \phi(p, m)$ and $\hat{x} = h(p, c)$.
   (a) Since x is feasible in the minimization problem, $p \cdot x \geq p \cdot \hat{x}$. Hence $\hat{x}$ is feasible in the maximization problem. Since $u(\hat{x}) \geq c = u(x)$, $\hat{x}$ is a solution in the maximization problem, hence (since the solution is unique), $\hat{x} = x$.
   (b) By Property 4 in Theorem 1 (Walras's Law), $p \cdot x = m$. Thus $e(p, v(p, m)) = p \cdot \hat{x} = p \cdot x = m$.
2. Let $m = e(p, c)$ and let $x = \phi(p, m)$ and $\hat{x} = h(p, c)$.
   (a) Since $\hat{x}$ is feasible in the maximization problem, $u(\hat{x}) \leq u(x)$. Hence x is feasible in the minimization problem. Since $p \cdot x \leq m = p \cdot \hat{x}$, x is actually a solution in the minimization problem, hence (since the solution is unique) $x = \hat{x}$.
   (b) By Property 4 in Theorem 4, $u(h(p, c)) = c$. Thus $v(p, e(p, c)) = u(x) = u(\hat{x}) = c$.

■

One practical implication of this is that it is possible to invert e to get v and vice versa. Thus, if you are given, say e, then you can compute h via Theorem 6 and also invert $e(p, c) = m$, setting $c = v(p, m)$, to get the indirect utility function v. You can then get ϕ from h via Theorem 8, Property 1a, or you can get ϕ from v via Roy's Identity, Theorem 3.

5 The Slutsky decomposition.

Theorem 9. *Under the standard assumptions, if $x^* = \phi(p^*, m^*) \gg 0$, then letting $c^* = v(p^*, m^*)$,*

$$D_p\phi(p^*, m^*) = D_p h(p^*, c^*) - D_m\phi(p^*, m^*)x^{*\prime}.$$

Proof. From Theorem 8, $h(p, c) = \phi(p, e(p, c))$ for any p in an open set containing p^*. The remainder of the argument is via the Chain Rule, which, strictly speaking, requires a explicit composite function. Define functions s and g by $s(p) = (p, e(p, c^*))$ and $g(p) = \phi(s(p))$. Then $h(p, c^*) = g(p)$, hence

$$D_p h(p^*, c^*) = Dg(p^*).$$

On the other hand, by the Chain Rule,

$$\begin{aligned} Dg(p^*) &= D\phi(p^*, m^*)Ds(p^*) \\ &= \begin{bmatrix} D_p\phi(p^*, m^*) & D_m\phi(p^*, m^*) \end{bmatrix} \begin{bmatrix} I_{N\times N} \\ D_p e(p^*, c^*) \end{bmatrix} \\ &= D_p\phi(p^*, m^*) + D_m\phi(p^*, m^*)D_p e(p^*, c^*) \\ &= D_p\phi(p^*, m^*) + D_m\phi(p^*, m^*)x^{*\prime} \end{aligned}$$

where the last equality comes from the fact that, by Theorem 6 and Theorem 8, $D_p e(p^*, c^*) = h(p^*, c^*)' = \phi(p^*, m^*)' = x^{*\prime}$. Thus

$$D_p h(p^*, c^*) = D_p\phi(p^*, m^*) + D_m\phi(p^*, m^*)x^{*\prime}.$$

Rearranging this yields the result. ■

The intuition is that if the price of good n goes up then there are two effects on the demand for good ℓ. One effect is the *substitution* effect. If the price of good n goes up by Δp_n then there is a change in relative prices causing a change in demand for good ℓ of approximately

$$D_{p_n} h_\ell(p^*, c^*)\Delta p_n.$$

This captures movement along the c^* indifference curve. The other effect is the *income effect* (or *weatlh effect*). If the price of good n goes up by Δp_n, then it would take an increase of wealth of $\Delta p_n x_n^*$ to allow x^* to remain affordable. Since nominal wealth does not, in fact, rise, real wealth effectively falls by $\Delta p_n x_n^*$, forcing movement to a new, lower indifference curve. The demand for good ℓ falls by approximately the change in real wealth times the marginal propensity to consume good ℓ out of wealth:

$$D_m\phi_\ell(p^*, m^*)x_n^*\Delta p_n$$

Combining all this, dividing by Δp_n and letting Δp_n go to zero yields

$$D_{p_n}\phi_\ell(p^*, m^*) = D_{p_n}h_\ell(p^*, c^*) - D_m\phi_\ell(p^*, m^*)x_n^*.$$

The expressions

$$D_p\phi(p^*, m^*) = D_p h(p^*, c^*) - D_m\phi(p^*, m^*)x^{*\prime}.$$

and

$$D_p h(p^*, c^*) = D_p\phi(p^*, m^*) + D_m\phi(p^*, m^*)x^{*\prime}.$$

are both referred to as the Slutsky Decomposition. They have different uses.

The form

$$D_p\phi(p^*, m^*) = D_p h(p^*, c^*) - D_m\phi(p^*, m^*)x^{*\prime}$$

encodes all that we know about how price changes affect demand. In particular, all information we have about the Law of Demand for ϕ is present in this expression.

The form

$$D_p h(p^*, c^*) = D_p\phi(p^*, m^*) + D_m\phi(p^*, m^*)x^{*\prime}.$$

allows us to write an unobservable, $D_p h$, in terms of objects that are observable in principle. This gives us a way to test demand theory, in principle. In particular, from Theorem 7,

$$D_p\phi(p^*, m^*) + D_m\phi(p^*, m^*)x^{*\prime}$$

is symmetric, negative semi-definite, and "almost negative definite."