The Envelope Theorem

In an optimization problem we often want to know how the value of the objective function will change if one or more of the parameter values changes. Let's consider a simple example: a price-taking firm choosing the profit-maximizing amount to produce of a single product. If the market price of its product changes, how much will the firm's profit increase or decrease? At first, the answer seems obvious: if the price change is $\Delta p > 0$ — an increase — then profit will increase by Δp on each unit the firm is producing, *i.e.*, by $(\Delta p)q$ if the firm is producing q units.

But on second thought, this doesn't seem quite right. When the price increases, the firm will probably change the amount q that it produces, so it seems as if we can't just apply the price change to the q units it *was* producing, but we apparently need to also take account of the change in q as well. Let's check this out.

The firm's problem is

$$(\mathrm{P}_1) \qquad \max_{q \in \mathbb{R}_+} \pi(q; p) = pq - C(q),$$

where $C(q)$ is the firm's cost function, which we'll assume is strictly convex and differentiable. For the firm, this is a problem with just one decision variable, q, and one parameter, p. The first-order condition for (P_1) is

$$\frac{\partial \pi}{\partial q}(q) = 0 \qquad i.e., \;\; C'(q) = p$$

— the firm chooses the output level q at which its marginal cost is equal to the market price (which the firm takes as given: we assumed it's a price-taking firm).

Let's write the solution function for (P_1) as $\hat{q}(p)$; then the value function for (P_1) is

$$v(p) = \pi\big(\hat{q}(p), p\big) = p\hat{q}(p) - C\big(\hat{q}(p)\big),$$

and we have

$$v'(p) = \frac{\partial \pi}{\partial p} + \frac{\partial \pi}{\partial q}\frac{\partial \hat{q}}{\partial p} = q + \frac{\partial \pi}{\partial q}\frac{\partial \hat{q}}{\partial p},$$

which *does* take account of the firm's output response to a price change, $\dfrac{\partial \hat{q}}{\partial p}$. But this response is multiplied by the term $\dfrac{\partial \pi}{\partial q}$, the change in profit that results from the change in q. And at the optimal q, that's zero: the first-order condition is $\dfrac{\partial \pi}{\partial q} = 0$. (To put it another way, since $MC = p$, an increase in q will increase cost and revenue by approximately the same amount, leaving profit virtually unchanged.) So it turns out that we do have $v'(p) = q$ after all, just as we first conjectured.

We've just proved the Envelope Theorem — for an optimization problem with one decision variable, one parameter, and no constraints:

The Envelope Theorem: For the maximization problem $\max_{x\in\mathbb{R}} f(x;\theta)$, if $f:\mathbb{R}^2 \to \mathbb{R}$ is a C^1-function and if the solution function for the maximization problem is differentiable at $\overline{\theta}$, then the derivative of the value function satisfies $v'(\theta) = \dfrac{\partial f}{\partial \theta}$ at $\overline{\theta}$. ∎

Corollary: If the solution function is differentiable on an open subset $U \subseteq \mathbb{R}$, then $v'(\theta) = \dfrac{\partial f}{\partial \theta}$ on U. ∎

Example 1: Let's consider the following maximization problem:

$$\begin{aligned}\max_{x\in\mathbb{R}} f(x,\theta) &= \theta - (x-\theta)^2, \quad \theta\in\mathbb{R}\\ &= \theta - x^2 + 2\theta x - \theta^2\\ &= 2\theta x - x^2 + \theta - \theta^2.\end{aligned}$$

Before doing the mathematics, let's place the maximization problem in an economic context. Suppose θ is the amount of capital a firm is using, and x is the firm's level of output. Assume that $f(x,\theta)$ is the firm's profit when it produces x units using θ units of capital. If the level of capital is taken as given (a parameter) by the firm — for example, if the firm is making a short-run output decision and it's unable to vary its capital in the short run — then the solution function $\hat{x}(\cdot)$ for this problem tells us the profit-maximizing output level, $\hat{x}(\theta)$, for the given amount of capital, θ; and the value function, $v(\cdot)$, gives us the maximum profit that can be achieved, $v(\theta)$, with the amount θ of capital. In other words, $\hat{x}(\theta)$ and $v(\theta)$ are the short-run optimal output and the short-run maximum profit.

How much can the firm increase its profit when, in the longer run, it can vary θ? The Envelope Theorem will tell us the answer. First note that $\dfrac{\partial f}{\partial \theta} = 2x + 1 - 2\theta$. Therefore, according to the Envelope Theorem, we have $v'(\theta) = \dfrac{\partial f}{\partial \theta} = 2x + 1 - 2\theta$. This expression depends on x as well as on θ; in order to determine $v'(\theta)$ from θ alone, we need to determine how x depends on θ — we need to determine the solution function. The first-order condition for the maximization problem is

$$\begin{aligned}\frac{\partial f}{\partial x} &= 2\theta - 2x = 2(\theta - x)\\ &= 0 \;\text{ at } x = \theta.\end{aligned}$$

Therefore the solution function is $\hat{x}(\theta) = \theta$, and therefore $v'(\theta) = 2\theta + 1 - 2\theta = 1$.

For this problem, of course, it's simple to obtain the value function directly:

$$v(\theta) = f\big(\hat{x}(\theta),\theta\big) = f(\theta,\theta) = \theta - (\theta-\theta)^2 = \theta.$$

Therefore, again, $v'(\theta) = 1$. We'll return to this example below.

Here's the Envelope Theorem for n variables and m paramters:

The Envelope Theorem: Assume that $f : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}$ is a C^1-function, and consider the maximization problem $\max_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{x}; \theta)$. If the solution function $\hat{\mathbf{x}} : \mathbb{R}^m \to \mathbb{R}^n$ is differentiable on an open set $U \subseteq \mathbb{R}^m$, then the partial derivatives of the value function satisfy

$$\frac{\partial v}{\partial \theta_i} = \frac{\partial f}{\partial \theta_i} \text{ on } U, \quad i = 1, \dots, m.$$

Proof: At any $\theta \in U$, for each $i = 1, \dots, m$ we have

$$\begin{aligned}
\frac{\partial v}{\partial \theta_i} &= \frac{\partial}{\partial \theta_i} f\big(\hat{\mathbf{x}}(\theta), \theta\big) \\
&= \frac{\partial f}{\partial \theta_i} + \sum_{j=1}^{n} \frac{\partial f}{\partial x_j} \frac{\partial \hat{x}_j}{\partial \theta_i} \\
&= \frac{\partial f}{\partial \theta_i}, \quad \text{because } \frac{\partial f}{\partial x_j} = 0 \text{ at } \hat{\mathbf{x}}(\theta),\ j = 1, \dots, n, \quad \text{according to the FOC.} \quad \blacksquare
\end{aligned}$$

The Envelope Theorem for Constrained Optimization

Now let's add a constraint to the maximization problem:

$$(\mathrm{P}_2) \qquad \max_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{x}; \theta) \text{ subject to } G(\mathbf{x}; \theta) \leqq 0,$$

where $f : \mathbb{R}^n \times \mathbb{R} \to \mathbb{R}$ and $G : \mathbb{R}^n \times \mathbb{R} \to \mathbb{R}$ are C^1-functions on an open set $U \subseteq \mathbb{R}^n \times \mathbb{R}$, and we assume that the solution function $\hat{\mathbf{x}}(\theta)$ exists and is differentiable for all $(\mathbf{x}, \theta) \in U$.

The first-order condition for (P_2) is that for some $\lambda \geqq 0$,

$$\nabla f(\mathbf{x}; \theta) = \lambda \nabla G(\mathbf{x}; \theta) \qquad i.e., \ \frac{\partial f}{\partial x_j}(\mathbf{x}; \theta) = \lambda \frac{\partial G}{\partial x_j}(\mathbf{x}; \theta),\ j = 1, \dots, n.$$

Let's assume that $\lambda > 0$; otherwise either the constraint is not binding or $\nabla f(\mathbf{x}; \theta) = \mathbf{0}$, so we'd be back in the unconstrained case.

Because the constraint is binding, the solution function $\hat{\mathbf{x}}(\cdot)$ satisfies $G\big(\hat{\mathbf{x}}(\theta), \theta\big) = 0$ for all θ such that $\big((\hat{\mathbf{x}}(\theta), \theta\big) \in U$. We therefore have the following proposition, which will be instrumental in the proof of the Envelope Theorem.

Proposition: If $G : \mathbb{R}^n \times \mathbb{R} \to \mathbb{R}$ is a C^1-function and $\hat{\mathbf{x}}(\cdot) : \mathbb{R} \to \mathbb{R}^n$ satisfies $G\big(\hat{\mathbf{x}}(\theta), \theta\big) = 0$ for all $(\mathbf{x}, \theta) \in U$, then

$$\sum_{j=1}^{n} \frac{\partial G}{\partial x_j} \frac{\partial \hat{x}_j}{\partial \theta} = -\frac{\partial G}{\partial \theta}.$$

Proof: The 1st-degree Taylor polynomial equation for $\Delta G(\Delta \mathbf{x}, \Delta \theta) = 0$ is

$$\frac{\partial G}{\partial x_1} \Delta x_1 + \cdots + \frac{\partial G}{\partial x_n} \Delta x_n + \frac{\partial G}{\partial \theta} \Delta \theta + R(\Delta \mathbf{x}, \Delta \theta) = 0,$$

where $\displaystyle\lim_{(\Delta \mathbf{x}, \Delta \theta) \to \mathbf{0}} \frac{1}{\|(\Delta \mathbf{x}, \Delta \theta)\|} R(\Delta \mathbf{x}, \Delta \theta) = 0$. Therefore

$$\begin{aligned}
& \lim_{\Delta \theta \to 0} \left[\sum_{j=1}^{n} \frac{\partial G}{\partial x_j} \frac{\Delta x_j}{\Delta \theta} + \frac{\partial G}{\partial \theta} \right] &= 0 \\
\textit{i.e.,} \quad & \sum_{j=1}^{n} \frac{\partial G}{\partial x_j} \lim_{\Delta \theta \to 0} \frac{\Delta x_j}{\Delta \theta} &= -\frac{\partial G}{\partial \theta} \\
\textit{i.e.,} \quad & \sum_{j=1}^{n} \frac{\partial G}{\partial x_j} \frac{\partial \hat{x}_j}{\partial \theta} &= -\frac{\partial G}{\partial \theta}. \quad \blacksquare
\end{aligned}$$

Now the proof of the Envelope Theorem, for a one-constraint maximization problem, is straightforward:

The Envelope Theorem: Assume that $f : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}$ and $G : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}$ are C^1-functions, and consider the maximization problem

$$(\mathrm{P}_2) \qquad \max_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{x}; \theta) \ \text{ subject to } \ G(\mathbf{x}, \theta) \leqq 0.$$

Assume that the solution function $\hat{\mathbf{x}}(\cdot) : \mathbb{R}^m \to \mathbb{R}^n$ is differentiable at $\overline{\theta} \in \mathbb{R}^m$ and that λ is the value of the Lagrange multiplier in the first-order condition for (P_2) at $\big(\hat{\mathbf{x}}(\overline{\theta}), \overline{\theta}\big)$. Then the partial derivatives of the value function satisfy

$$\frac{\partial v}{\partial \theta_i} = \frac{\partial f}{\partial \theta_i} - \lambda \frac{\partial G}{\partial \theta_i} \ \text{ at } \overline{\theta}, \quad i = 1, \ldots, m,$$

where the partial derivatives of f and G are evaluated at $\big(\hat{\mathbf{x}}(\overline{\theta}), \overline{\theta}\big)$.

Proof: For each $i = 1, \ldots, m$, and evaluating all partial derivatives of f and G at $\big(\hat{\mathbf{x}}(\overline{\theta}), \overline{\theta}\big)$, we have

$$\begin{aligned}
\frac{\partial v}{\partial \theta_i}(\overline{\theta}) &= \frac{\partial f}{\partial \theta_i} + \sum_{j=1}^{n} \frac{\partial f}{\partial x_j} \frac{\partial \hat{x}_j}{\partial \theta}(\overline{\theta}) \\
&= \frac{\partial f}{\partial \theta_i} + \sum_{j=1}^{n} \lambda \frac{\partial G}{\partial x_j} \frac{\partial \hat{x}_j}{\partial \theta}(\overline{\theta}) \ \text{ from the FOC for } (\mathrm{P}_2) \\
&= \frac{\partial f}{\partial \theta_i} - \lambda \frac{\partial G}{\partial \theta_i}, \ \text{ by the proposition above.} \quad \blacksquare
\end{aligned}$$

Now it's an easy exercise to extend the proof to maximization problems with multiple constraints. So we have the general version of the Envelope Theorem:

The Envelope Theorem: Assume that $f : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}$ and $G : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}^\ell$ are C^1-functions, and consider the maximization problem

$$\text{(P)} \qquad \max_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{x};\theta) \;\; \text{subject to} \;\; G(\mathbf{x},\theta) \leqq \mathbf{0},$$

i.e.,

$$\text{(P)} \qquad \max_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{x};\theta) \;\; \text{subject to} \;\; G_1(\mathbf{x},\theta) \leqq 0, \ldots, G_\ell(\mathbf{x},\theta) \leqq 0.$$

Assume that the solution function $\hat{\mathbf{x}}(\cdot) : \mathbb{R}^m \to \mathbb{R}^n$ is differentiable at $\overline{\theta} \in \mathbb{R}^m$ and that for each $k = 1, \ldots, \ell$, λ_k is the value of the Lagrange multiplier for constraint $G_k(\mathbf{x},\theta) \leqq 0$ in the first-order condition for (P) at $\big(\hat{\mathbf{x}}(\overline{\theta}), \overline{\theta}\big)$. Then the partial derivatives of the value function satisfy

$$\frac{\partial v}{\partial \theta_i} = \frac{\partial f}{\partial \theta_i} - \sum_{k=1}^{\ell} \lambda_k \frac{\partial G_k}{\partial \theta_i} \;\; \text{at } \overline{\theta}, \;\; i = 1, \ldots, m,$$

where the partial derivatives of f and G are evaluated at $\big(\hat{\mathbf{x}}(\overline{\theta}), \overline{\theta}\big)$. ■

Example 1, continued:

Let's take a look at Figure 1, where we have the graphs of the functions $f(\cdot,\theta)$ in Example 1 for several values of θ, and also the graph of the value function $v(\theta)$. Note that the graph of $v(\theta)$ is *not* the envelope of the various graphs of the functions $f(\cdot,\theta)$. In fact, it can't be: the $v(\theta)$ graph is the locus of the "peaks" of all the $f(\cdot,\theta)$ graphs. So it's only when $v(\theta)$ is a constant function that it will be the envelope of the $f(\cdot,\theta)$ graphs, a pretty uninteresting case.

Figure 2 depicts the graph that *is* the envelope of the graphs of $f(x,\theta)$ for various values of θ. It's the graph of a function we'll denote by $\tilde{v}$; the function is $\tilde{v}(x) = x + \frac{1}{4}$.

Now let's change the question we ask about this firm. Let's ask, for any output level x, what's the optimal amount of capital with which to produce x units? That is, what amount of capital, θ, will enable the firm to generate the most profit *by producing the output amount* x? This is a new maximization problem,

$$(\widetilde{\text{P}}) \qquad \max_{\theta \in \mathbb{R}} f(x,\theta) = 2\theta x - x^2 + \theta - \theta^2 \;\; \text{for a given value of } x.$$

Exercise: Show that the solution function for the problem $(\widetilde{\text{P}})$ is $\tilde{\theta}(x) = x + \frac{1}{2}$ and the value function is $\tilde{v}(x) = x + \frac{1}{4}$.

Now look at Figure 2 again. At $x = \frac{1}{2}$, for example, the maximum possible profit is indeed $\tilde{v}(x) = x + \frac{1}{4} = \frac{3}{4}$. And the maximum profit when $x = \frac{1}{2}$ is achieved on the graph of the function $f(x, 1)$ — *i.e.*, the capital required in order to achieve maximum profit by producing $x = \frac{1}{2}$ is $\theta = 1 = x + \frac{1}{2} = \tilde{\theta}(x)$ for $x = \frac{1}{2}$. Moreover, the point $(x, \tilde{v}(x))$ is a point of tangency between the graph of $f(x, 1)$ and the graph of $v(\theta)$ — not $\tilde{v}(x)$. The same thing is true at other values of x: the maximum profit attainable by producing x units is with capital amount $x + \frac{1}{2}$ and this occurs at a tangency between the graph of $f(x, \theta)$ for $\theta = \tilde{\theta}(x)$ and the graph of $v(\theta)$.

So what we find is that, while the graph of $v(\theta)$, the value function for the problem (P), is not the envelope of the graphs of the functions being maximized for the various values of θ, it *is* the envelope of the graphs of the functions being maximized in the problem $(\widetilde{\mathrm{P}})$. However, this is a different maximization problem — it's the *reverse* problem, so to speak, in which we've switched the decision variable and the parameter. So it's really easy to be misled by the term "Envelope Theorem." The theorem is about the value function of a maximization (or minimization) problem where the solution function is, say, $\hat{x}(\theta)$ for a parameter θ, but the value function $v(\theta)$ for this problem is actually the envelope of the functions $f(x, \theta)$ in the reverse problem, where the solution function is $\hat{\theta}(x)$.

The Lagrange Multiplier as Shadow Value

Suppose the parameter in our maximization problem is the right-hand side of a constraint:

$$\max_{x \in \mathbb{R}^n} f(\mathbf{x}) \text{ subject to } g(\mathbf{x}) \leqq b.$$

Rewrite the problem as

$$\max f(\mathbf{x}) \text{ subject to } G(\mathbf{x}; b) \leqq 0, \quad \text{where } G(\mathbf{x}; b) = g(\mathbf{x}) - b.$$

The Envelope Theorem tells us that the derivative of the value function is

$$v'(b) = \frac{\partial f}{\partial b} - \lambda \frac{\partial G}{\partial b} = 0 - (\lambda)(-1) = \lambda.$$

Therefore the value of the Lagrange multiplier at a solution of the constrained maximization problem is the amount by which the objective value will be increased as a result of increasing the RHS of the constraint by one unit (i.e., "relaxing" the constraint by a unit). And, of course, it's also the amount by which the objective value will be decreased if we decrease the RHS by a unit. So λ is the marginal value, or "shadow value", of a change in the RHS of the constraint. (It's important to note that the value of λ depends on b.)

Example 2: Hotelling's Lemma for a Profit-Maximizing Firm

A firm produces a single product and is a price-taker in both its product and input markets. If the firm uses only two inputs, then its profit-maximization problem is

$$\max_{q,x_1,x_2} \pi(q,x_1,x_2;p,w_1,w_2) = pq - w_1x_1 - w_2x_2 \text{ subject to } q \leqq f(x_1,x_2),$$

where q is the amount of the product produced via the strictly concave production function $f(x_1,x_2)$ and x_1 and x_2 are the amounts of the two inputs used. The market prices are p per unit for the product and w_1 and w_2 for the inputs. Let $\theta = (p,w_1,w_2)$ and let's rewrite the constraint as $G(q,x_1,x_2) \leqq 0$, where $G(q,x_1,x_2) = q - f(x_1,x_2)$.

The value function $v(\theta)$ for the firm's maximization problem tells us the firm's profit as a function of the parameters (the prices) p, w_1, and w_2. The Envelope Theorem tells us immediately that

$$\begin{aligned}
\frac{\partial v}{\partial p} &= \frac{\partial \pi}{\partial p} - \lambda\frac{\partial G}{\partial p} = \frac{\partial \pi}{\partial p} = q = \hat{q}(p,w_1,w_2) \\
\frac{\partial v}{\partial w_1} &= \frac{\partial \pi}{\partial w_1} - \lambda\frac{\partial G}{\partial w_1} = \frac{\partial \pi}{\partial w_1} = -x_1 = -\hat{x}_1(p,w_1,w_2) \\
\frac{\partial v}{\partial w_2} &= \frac{\partial \pi}{\partial w_2} - \lambda\frac{\partial G}{\partial w_2} = \frac{\partial \pi}{\partial w_2} = -x_2 = -\hat{x}_2(p,w_1,w_2).
\end{aligned}$$

If there are n inputs, we have $\dfrac{\partial v}{\partial w_i} = -x_i = -\hat{x}_i(p,w_1,w_2)$ for each $i = 1,\ldots,n$.

This result is **Hotelling's Lemma**. Note that if we can observe the amounts $q, x_1, \ldots, x_n$, then we don't need to obtain the solution function, or even obtain the first-order conditions, to use this result. However, if we don't know these amounts, then we would have to obtain the solution function that gives us $\hat{q}(p,\mathbf{w})$ and/or $\hat{x}_i(p,\mathbf{w})$, depending on which derivatives of $v(\cdot)$ we're trying to obtain.

Example 3: Shephard's Lemma for a (Conditional) Cost-Minimizing Firm

A firm produces a single product and is a price-taker in its input markets. At whatever level of output it produces, we assume that it minimizes the cost of doing so. The optimization problem is

$$\min_{x_1,x_2} E(x_1,x_2;w_1,w_2) = w_1x_1 + w_2x_2 \text{ subject to } f(x_1,x_2) \geqq q,$$

where x_1 and x_2 are the amounts of the two inputs used, and q is the amount of the product to be produced via the production function $f(x_1,x_2)$. The market prices of the inputs are w_1 and w_2 per unit. The firm need not be a profit maximizer.

The Solution Function: $\hat{\mathbf{x}}(w_1,w_2,q)$ — *i.e.*, $\hat{x}_1(w_1,w_2,q)$ and $\hat{x}_2(w_1,w_2,q)$ — the firm's **conditional factor demand function(s)**, which we obtain from the first-order conditions for the minimization problem.

The Value Function: $v(w_1, w_2, q)$, the minimum cost of producing q units at input prices w_1 and w_2. We generally write the value function for this problem as $C(w_1, w_2, q)$ instead of $v(w_1, w_2, q)$. It's the firm's **cost function**. If we assume w_1 and w_2 are fixed, then we would write the cost function as simply $C(q)$. Of course, $C(w_1, w_2, q) = w_1\hat{x}_1(w_1, w_2, q) + w_2\hat{x}_2(w_1, w_2, q)$.

Note that while w_1, w_2, and q are all parameters in this problem, they're qualitatively different in the problem's interpretation. The parameters w_1 and w_2 are exogenous to the firm — the firm can't affect them by its actions. The parameter q is also exogenous *in the minimization problem* (that's what we mean by a parameter), but it's actually endogenous to the firm: the firm can choose the level q of its output — perhaps to maximize profit, or perhaps to maximize market share, or perhaps according to some other criterion. But we assume that it minimizes the cost of producing that q. It's for this reason that we say the factor demand functions $\hat{x}_i(q, w_1, w_2)$ are *conditional* factor demands (and the firm is conditionally minimizing its cost). The firm's overall objective isn't to minimize its cost (it could do that by producing $q = 0$); but *conditional on producing* q, it wants to minimize its cost.

Shephard's Lemma: $\dfrac{\partial C}{\partial w_i} = x_i = \hat{x}_i(\mathbf{w}, q).$

This is an immediate implication of the Envelope Theorem. And if we can observe the usage x_i of input i, Shephard's Lemma gives us the value of $\dfrac{\partial C}{\partial w_i}$ without obtaining the first-order conditions or knowing the conditional factor demand functions.

Marginal Cost: The firm's marginal cost is the derivative of its cost function, *i.e.*, the derivative of the value function with respect to q; and q is the right-hand side of the constraint in the minimization problem. So the marginal cost is the value of the Lagrange multiplier λ, which of course depends on $(q, w_1, w_2) : MC(q, w_1, w_2) = \lambda(q, w_1, w_2)$.

While this example has only two inputs, everything we've done goes through in the same way for any number of inputs.

Example 4: Roy's Identity

In demand theory, the consumer's demand function can be expressed as a ratio of derivatives of the indirect utility function, a result known as Roy's Identity.

Theorem (Roy's Identity): Let $u : \mathbb{R}^\ell_{++} \to \mathbb{R}$ be a differentiable and strictly quasiconcave utility function. Then the demand function $\hat{\mathbf{x}}(\cdot)$ and the indirect utility function $v(\cdot)$ satisfy the equations

$$\hat{x}_k(\mathbf{p}, w) = -\frac{\dfrac{\partial v}{\partial p_k}(\mathbf{p}, w)}{\dfrac{\partial v}{\partial w}(\mathbf{p}, w)}, \quad k = 1, \ldots, \ell.$$

Proof: For each $k = 1, \ldots, \ell$ the Envelope Theorem yields

$$\frac{\partial v}{\partial p_k} = \frac{\partial u}{\partial p_k} - \lambda \frac{\partial \mathbf{p} \cdot \mathbf{x}}{\partial p_k} = -\lambda x_k$$

at each $(\mathbf{p}, w) \in \mathbb{R}^{\ell+1}_{++}$, where $\lambda(\mathbf{p}, w)$ is the Lagrange multiplier in the utility-maximization problem's first-order condtion. The result then follows from the fact that $\lambda(\mathbf{p}, w) = \dfrac{\partial v}{\partial w}(\mathbf{p}, w)$. ■

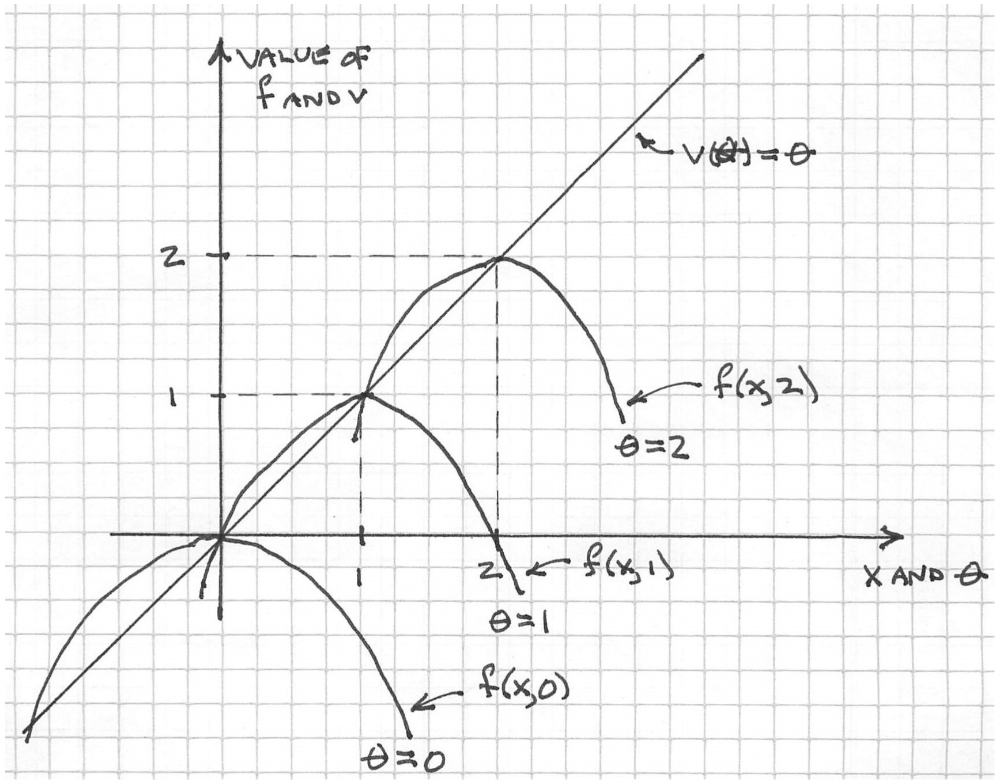

Figure 1

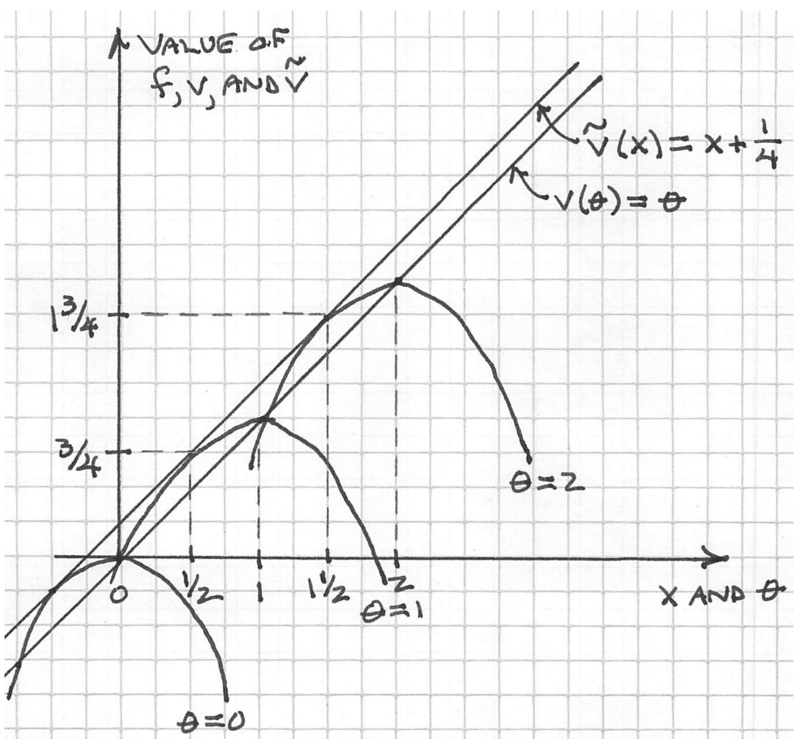

Figure 2