Legendre-Fenchel transforms in a nutshell

Hugo Touchette
School of Mathematical Sciences, Queen Mary, University of London, London E1 4NS, UK

Started: July 11, 2005; last compiled: August 14, 2007

The aim of this report is to list and explain the basic properties of the Legendre-Fenchel transform, which is a generalization of the Legendre transform commonly encountered in physics. The precise way in which the Legendre-Fenchel transform generalizes the Legendre transform is carefully explained and illustrated with many examples and pictures. The understanding of the difference between the two transforms is important because the general transform which arises in statistical mechanics is the Legendre-Fenchel transform, not the Legendre transform.

All the results contained in this report can be found with much more mathematical details and rigor in [1]. The proofs of these results can also be found in that reference.

1. Definitions

Consider a function $f(x) : \mathbb{R} \to \mathbb{R}$. We define the **Legendre-Fenchel** (LF) transform of $f(x)$ by the variational formula

$$f^*(k) = \sup_{x \in \mathbb{R}} \{kx - f(x)\}. \tag{1}$$

We also express this transform by $f^*(k) = (f(x))^*$ or, more compactly, by $f^* = (f)^*$, where the star stands for the LF transform.

The LF transform of $f^*(k)$ is

$$f^{**}(x) = \sup_{k \in \mathbb{R}} \{kx - f^*(k)\}. \tag{2}$$

This corresponds also to the **double LF transform** of $f(x)$. The double-star notation comes obviously from our compact notation for the LF transform:

$$f^{**} = (f^*)^* = ((f)^*)^*. \tag{3}$$

Remark 1. LF transforms can also be defined using an infimum (min) rather than a supremum (max):

$$g^*(k) = \inf_{x\in\mathbb{R}}\{kx - g(x)\}. \tag{4}$$

Transforming one version of the LF transform to the other is just a matter of introducing minus signs at the right place. Indeed,

$$-f^*(k) = -\sup_x\{kx - f(x)\} = \inf_x\{-kx + f(x)\}, \tag{5}$$

so that

$$g^*(q) = \inf_x\{qx - g(x)\}, \tag{6}$$

making the transformations $g(x) = -f(x)$ and $g^*(q) = -f^*(k = -q)$.

Remark 2. The Legendre-Fenchel transform is often referred to in physics as the **Legendre transform**. This does not do justice to Fenchel who explicitly studied the variational formula (1), and applied it to nondifferentiable as well as nonconvex functions. What Legendre actually considered is the transform defined by

$$f^*(k) = kx_k - f(x_k) \tag{7}$$

where x_k is determined by solving

$$f'(x) = k. \tag{8}$$

This form is more limited in scope than the LF transform, as it applies only to differentiable functions, and, we shall see later, convex functions. In this sense, the LF transform is a generalization of the Legendre transform, which reduces to the Legendre transform when applied to convex, differentiable functions. We shall comment more on this later.

Remark 3. The LF transform is not necessarily **self-inverse** (we also say **involutive**); that is to say, f^{**} need not necessarily be equal to f. The equality $f^{**} = f$ is taken for granted too often in physics; we'll see later in which cases it actually holds and which other cases it does not.

Remark 4. The definition of the LF transform can trivially be generalized to functions defined on higher-dimensional spaces (i.e., functions $f(\mathbf{x}) : \mathbb{R}^n \to \mathbb{R}$, with n a positive integer) by replacing the normal real-number product kx by the scalar product $\mathbf{k}\cdot\mathbf{x}$, where $\mathbf{k}$ is a vector having the same dimension as $\mathbf{x}$.

Remark 5. (Steepest-descent or Laplace approximation). Consider the definite integral

$$F(k,n) = \int_{\mathbb{R}} e^{n[kx - f(x)]} dx. \tag{9}$$

In the limit $n \to \infty$, it is possible to approximate this integral using Laplace Method (or steepest-descent method if $x \in \mathbb{C}$) by locating the maximum value of the integrand corresponding to the maximum value of the exponent $kx - f(x)$ (assume there's only one such value). This yields,

$$F(k,n) \approx \exp\left(n \sup_{x\in\mathbb{R}} \{kx - f(x)\} \right). \tag{10}$$

It can be proved that the corrections to this approximation are subexponential in n, i.e.,

$$\ln F(k,n) = n \sup_{x\in\mathbb{R}} \{kx - f(x)\} + o(n). \tag{11}$$

Remark 6. (The LF transform in statistical mechanics). Let U be the energy function of an n-body system. In general, the density $\Omega_n(u)$ of microscopic states of the system having a mean energy $u = U/n$ scales exponentially with n, which is to say that

$$\ln \Omega_n = ns(u) + o(n), \tag{12}$$

where $s(u)$ is the microcanonical entropy function of the system. (This can be taken as a definition of the microcanonical entropy.) Defining the canonical partition function of the system in the usual way, i.e.,

$$Z_n(\beta) = \int \Omega_n(u) e^{-n\beta u} du, \tag{13}$$

we can use Laplace Method to write

$$\varphi(\beta) = \lim_{n\to\infty} -\frac{1}{n} \ln Z_n(\beta) = \inf_u \{\beta u - s(u)\}. \tag{14}$$

Physically, $\varphi(\beta)$ represents the free energy of the system in the canonical ensemble. So, what the above result shows is that the canonical free energy is the LF transform of the microcanonical entropy ($\varphi = s^*$). The inverse result, namely $s = \varphi^*$, is not always true, as will become clear later.

2. Theory of LF transforms

The theory of LF transforms deals mainly with two questions:

Q1: How is the shape of $f^*(k)$ determined by the shape of $f(x)$, and vice versa?

Q2: When is the LF transform involutive? That is, when does $f^{**} = ((f)^*)^* = f$?

We'll see next that these two questions are answered by using a fundamental concept of convex analysis known as a supporting line.

2.1. Supporting lines

We say that the function $f : \mathbb{R} \to \mathbb{R}$ has or admits a **supporting line** at $x \in \mathbb{R}$ if there exists $\alpha \in \mathbb{R}$ such that

$$f(y) \geq f(x) + \alpha(y - x), \tag{15}$$

for all $y \in \mathbb{R}$. The parameter α is the slope of the supporting line. We further say that a supporting line is **strictly supporting** at x if

$$f(y) > f(x) + \alpha(y - x) \tag{16}$$

holds for all $y \neq x$. For these definitions to make sense, we need obviously to have $f < \infty$.

Remark 7. For convenience, it is useful to replace the expression "f admits a supporting line at x" by "f is **convex** at x". So, from now on, the two expressions mean the same (this is a definition). If f does not admit a supporting line at x, then we shall say that f is **nonconvex** at x.

The geometrical interpretation of supporting lines is shown in Figure 1. In this figure, we see that

- The point a admits a supporting line (f is convex at a). The supporting line has the property that it touches f at the point $(a, f(a))$ and lies *beneath* the graph of $f(x)$ for all x; hence the term "supporting".

- The supporting line at a is strictly supporting because it touches the graph of $f(x)$ only at a. In this case, we say that f is strictly convex at a.

- The point b does not admit any supporting lines; any lines passing through $(b, f(b))$ must cross the graph of $f(x)$ at some point. In this case, we also say that f is nonconvex at b.

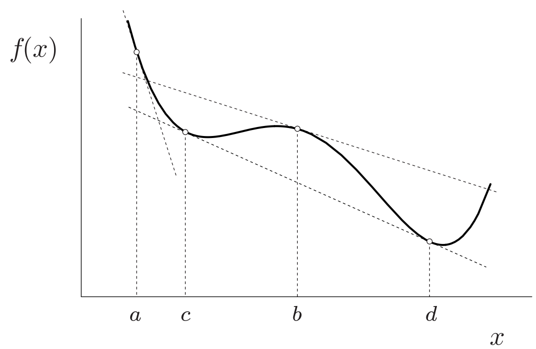

Figure 1: Geometric interpretation of supporting lines.

- The point c admits a supporting line which is non-strictly supporting, as it touches another point (d) of the graph of $f(x)$. (The points c and d share the same supporting line.)

From this picture, we easily deduce the following result:

Proposition 1. If f admits a supporting line at x and $f'(x)$ exists, then the slope α of the supporting line must be equal to $f'(x)$. In other words, for differentiable functions, a supporting line is also a **tangent** line.

2.2. Convexity properties

Before answering Q1 and Q2, let us pause briefly for two important results, which we state without proofs.

Theorem 2. $f^*(k)$ is an always convex function of k (independently of the shape of f).

Corollary 3. $f^{**}(x)$ is an always convex function of x (again, independently of the shape of f).

The precise meaning of convex here is that f^* (or f^{**}) admits a supporting line at all k (all x, respectively). More simply, it means that f^* and f^{**} are ∪-shaped.[1]

Note that these results tell us already that the LF transform cannot always be involutive. Indeed, $f^{**}(x)$ is convex even if $f(x)$ is not, so that $f \neq f^{**}$ if f is not everywhere convex. We'll see later that this is the only problematic case.

[1] There seems to be some confusion in the literature about the definitions of "concave" and "convex." The Webster (7th Edition), for one, defines a ∪-shaped function to be *concave* rather than convex. However, most mathematical textbooks will agree in defining the same function to be convex.

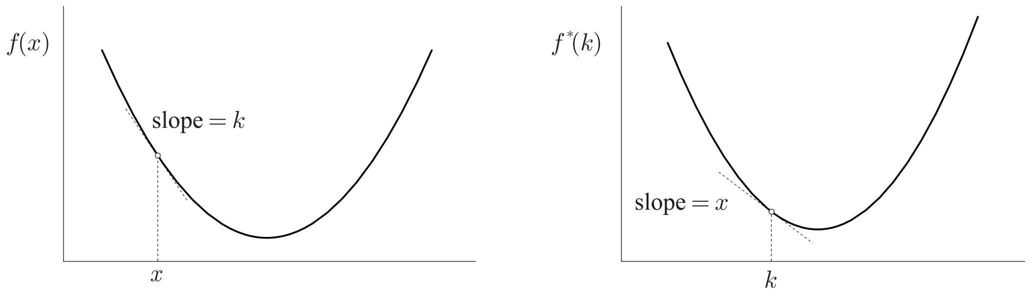

Figure 2: Illustration of the duality property for supporting lines: points of f are transformed into slopes of f^*, and slopes of f are transformed into points of f^*.

2.3. Supporting line duality

We now answer our first question (Q1): How is the shape of $f^*(k)$ determined by the shape of $f(x)$, and vice versa? A partial answer is provided by the following result:

Theorem 4. If f admits a supporting line at x with slope k, then f^* at k admits a supporting line with slope x.

This theorem is illustrated in Figure 2. The next theorem covers the special case of strict convexity.

Theorem 5. If f admits a strict supporting line at x with slope k, then f^* admits a tangent supporting line at k with slope $f^{*\prime}(k) = x$. (Hence f^* is differentiable in this case in addition to admit a supporting line.)

2.4. Inversion of LF transforms

The answer to Q2 ($f \stackrel{?}{=} f^{**}$) is provided by the following result:

Theorem 6. $f(x) = f^{**}(x)$ if and only if f admits a supporting line at x.

Thus, from the point of view of $f(x)$, we have that the LF transform is involutive at x if and only if f is convex at x (in the sense of supporting lines). Changing our point of view to $f^*(k)$, we have the following:

Theorem 7. If f^* is differentiable at k, then $f = f^{**}$ at $x = f^{*\prime}(k)$.

We'll see later with a specific example that the differentiability property of f^* is sufficient (as stated) but non-necessary for $f = f^{**}$. For now, we note the following obvious corollary:

Corollary 8. If $f^*(k)$ is everywhere differentiable, then $f(x) = f^{**}(x)$ for all x.

This says in words that the LF transform is completely involutive if $f^*(k)$ is everywhere differentiable.

We end this section with another corollary and a result which helps us visualize the meaning of $f^{**}(x)$.

Corollary 9. A convex function can always be written as the LF transform of another function. (This is not true for nonconvex functions.)

Theorem 10. $f^{**}(x)$ is the largest convex function satisfying $f^{**}(x) \leq f(x)$.

Because of this result, we call $f^{**}(x)$ the **convex envelope** or **convex hull** of $f(x)$. We'll precise the meaning of these expressions in the next section.

3. Some particular cases

We consider in this section a number of examples which help visualize the meaning and application of all the results presented in the previous section. All of the examples considered arise in statistical mechanics.

3.1. Differentiable, convex functions

The LF transform

$$f^*(k) = \sup_x \{kx - f(x)\} \tag{17}$$

is in general evaluated by finding the critical points x_k (there could be more than one) which maximize the function

$$F(x,k) = kx - f(x). \tag{18}$$

In mathematical notation, we express x_k in the following manner:

$$x_k = \arg\sup_x F(x,k) = \arg\sup_x \{kx - f(x)\}, \tag{19}$$

where "arg sup" reads "arguments of the supremum," and mean in words "points at which the maximum occurs."

Now, assume that $f(x)$ is everywhere differentiable. Then, we can find the maximum of $F(x,k)$ using the common rules of calculus by solving

$$\frac{\partial}{\partial x} F(x,k) = 0, \tag{20}$$

for a fixed value of k. Given the form of $F(x,k)$, this is equivalent to solving

$$k = f'(x) \tag{21}$$

for x given k. As noted before, there could be more than one critical points of $F(x,k)$ that would solve here the above differential equation. To make sure that there is actually only one solution for every $k \in \mathbb{R}$, we need to impose the following two conditions on f:

1. $f'(x)$ is continuous and monotonically increasing for increasing x;

2. $f'(x) \to \infty$ for $x \to \infty$ and $f'(x) \to -\infty$ for $x \to -\infty$.

Given these, we are assured that there exists a unique value x_k for each $k \in \mathbb{R}$ satisfying $k = f'(x_k)$ and which maximizes $F(x,k)$. As a result, we can write

$$f^*(k) = kx_k - f(x_k), \tag{22}$$

where

$$f'(x_k) = k. \tag{23}$$

These two equations define precisely what the Legendre transform of $f(x)$ is (as opposed to the LF transform, which is defined with the sup formula). Accordingly, we have proved that the LF transform reduces to the Legendre transform for differentiable and strictly convex functions. (The strictly convex property results from the monotonicity of $f'(x)$.) Since $f(x)$ at this point is convex by assumption, we must have $f = f^{**}$ for all x. Therefore, the Legendre transform must be involutive (always), and the inverse Legendre transform is the Legendre transform itself; in symbol,

$$f(x) = k_x x - f^*(k_x), \tag{24}$$

where k_x is the unique solution of

$$f^{*\prime}(k) = x. \tag{25}$$

3.2. Function having a nondifferentiable point

What happens if $f(x)$ has one or more nondifferentiable points? Figure 3 shows a particular example of a function $f(x)$ which is nondifferentiable at x_c. What does its LF transform $f^*(k)$ look like?

The answer is provided by what we have learned about supporting lines. Let's consider the differentiable and nondifferentiable parts of $f(x)$ separately:

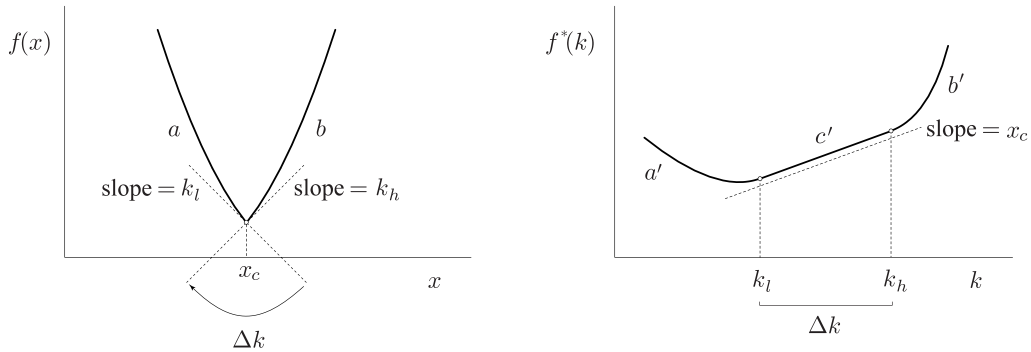

Figure 3: Function having a nondifferentiable point; its LF transform is affine.

- Differentiable points of f: Each point $(x,\ f(x))$ on the differentiable branches of $f(x)$ admits a strict supporting line with slope $f'(x) = k$. From the results of the previous section, we then know that these points are transformed at the level of $f^*(k)$ into points $(k,\ f^*(k))$ admitting supporting line of slopes $f^{*\prime}(k) = x$. For example, the differentiable branch of $f(x)$ on the left (branch a in Figure 3) is transformed into a differentiable branch of $f^*(k)$ (branch a') which extends over all $k \in (-\infty, k_l]$. This range of k-values arises because the slopes of the left-branch of $f(x)$ ranges from $-\infty$ to k_l. Similarly, the differentiable branch of $f(x)$ on the right (branch b) is transformed into the right differentiable branch of $f^*(k)$ (branch b'), which extends from k_h to $+\infty$. (Note that, for the two differentiable branches, the LF transform reduces to the Legendre transform.)

- Nondifferentiable point of f: The nondifferentiable point x_c admits not one but infinitely many supporting lines with slopes in the range $[k_l, k_h]$. As a result, each point of $f^*(k)$ with $k \in [k_l, k_h]$ must admit a supporting line with constant slope x_c (branch c'). That is, $f^*(k)$ must have a constant slope $f^{*\prime}(k) = x_c$ in the interval $[k_l, k_h]$. We say in this case that $f^*(k)$ is **affine** or **linear** over (k_l, k_h). (The affinity interval is always the open version of the interval over which f^* has constant slope.)

The case of functions having more than one nondifferentiable point is treated similarly by considering each nondifferentiable point separately.

3.3. Affine function

Since $f(x)$ in the previous example is convex, $f(x) = f^{**}(x)$ for all x, and so the roles of f and f^* can be inverted to obtain the following: a convex function $f(x)$ having an

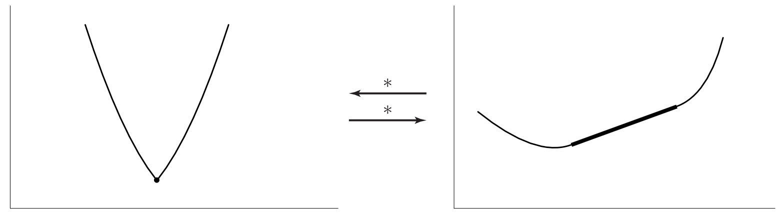

Figure 4: Nondifferentiable points are transformed into affine parts under the action of the LF transform and vice versa.

affine part has a LF transform $f^*(k)$ having one nondifferentiable point; see Figure 4. More precisely, if $f(x)$ is affine over (x_l, x_h) with slope k_c in that interval, then $f^*(k)$ will have a nondifferentiable point at k_c with left- and right-derivatives at k_c given by x_l and x_h, respectively.

3.4. Bounded-domain function with infinite slopes at boundaries

Consider the function $f(x)$ shown in Figure 5. This function has the particularity to be defined only on a bounded domain of x-values, which we denote by $[x_l, x_h]$. Furthermore, $f'(x) \to \infty$ as $x \to x_l + 0$ and $x \to x_h - 0$ (the derivative of f blows up near at the boundaries). Outside the interval of definition of $f(x)$, we formally set $f(x) = \infty$.

To determine the shape of $f^*(k)$, we use again what we know about supporting lines of f and f^*. All points $(x, f(x))$ with $x \in (x_l, x_h)$ admit a strict supporting line with slope $k(x)$. These points are represented at the level of f^* by points $(k(x), f^*(k(x)))$ having a supporting line of slope x. As x approaches x_l from the right, the slope of $f(x)$ diverges to $-\infty$. At the level of f^*, this implies that the slope of the supporting line of f^* reaches x_l as $k \to -\infty$. Similarly, since the slope of $f(x)$ goes to $+\infty$ as $x \to x_h$, the slope of the supporting line of f^* reaches the value x_h as $k \to +\infty$; see Figure 5.

Note, finally, that $f = f^{**}$ since f is convex. This means that we can invert the roles of f and f^* in this example just like in the previous one to obtain the following: the LF transform of a convex function which is asymptotically linear is a convex function which is finite on a bounded domain with diverging slopes at the boundaries.

3.5. Bounded-domain function with finite slopes at boundaries

Consider now a variation of the previous example. Rather than having diverging slopes at the boundaries x_l and x_h, we assume that $f(x)$ has finite slopes at these points. We denote the right-derivative of f at x_l by k_l and its left-derivative at x_h by k_h.

For this example, everything works as in the previous example except that we have to be careful about the boundary points. As in the case of nondifferentiable points, f at

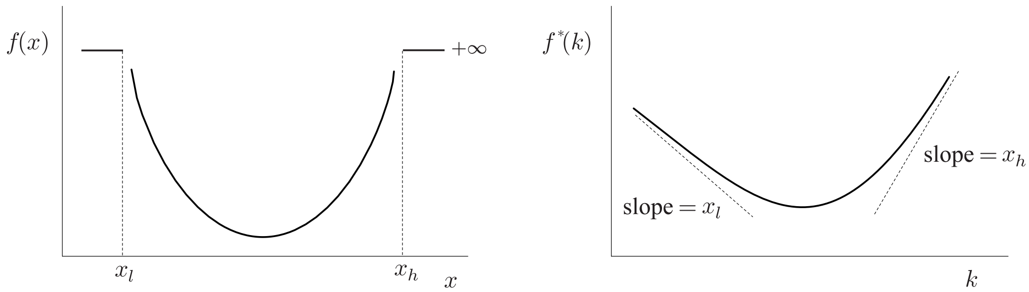

Figure 5: Function defined on a bounded domain with diverging slopes at boundaries; its LF transform is asymptotically linear as $|k| \to \pm\infty$.

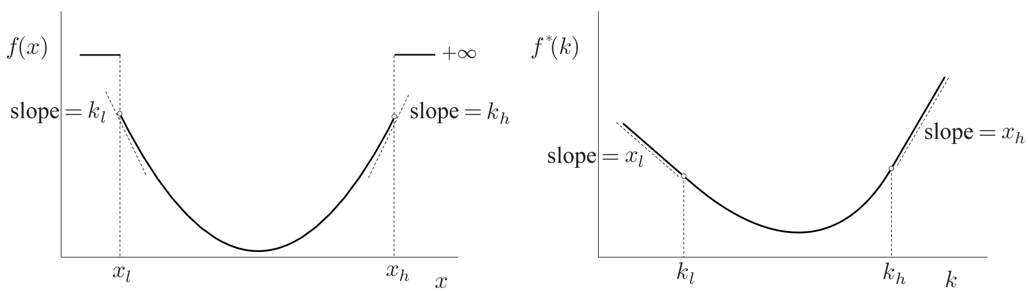

Figure 6: Function defined on a bounded domain with finite slopes at boundaries; its LF transform has affine parts outside some interior domain.

x_h admits not one but infinitely many supporting lines with slopes taking values in the range $[k_h, \infty)$. At the level of f^*, this means that all points $(k, f^*(k))$ with $k \in [k_h, \infty)$ have supporting lines with constant slope x_h; that is, $f^*(k)$ is affine past k_h with slope x_h. Likewise, f at x_l admits an infinite number of supporting lines with slopes now ranging from $-\infty$ to k_l. As a consequence, f^* must be affine over the range $(-\infty, k_l)$ with constant slope x_l; see Figure 6.

3.6. Nonconvex function

Our last example is quite interesting, as it illustrates the precise case for which the LF transform is not involutive, namely nonconvex functions.

The function that we consider is shown in Figure 7; it has three branches having the following properties:

- Branch a: The points on this branch, which extends from $x = -\infty$ to x_l, admit strict supporting lines. This branch is thus transformed into a differentiable branch

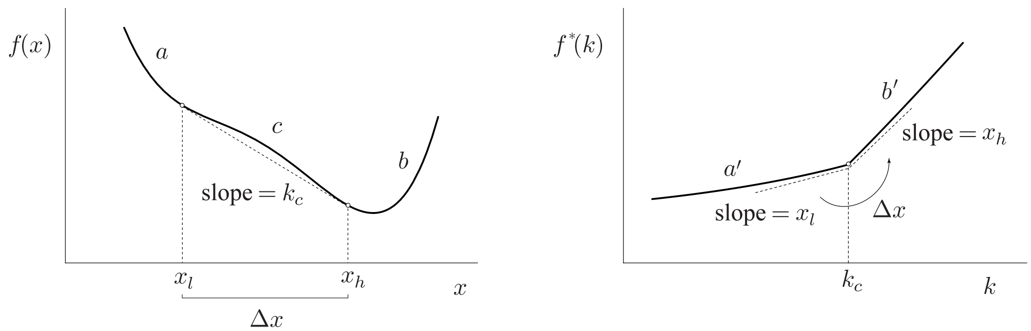

Figure 7: Nonconvex function; its LF transform has a nondifferentiable point.

at the level of f^* (branch a').

- Branch b: Similarly as for branch a.
- Branch c: None of the points on this branch, which extends from (x_l, x_h), admit supporting lines. This means that these points are not represented at the level of f^*. In other words, there is not one point of f^* which admits a supporting line with slope in the range (x_l, x_h). (That would contradict the fact that f^* has a supporting line at k with slope x if and only if f admits a supporting line at x with slope k.)

These three observations have two important consequences (see Figure 7):

1. $f^*(k)$ must have a nondifferentiable point at k_c, with k_c equal to the slope of the supporting line connecting the two points $(x_l, f(x_l))$ and $(x_h, f(x_h))$. This follows since x_l and x_h share the same supporting line of slope k_c. Thus, in a way, f^* must have two slopes at k_c.
2. Define the **convex extrapolation** of $f(x)$ to be the function obtained by replacing the nonconvex branch of $f(x)$ (branch c) by the supporting line connecting the two convex branches of f (a and b). Then, both the LF transforms of f and its convex extrapolation yield f^*. This is evident from our previous working of nondifferentiable and affine functions. It should also be evident from the example of nondifferentiable functions that the convex extrapolation of f is nothing but f^{**}, the double LF transform of f. This explains why we call f^{**} the convex envelope of f.

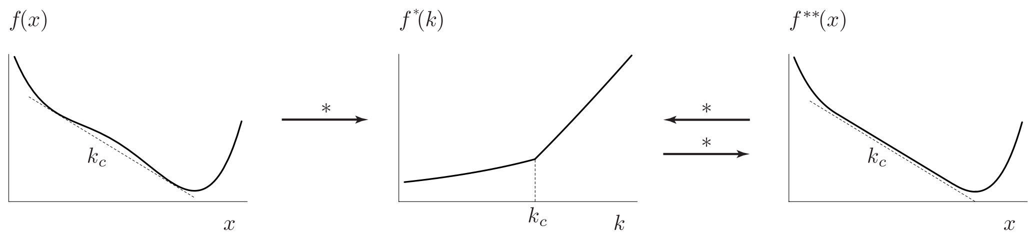

Figure 8: Structure of the LF transform for nonconvex functions.

To summarize, note that, as a result of Point 2 above, we have

$$(f^{**})^* = (f)^* = f^*. \tag{26}$$

Also, for the example considered, we have

$$(f^*)^* = f^{**} \neq f. \tag{27}$$

Overall, this means that the LF transform has the following structure:

$$f \rightarrow f^* \rightleftharpoons f^{**}, \tag{28}$$

where the arrows stand for the LF transform; see Figure 8. This diagram clearly shows that the LF transform is non-involutive in general. For convex functions, i.e., functions admitting supporting lines everywhere, the diagram reduces to

$$f \rightleftharpoons f^*. \tag{29}$$

That is, in this case, the LF transform is involutive (see Theorem 2.4).

4. Important results to remember

- The LF transform yields only convex functions: $f^* = (f)^*$ is convex and so is $f^{**} = (f^*)^*$.
- The shape of f^* is determined from the shape of f by using the duality relationship which exists between the supporting lines of f^* and those of f.
  - Points of f are transformed into slopes of f^*, and slopes of f are transformed into points of f^*.
  - Nondifferentiable points of f are transformed, through the action of the LF transform, into affine branches of f^*.

  - Affine or nonconvex branches of f are transformed into nondifferentiable points of f^*. These are the only two cases producing nondifferentiable points.

- The involution (self-inverse) property of the LF transform is determined from the supporting line properties of f or from the differentiability properties of f^*.

  - $f = f^{**}$ at x if and only if f admits a supporting line at x.
  - If f^* is differentiable at k, then $f = f^{**}$ at $x = f^{*\prime}(k)$.

- The double LF transform f^{**} of f corresponds to the convex envelope of f.

- The complete structure of the LF transform for general functions goes as follows:

$$f \to f^* \rightleftharpoons f^{**}, \tag{30}$$

  where the arrows denote the LF transform. For convex functions ($f = f^{**}$), this reduces to

$$f \rightleftharpoons f^*; \tag{31}$$

  i.e., in this case, the LF transform is involutive.

- The LF transform is more general that the Legendre transform because it applies to nonconvex functions as well as nondifferentiable functions.

- The LF transform reduces to the Legendre transform in the case of convex, differentiable functions.

Acknowledgments

This report was written while visiting the Rockefeller University, New York, in July 2005. I wish to thank E.G.D. Cohen for his invitation to visit Rockefeller, as well as for his many comments on the content and presentation of this report, and for his contagious desire to make things clear. My work is supported by the Natural Sciences and Engineering Research Council of Canada (Post-Doctoral Fellowship).

References

[1] R. T. Rockafellar. *Convex Analysis*. Princeton University Press, Princeton, 1970.