Chapter 1

CONVEX ANALYSIS IN THE CALCULUS OF VARIATIONS

R. T. Rockafellar

Abstract Convexity properties are vital in the classical calculus of variations, and many notions of convex analysis, such as the Legendre-Fenchel transform, have their origin there. Conceptual developments in contemporary convex analysis have in turn enriched that venerable subject by making it possible to treat a vastly larger class of problems effectively in a "neoclassical" framework of extended-real-valued functions and their subgradients.

Keywords: convex analysis, calculus of variations, generalized Euler-Lagrange equations, Legendre-Fenchel transform, generalized Hamiltonian equations, neoclassical problems of Bolza, optimal control

1. CLASSICAL LAGRANGE PROBLEMS

The calculus of variations is the oldest branch of optimization, dating from over three hundred years ago. Most books on that subject are steeped in traditional thinking and show little influence of convex analysis as we now know it elsewhere in optimization. Nevertheless convexity properties have an essential role in the calculus of variations. They can be seen as deeply affecting the statement of necessary and sufficient conditions for optimality and, through investigations of the existence of solutions, the very way that problems ought to be posed. Concepts emerging from convex analysis have furthermore led to the development of a mathematical framework so sturdy and broad that traditional statements in the calculus of variations can be stretched in their interpretation so as to cover problems of optimal control and much more than ever might have been imagined. This article aims at explaining in basic terms why—and how—that has come about.

The fundamental *Lagrange problem* in the calculus of variations, with respect to a fixed interval $[\tau_0, \tau_1]$ and fixed endpoints ξ_0 and ξ_1, has the form:

$$(\mathcal{P}_0) \qquad \begin{array}{c} \text{minimize } \; J_0[x] := \int_{\tau_0}^{\tau_1} L(t, x(t), \dot{x}(t))dt \\ \text{subject to } \; x(\tau_0) = \xi_0, \; x(\tau_1) = \xi_1, \end{array}$$

for a function $L : [\tau_0, \tau_1] \times I\!R^n \times I\!R^n \to I\!R$ which is assumed to be continuously differentiable. The minimization takes place in a space of functions $x : [\tau_0, \tau_1] \to I\!R^n$, called *arcs*, with $\dot{x}(t)$ denoting the derivative of $x(t)$. But what space should that be? The simplest and seemingly most natural choice is

$$\mathcal{C}_n^1[\tau_0, \tau_1] \; = \; \text{space of continuously differentiable arcs.}$$

The integral functional J_0 is well defined on this space and even continuously Fréchet differentiable under a standard choice of norm. In terms of the gradients of $L(t, x, v)$ with respect to x and v, the directional derivatives are given by

$$J_0'[x; w] = \int_{\tau_0}^{\tau_1} \Big[\langle \nabla_x L(t, x(t), \dot{x}(t)), w(t) \rangle + \langle \nabla_v L(t, x(t), \dot{x}(t)), \dot{w}(t) \rangle \Big] dt, \tag{1}$$

where $\langle \cdot, \cdot \rangle$ denotes the canonical inner product in $I\!R^n$. A classical argument yields the fact that if an arc x is optimal in $(\mathcal{P}_0)$ (where for purposes of this article we always interpret optimality as *global* optimality), the *Euler-Lagrange condition* must hold:

$$\frac{d}{dt} \nabla_v L(t, x(t), \dot{x}(t)) = \nabla_x L(t, x(t), \dot{x}(t)), \tag{2}$$

along with the *Weierstrass condition*:

$$L(t, x(t), v) \geq L(t, x(t), \dot{x}(t)) + \langle \nabla_v L(t, x(t), \dot{x}(t)), \; v - \dot{x}(t) \rangle \; \text{ for all } \; v, \tag{3}$$

which in turn entails the *Legendre condition*:

$$\nabla_{vv}^2 L(t, x(t), \dot{x}(t)) \; \text{ is positive semi-definite} \tag{4}$$

when $L(t, x(t), v)$ is twice differentiable with respect to v, so that the Hessian matrix $\nabla_{vv}^2 L(t, x(t), \dot{x}(t))$ exists. Already here, convexity enters implicitly in important ways.

First and most obvious is the connection between the Legendre condition and possible convexity of the function $f(v) = L(t, x(t), v)$. When f is twice differentiable, it is a convex function if and only if its Hessian $\nabla^2 f(v)$ is positive semi-definite for all v. Of course, the Legendre condition (4) asserts positive semi-definiteness only for $v = \dot{x}(t)$, but still it has suggestive potential. The strengthened form of the Legendre condition in which positive semi-definiteness is replaced by positive definiteness, would, in the presence of continuous twice differentiability, imply that f is strongly convex on a neighborhood of $\dot{x}(t)$. In this sense, convexity is close by. In fact the Weierstrass condition, which supersedes the Legendre condition, brings it even closer.

To see this aspect of the Weierstrass condition, let $y(t) = \nabla_v L(t, x(t), \dot{x}(t))$ and again consider $f(v) = L(t, x(t), v)$ for fixed t, so that $y(t) = \nabla f(\dot{x}(t))$ and (3) can be written as

$$f(v) \;\geq\; f(\dot{x}(t)) + \langle y(t),\, v - \dot{x}(t) \rangle \;\text{ for all }\; v. \tag{5}$$

The right side of the inequality is an affine function $a(v)$; the inequality says that $f(v) \geq a(v)$ everywhere, and that these functions agree at $v = \dot{x}(t)$. Unlike the Legendre condition, this is not just a one-point property but a global property. Convexity is brought in because (5) would automatically hold if f is a convex function, whereas it seems puzzlingly strong if f is not convex. Indeed, let

$$\bar{f} \;=\; \text{convex hull of }\; f,$$

this being the greatest convex function majorized by f. From our knowledge that (5) holds for $y(t) = \nabla f(\dot{x}(t))$, it can be verified that $\bar{f}$ is a finite convex function which is differentiable at $\dot{x}(t)$ and has the same gradient there as f, i.e., the vector $y(t)$. Thus, the real content of the Weierstrass condition (3) is in the relations

$$f(\dot{x}(t)) = \bar{f}(\dot{x}(t)), \qquad \nabla f(\dot{x}(t)) = \nabla \bar{f}(\dot{x}(t)). \tag{6}$$

In other words, it says something about how the optimality of an arc x implies that L acts along x as if it were replaced by its convex hull in the v argument.

The issue of convexity or convexification of $L(t, x, v)$ with respect to v has significance also from a different perspective. When can we be sure that an optimal arc even exists in problem $(\mathcal{P}_0)$? It may be anticipated that the existence of an optimal arc depends on some kind of compactness property of the sublevel sets $\{x \mid J_0[x] \leq \alpha\}$, $\alpha \in I\!R$, but such compactness must be tuned to the space over which the minimization takes place, and $\mathcal{C}^1_n[\tau_0, \tau_1]$ is for that reason a poor choice. Often in textbooks, $\mathcal{C}^1_n[\tau_0, \tau_1]$ is replaced by the space $\bar{\mathcal{C}}^1_n[\tau_0, \tau_1]$ consisting of the arcs $x : [\tau_0, \tau_1] \to I\!R^n$ that are *piecewise* continuously differentiable; elementary examples abound in which solutions can be found in $\bar{\mathcal{C}}^1_n[\tau_0, \tau_1]$ but not in $\mathcal{C}^1_n[\tau_0, \tau_1]$, and this has long been recognized. But $\bar{\mathcal{C}}^1_n[\tau_0, \tau_1]$ fares no better than $\mathcal{C}^1_n[\tau_0, \tau_1]$ in attempts to establish the existence of solutions in general. We know now that to get existence we have to pass to a far more general space and introduce both coercivity and *convexity* assumptions on the behavior of $L(t, x, v)$ with respect to v.

Let $|v|$ stand for the Euclidean norm of $v \in I\!R^n$. The function L is said to satisfy the *Tonelli growth condition* if

$$L(t, x, v) \geq \theta(|v|) \;\text{ for all }\; (t, x, v) \in [\tau_0, \tau_1] \times I\!R^n \times I\!R^n \tag{7}$$

for a nondecreasing function $\theta : [0, \infty) \to (-\infty, \infty)$ that is coercive, i.e., has $\theta(s)/s \to \infty$ as $s \to \infty$. For $p \in [1, \infty]$ let

$$\mathcal{A}^p_n[\tau_0, \tau_1] \;=\; \text{space of absolutely continuous arcs } \; x \; \text{ with } \; \dot{x} \in \mathcal{L}^p_n[\tau_0, \tau_1].$$

Thus, $\mathcal{A}^1_n[\tau_0, \tau_1]$ consists simply of all absolutely continuous $x : [\tau_0, \tau_1] \to I\!\!R^n$, whereas $\mathcal{A}^\infty_n[\tau_0, \tau_1]$ consists of all Lipschitz continuous $x : [\tau_0, \tau_1] \to I\!\!R^n$. Of course an arc x in any of these spaces is only differentiable almost everywhere, so that the integrand $L(t, x(t), \dot{x}(t))$ itself is only defined almost everywhere—but that is enough to make sense of the functional J_0, provided that its value is interpreted as ∞ in the eventuality that $L(t, x(t), \dot{x}(t))$ is not bounded from above by an integrable function of t.

Theorem 1. *Let problem $(\mathcal{P}_0)$ be placed in the arc space $\mathcal{A}^1_n[\tau_0, \tau_1]$. If the function L satisfies the Tonelli growth condition, and $L(t, x, v)$ is convex with respect to v, then the minimum in $(\mathcal{P}_0)$ is attained.*

This fact suggests that $\mathcal{A}^1_n[\tau_0, \tau_1]$ should be adopted as the arc space without further ado, but unfortunately there are impediments to that. On this larger space the functional J_0 lacks adequate differentiability, except in very special circumstances, and the traditional derivation of optimality conditions collapses. A fallback instead of $\mathcal{A}^1_n[\tau_0, \tau_1]$ is $\mathcal{A}^\infty_n[\tau_0, \tau_1]$, on which the directional derivative formula (1) persists under our working assumptions on L.

Theorem 2. *Let problem $(\mathcal{P}_0)$ be placed in the arc space $\mathcal{A}^\infty_n[\tau_0, \tau_1]$. Then, for an arc x that is optimal in $(\mathcal{P}_0)$, the Euler-Lagrange equation (2) and the Weierstrass condition (3) must hold in the following interpretation: there is an arc y, likewise in $\mathcal{A}^\infty_n[\tau_0, \tau_1]$, such that for almost every $t \in [\tau_0, \tau_1]$ one has*

$$y(t) = \nabla_v L(t, x(t), \dot{x}(t)), \qquad \dot{y}(t) = \nabla_x L(t, x(t), \dot{x}(t)), \tag{8}$$

$$L(t, x(t), v) \;\geq\; L(t, x(t), \dot{x}(t)) + \langle y(t),\, v - \dot{x}(t) \rangle \;\; \textit{for all} \;\; v \tag{9}$$

Note that the formulation of the Euler-Lagrange condition in (8) is preferable to the version in (2) even when $(\mathcal{P}_0)$ is placed in $\mathcal{C}^1_n[\tau_0, \tau_1]$, in which case y too belongs to $\mathcal{C}^1_n[\tau_0, \tau_1]$. It makes explicit the differentiability property of $\nabla_v L(t, x(t), \dot{x}(t))$ in t that is tacit in the original statement (2). The arc y, said to be *adjoint* to x, deserves the added attention anyway.

Note also that the Weierstrass condition (9), which still implies the Legendre condition (4), once more holds automatically when $L(t, x, v)$ is convex in v, for the reasons already noted. Without such convexity, (9) can again be seen as a condition relating the function $f(v) = L(t, x(t), v)$ to its convex hull. But the discussion of this issue is incomplete without bringing the so-called canonical from of the Euler-Lagrange condition into the picture, which we do next.

In further analysis of the first equation in (8), let us step back from the consideration of arcs x and y and look at the equation $y = \nabla_v L(t, x, v)$ as a relation between vector variables y and v in $I\!R^n$ for given $(t, x) \in [\tau_0, \tau_1] \times I\!R^n$. Suppose we are able to invoke the implicit function theorem in order to solve this equation uniquely for v as a function of (t, x, y); the technical assumption needed for this locally is that the mapping $v \mapsto \nabla_v L(t, x, v)$ is continuously differentiable and has nonsingular Jacobian, or in other words, that $L(t, x, v)$ is twice differentiable in v, with the Hessian $\nabla^2_{vv} L(t, x, v)$ being nonsingular and depending continuously on (t, x, v). Let the solution mapping be denoted by $v = h(t, x, y)$, and let

$$H(t, x, y) = \langle y, v \rangle - L(t, x, v) \text{ with } h(t, x, y) \text{ substituted for } v. \tag{10}$$

Formulas for the derivatives of H can be obtained from the implicit function theorem; one has

$$\begin{cases} \nabla_y H(t, x, y) = h(t, x, y), \\ \nabla_x H(t, x, y) = -\nabla_x L(t, x, h(t, x, y)). \end{cases}$$

Then, on returning to the context of arcs, we can express the Euler-Lagrange condition (8) in the form

$$\dot{x}(t) = \nabla_y H(t, x(t), y(t)), \qquad \dot{y}(t) = -\nabla_x H(t, x(t), y(t)), \tag{11}$$

and recognize that this constitutes an ordinary differential equation in the form of $\dot{z}(t) = F(t, z(t))$ for $z(t) = (x(t), y(t))$. The conversion of (8) to an ODE is obviously both illuminating and advantageous.

The function H is known as the *Hamiltonian* associated with L, which in contrast is itself often called the *Lagrangian*; the process of obtaining $H(t, x, y)$ from $L(t, x, v)$ by way of the implicit function theorem in the manner described is the *Legendre transform* (with respect to the v variable with (t, x) fixed). The relations (11) are the *Hamiltonian equations*.

Unfortunately, in this traditional development there is a serious flaw which textbooks seldom address. The nonsingularity assumption on $\nabla^2_{vv} L(t, x, v)$ on which the Legendre transform depends is local. Indeed, it is often envisioned merely along the x trajectory being tested for optimality, where it comes down to the nonsingularity of $\nabla^2_{vv} L(t, x(t), \dot{x}(t))$ and therefore the positive definiteness of this Hessian, since we already know from the Legendre condition (4) that $\nabla^2_{vv} L(t, x(t), \dot{x}(t))$ must be positive semi-definite. Thus, if this approach through the Legendre transform is adopted, the Hamiltonian H threatens to be defined only in a local sense in relation to a particular x trajectory, and only to the extent that $L(t, x(t), v)$ is strongly convex in v around $v = \dot{x}(t)$. Furthermore, it can be seen from such local convexity that the prescription for H in (10) then implies

$$H(t, x(t), y(t)) = \text{ local max of } \langle y(t), v \rangle - L(t, x(t), v) \text{ in } v \text{ around } \dot{x}(t),$$

with the local max being attained when $v = \dot{x}(t)$. Another angle, however, is that the Weierstrass condition says

$$\text{global max of } \langle y(t), v\rangle - L(t, x(t), v) \text{ in } v \text{ is attained at } \dot{x}(t).$$

Obviously this situation is perplexing and unsatisfactory. An answer in part could be to insist that $\nabla^2_{vv} L(t, x, v)$ be positive definite for all $(t, x, v) \in [\tau_0, \tau_1] \times I\!R^n \times I\!R^n$, so as to avoid the awkward localization, but that would be unpleasantly restrictive.

Convex analysis will come to the rescue with a superior definition of the Hamiltonian function H. That definition will rely instead on the Legendre-Fenchel transform, which reduces to the Legendre transform in the presence of continuous second-order differentiability with positive definite Hessians, but actually makes sense even without any differentiability. Before getting to that, however, it is good to complement the discussion of necessary conditions by taking a brief look at how convexity supports sufficient conditions for optimality.

Theorem 3. *Let problem $(\mathcal{P}_0)$ be placed in the arc space $\mathcal{A}^p_n[\tau_0, \tau_1]$ for any $p \in [1, \infty]$, and suppose that $L(t, x, v)$ is convex with respect to (x, v). If the Euler-Lagrange condition in form (8) holds for arcs x and y with $x \in \mathcal{A}^p_n[\tau_0, \tau_1]$ and $y \in \mathcal{A}^1_n[\tau_0, \tau_1]$, then x is optimal in $(\mathcal{P}_0)$.*

The proof of this result is elementary but also illuminating because it opens doors to generalization. The convexity of $L(t, x, v)$ not just in v but with respect to (x, v) yields the gradient inequality for the function $L(t, \cdot, \cdot)$ that, when used to compare the pair $(x(t), \dot{x}(t))$ from the arc x with a pair $(x'(t), \dot{x}'(t))$ coming from an alternative arc x', says in the context of (8) that

$$L(t, x'(t), \dot{x}'(t)) \geq L(t, x(t), \dot{x}(t)) + \langle \dot{y}(t), x'(t) - x(t)\rangle + \langle y(t), \dot{x}'(t) - \dot{x}(t)\rangle.$$

On integrating both sides of this inequality we get

$$J_0[x'] \geq J_0[x] + \langle y(\tau_1),\, x'(\tau_1) - x(\tau_1)\rangle - \langle y(\tau_0),\, x'(\tau_0) - x(\tau_0)\rangle.$$

When x', like x, satisfies the endpoint constraints $x'(\tau_0) = \xi_0$, $x'(\tau_1) = \xi_1$, we obtain $J_0[x'] \geq J_0[x]$ and thus the global optimality of x.

Observe that under the convexity hypothesis on L in Theorem 3, J_0 is a convex functional on the arc space in question.

Theorem 3 does not give the most general form of sufficient condition in the calculus of variations. In the classical theory that would be a hard-to-verify criterion known as the Jacobi condition, involving second derivatives of the functional J_0 for which there is an integral expression based on second derivatives of L. By contrast, the condition in Theorem 3, although not as comprehensive, is simple to use and still covers many applications. But books usually neglect to mention it.

Variants of the optimality conditions in Theorems 2 and 3 for problems beyond $(\mathcal{P}_0)$, having boundary conditions other than $x(\tau_0) = \xi_0$ and $x(\tau_1) = \xi_1$, are well known and easily obtained. If the first endpoint of x is fixed but the second endpoint is free, the condition $y(\tau_1) = 0$ has to be added, whereas if the free endpoint is the first instead of the second, the requirement is $y(\tau_0) = 0$. If both endpoints are free but a cost term $l(x(\tau_0), x(\tau_1))$ is added to the integral in the functional J_0 being minimized, one has to have

$$(y(\tau_0), -y(\tau_1)) = \nabla l(x(\tau_0), x(\tau_1)).$$

Mixtures of such possibilities can also be handled, involving an endpoint cost term in J_0 but also a system of constraints on $x(\tau_0)$ and $x(\tau_1)$, not only equations but perhaps also inequalities. In every such case, optimality entails a certain relation between $(x(\tau_0), x(\tau_1))$ and $(y(\tau_0), y(\tau_1))$, known historically as a *transversality condition*. Classical theory does not provide a universal way of expressing transversality conditions, but a simple, unifying scheme does come out of the modern developments described below.

Besides these variants on how to treat $x(\tau_0)$ and $x(\tau_1)$ in $(\mathcal{P}_0)$ there can be constraints of equality or inequality type, or both, imposed on $x(t)$ and $\dot{x}(t)$, separately or jointly. When all such complications are allowed simultaneously in the problem statement, one gets a formulation that marked the high point in the calculus of variations before the advent of optimal control, called a *Bolza problem*.

It is also possible to consider problems in which the interval $[\tau_0, \tau_1]$ is not fixed but subject to optimization as well, but we will skip over that in the interests of simplicity.

2. NEOCLASSICAL BOLZA PROBLEMS

The classical framework for the calculus of variations has led to many successes, most notably in applications to physics. Its shortcomings, such as have been pointed out above, were not really apparent until "optimization" arose in the 1950's as a practical subject with a numerical orientation. Inequality constraints came to be viewed not just as an occasional necessity in problem formulation, but rather as a dominant feature affecting the analysis of solutions and the design of algorithms for finding them. Penalty expressions with discontinuities in derivatives entered the stage that way too. Methods of differential calculus were no longer enough, whether in connection with such expressions or in coping with the geometry of sets specified by systems of inequality constraints. On the other hand, convexity properties of sets and functions took on a new significance, and fascinating phenomena of duality came to light.

For the calculus of variations, the most revolutionary development was the emergence of optimal control theory. A typical problem of optimal control is $(\mathcal{P}_1)$

$$\text{minimize } J_1[x,u] := \int_{\tau_0}^{\tau_1} f_0(t,x(t),u(t))dt + h(x(\tau_1)) \text{ subject to}$$
$$x(\tau_0)=\xi_0,\ \ \dot{x}(t)=f(t,x(t),u(t)),\ \ u(t)\in U(t,x(t)),\ \ x(\tau_1)\in E,$$

where the minimization takes place over pairs (x,u) comprised of arcs $x : [\tau_0,\tau_1] \to I\!R^n$ and *control* functions $u : [\tau_0,\tau_1] \to I\!R^m$ that satisfy the specified constraints. For instance, one could require $x \in \mathcal{A}_n^\infty[\tau_0,\tau_1]$ and $u \in \mathcal{L}_m^\infty[\tau_0,\tau_1]$, interpreting the constraints $\dot{x}(t) = f(t,x(t),u(t))$ and $u(t) \in U(t,x(t))$ as holding for almost every t. The sets $U(t,x(t))$ and E could be defined by equations and inequalities, but we need not go into that. There could likewise be a condition $x(t) \in X(t)$, called a state constraint, but we may regard that as implicit in the description of the region where $U(t,x(t)) \neq \emptyset$.

The challenge of working with such formulations involving controls led to new results like the Pontriagin "maximum principal," which expresses necessary conditions for optimality akin to the Euler-Lagrange equation and Weierstrass condition, but is limited to situations where $U(t,x(t))$ does not actually depend on $x(t)$.

The attitude in the early days of control theory was that problems in optimal control could be regarded as generalizations of problems in the calculus of variations in which the differential equation $\dot{x}(t) = f(t,x(t),u(t))$ provides additional interest and capabilities. In the elementary case where $f(t,x,u) \equiv u$, the differential equation reduces to $\dot{x}(t) = u(t)$, problem $(\mathcal{P}_1)$ comes down to a certain classical problem of Bolza.

The utterly different point of view that we explain next, about how optimal control and the calculus of variations can be seen as fitting together, emerged instead from discoveries made in applying convex analysis to optimization more generally. It focuses on the seemingly much simpler problem model

$$(\mathcal{P}) \qquad \text{minimize } J[x] := \int_{\tau_0}^{\tau_1} L(t,x(t),\dot{x}(t))dt + l(x(\tau_0),x(\tau_1)),$$

where the minimization is over a space of arcs, e.g. $x \in \mathcal{A}_n^p[\tau_0,\tau_1]$ for some $p \in [1,\infty]$. This is called a *generalized* Bolza problem, but at first that may sound paradoxical. Shouldn't $(\mathcal{P})$ rather be called a simplified Bolza problem, since no contraints are apparent in it?

The key source of generality in $(\mathcal{P})$ lies in the great breadth of the class of functions L and l that are now admitted. No longer is it tacitly assumed that these functions are differentiable to whatever order might be deemed useful. To the contrary, they need not be continuous or even finite everywhere. They are allowed to take $+\infty$ as a value, in particular as a means of representing constraints by imposing infinite penalties when they are violated.

For that reason, $(\mathcal{P})$ covers not only classical Bolza problems but also problems in optimal control like $(\mathcal{P}_1)$. To capture $(\mathcal{P}_1)$, one would take

$$l(x_0,x_1) = \begin{cases} g(x_1) & \text{if } x_0 = \xi_0 \text{ and } x_1 \in E, \\ \infty & \text{if not,} \end{cases} \tag{12}$$

and on the other hand

$$L(t,x,v) = \inf\{f_0(t,x,u) \,|\, u \in U(t,x) \text{ satisfying } f(t,x,u) = v\}, \tag{13}$$

where the convention is used that the infimum of an empty set of numbers is $+\infty$. The suppression of control variables in passing in this way from $(\mathcal{P}_1)$ to $(\mathcal{P})$ loses nothing because, once an arc x has been identified as solving $(\mathcal{P})$, a corresponding control function u can be obtained by selecting, for each t at which $\dot{x}(t)$ exists, a vector $u(t)$ in the set

$$\operatorname{argmin}\{f_0(t,x(t),u) \,|\, u \in U(t,x(t)) \text{ satisfying } f(t,x(t),u) = \dot{x}(t)\}. \tag{14}$$

Of course, in validating such an approach a number of technical issues have to be resolved. The functional J in $(\mathcal{P})$ has to be well defined when l and L come out of formulas like those in (12) and (13), and there has to be assurance that, when $u(t)$ is selected from the set in (14) this can be done so as to make the resulting function u belong to the right function space. But such issues have by now all been worked out to satisfaction.

The radical departure from classical theory in investigating generalized Bolza problems $(\mathcal{P})$ without the customary restrictions on L and l started in 1970 in [10] under the alternative assumption that $L(t,\cdot,\cdot)$ and l are convex, which makes the Bolza functional J itself be convex. By that time, researchers in convex analysis were well accustomed to working with infinite penalty representations of constraints (see e.g. [8]) and had realized that many conditions usually stated in terms of gradients could, for convex functions without differentiability, be articulated instead with "subgradients." It was a natural step to try this out in the calculus of variations by studying the fully convex case of $(\mathcal{P})$ and looking for subgradient versions of the Euler-Lagrange condition and the Hamiltonian equations.

In the intervening years, a huge effort has gone into expanding the territory of this type of "nonsmooth" analysis beyond convex functions, so that subgradient conditions in Bolza problems could be established more generally. The doctoral thesis of Clarke [1] gave the first big advance, which eventually led to his book [2]. Many others have also gotten involved, and much research is still ongoing; see [14] and its commentaries.

To explain this major development further, let us review some basic ideas that underlie that literature of convex and nonsmooth analysis, in which notation and terminology are oriented asymmetrically toward minimization. A function

$\varphi : I\!R^n \to \overline{I\!R}$ (where $\overline{I\!R} = [-\infty, \infty]$) has *effective domain*

$$\operatorname{dom}\varphi = \{v \,|\, \varphi(v) < \infty\}$$

and is called *proper* if it is finite on this set, and this set is nonempty (or in other words, if $\varphi(v) < \infty$ for some v and $\varphi(v) > \infty$ for every v). The proper functions are thus the extended-real-valued functions on $I\!R^n$ obtained by taking a finite function on a nonempty subset C of $I\!R^n$ and extending it by giving it the value ∞ everywhere outside of C.

The *epigraph* of φ is the set $\operatorname{epi}\varphi = \{(v, \alpha) \in I\!R^n \times I\!R \,|\, \varphi(v) \geq \alpha\}$. It is a closed set if and only if φ is lower semicontinuous (lsc) and, more importantly for indicating the motivation behind epigraphs, it is a convex set if and only if φ is a convex function. For extended-real-valued functions, the epigraph is a better carrier of geometric information than the graph, because it belongs to $I\!R^n \times I\!R$, a vector space, whereas the graph merely belongs to $I\!R^n \times \overline{I\!R}$.

Minimizing φ over $I\!R^n$ is equivalent to minimizing φ over $\operatorname{dom}\varphi$; this is the *principle of infinite penalization*. The minimum is attained if φ is proper and its sublevel sets $\{v \,|\, \varphi(x) \leq \alpha\}$ are compact. When φ is lsc, these sets are closed, so for them to be compact only boundedness is required, and that can be enforced by a growth condition. The same notions can be invoked for functions in infinite-dimensional spaces rather than just $I\!R^n$, but then a more sophisticated assessment of compactness is needed.

In $(\mathcal{P})$ it is natural to take $L(t, \cdot, \cdot)$ and l to be proper functions on $I\!R^n \times I\!R^n$ that are lsc, but how should $L(t, \cdot, \cdot)$ depend on t? This question, more subtle than might be anticipated, was the first big challenge in putting problem $(\mathcal{P})$ on a firm technical footing. For one thing, integrand $L(t, x(t), \dot{x}(t))$ has to be Lebesgue measurable as a function of t for any $x \in \mathcal{A}^1_n[\tau_0, \tau_1]$, but that can fail under the assumption merely that $L(t, x, v)$ is Lebesgue measurable in t for each $(x, v) \in I\!R^n \times I\!R^n$, or that L is Lebesgue measurable with respect to (t, x, v). The answer, confirmed from many different angles, has turned out to be that, in combination with $L(t, x, v)$ being lsc with respect to (x, v), L should be measurable with respect to the hybrid σ-field generated by the Lebesgue sets in the t argument and the Borel sets in the (x, v) argument. With these properties, L is said to be a *normal integrand*.

The concept of a normal integrand was developed originally in the case of convex functions dependent on an additional parameter belonging to a measure space. Thus, it is an innovation attributable to the rise of convex analysis. A full discussion of normal integrands, their history and properties (including various equivalent definitions), is given in Chap. 14 of [14].

With L taken to be a normal integrand, the functional J in $(\mathcal{P})$ can be given the following rigorous definition when recalling that, for any Lebesgue measurable function β on $[\tau_0, \tau_1]$ that is majorized by a function $\gamma \in \mathcal{L}^1_1[\tau_0, \tau_0]$ majorizing β, the integral $\int_{\tau_0}^{\tau_1} \beta(t)dt$ has a well defined value in $[-\infty, \infty)$. This mirrors the fact that the sum of two numbers in $[-\infty,\infty)$ is a well defined number in $[-\infty,\infty)$. Accordingly, for any $x \in \mathcal{A}^1_n[\tau_0,\tau_1]$, the value of $J[x]$ in $(\mathcal{P})$ is a well defined number in $[-\infty,\infty)$ as long as $L(t,x(t),\dot{x}(t))$ is majorized by an integrable function and $l(x(\tau_0),x(\tau_1)) < \infty$. When these conditions are not met, $J[x]$ is defined as equal to ∞. It is immediate then that

$$J[x] < \infty \quad \Longrightarrow \quad \begin{cases} (x(t),\dot{x}(t)) \in \operatorname{dom} L(t,\cdot,\cdot) \ \text{ for a.e. } t, \\ (x(\tau_0),x(\tau_1)) \in \operatorname{dom} l. \end{cases}$$

The principle of infinite penalization reveals therefore that the constraints on the right are implicit in the minimization in problem $(\mathcal{P})$.

Once $(\mathcal{P})$ has been interpreted in this manner, it is possible to move on to the question of whether a solution arc exists. In view of the classical result in Theorem 1, it may be expected that the right space for this is $\mathcal{A}^1_n[\tau_0,\tau_1]$ and that the convexity of $L(t,x,v)$ in v will be required. Something additional will be needed to reflect the relaxation of the fixed endpoint constraints in $(\mathcal{P}_0)$ by the endpoint term in $(\mathcal{P})$. On the other hand, while the Tonelli growth condition (7) might still be serviceable, it is really too severe for many of the targeted applications. For example, Lagrangian functions L coming from control problems, via (13), have a hard time satisfying it.

Theorem 4. *Let problem $(\mathcal{P})$ be placed in the arc space $\mathcal{A}^1_n[\tau_0,\tau_1]$. Suppose the function l is lsc and the function L is a normal integrand such that $L(t,x,v)$ is convex with respect to v and the following growth condition is fulfilled:*

$$\begin{cases} L(t,x,v) \geq \theta(\max\{0,|v|-\alpha(t)|x|-\beta(t)\}) - \gamma(t)|x| - \delta(t), \\ l(x_0,x_1) \geq \theta_0(x_0) - \theta_1(x_1), \end{cases} \tag{15}$$

for integrable functions $\alpha,\beta,\gamma,\delta : [\tau_0,\tau_1] \to [0,\infty]$ and nondecreasing functions $\theta,\theta_0,\theta_1 : [0,\infty) \to [0,\infty]$ such that

$$\lim_{s\to\infty} \frac{\theta(s)}{s} = \infty, \qquad \lim_{s\to\infty} \frac{\theta_0(s)}{s} = \infty, \qquad \lim_{s\to\infty} \frac{\theta_1(s)}{s} < \infty. \tag{16}$$

If at least one arc $x \in \mathcal{A}^1_n[\tau_0,\tau_1]$ exists with $J[x] < \infty$, then the minimum in $(\mathcal{P})$ is finite and attained.

This result comes out of [9], where in fact the growth condition on L is stated in a slightly broader dual form which allows θ to depend on t. Note that this condition imposes a coercive penalty on $\dot{x}(t)$ to the extent that a bound of the form $|\dot{x}(t)| \leq \alpha(t)|x(t)| + \beta(t)$ is transgressed by the arc x. It covers the Tonelli condition (7) as a very special case, namely the one where the functions α, β, γ and δ are $\equiv 0$. In the growth condition on l, the coercivity in the first argument but counter-coercivity in the second argument could be reversed, i.e., one could assume instead that $l(x_0,x_1) \geq \theta_0(x_1) - \theta_1(x_0)$ without undermining the conclusions.

The cited paper [9] furthermore provides criteria under which the generalized Bolza problem $(\mathcal{P})$ can rightly be considered to reflect, through formulas like (13), an optimal control problem with respect to arcs x and control functions u. The means of recovering an optimal u from an optimal x are furnished there as well.

For a "Lagrangian" function L in the vast class envisioned for $(\mathcal{P})$, is there an associated "Hamiltonian" function H? A powerful and convincing answer to this question is available through convex analysis. To understand it in clearest terms, let us begin by considering again an arbitrary function $\varphi : I\!R^n \to \overline{I\!R}$. The *Legendre-Fenchel transform* of convex analysis assigns to φ, as its *conjugate* the function $\varphi^* : I\!R^n \to \overline{I\!R}$ defined by

$$\varphi^*(y) = \sup_v \{\langle v, y\rangle - \varphi(v)\},$$

and as its *biconjugate* the function conjugate to φ^*, which is

$$\varphi^{**}(v) = \sup_y \{\langle v, y\rangle - \varphi^*(y)\}.$$

Regardless of any assumptions on φ, both φ^* and φ^{**} are convex and lsc. If $\varphi \equiv \infty$, then obviously $\varphi^* \equiv -\infty$ and $\varphi^{**} \equiv \infty$, whereas if $\varphi \not\equiv \infty$ but φ fails to majorize any affine function, one has $\varphi^* \equiv \infty$ and $\varphi^{**} \equiv -\infty$. In the remaining and most important case, where φ is proper and majorizes an affine function (the latter being known to follow from the former when φ is convex), both φ^* and φ^{**} are proper and moreover φ^{**} is the greatest proper, lsc, convex function majorized by φ. Thus in particular, if φ is lsc, proper and convex, the same holds for φ^*, and then $\varphi^{**} = \varphi$. In this manner the Legendre-Fenchel transform induces a one-to-one correspondence within the collection of all lsc, proper, convex functions on $I\!R^n$. A further observation is that if φ is twice continuously differentiable with Hessian matrices that are positive definite, the Legendre-Fenchel transform reduces essentially to the older Legendre transform. The details behind these facts can be found in [8] (or [14]).

It is natural from this perspective to define, for any $L : [\tau_0, \tau_1] \times I\!R^n \times I\!R^n \to \overline{I\!R}$, the associated Hamiltonian $H : [\tau_0, \tau_1] \times I\!R^n \times I\!R^n \to \overline{I\!R}$ by

$$H(t, x, y) = \sup_v \{\langle v, y\rangle - L(t, x, v)\}. \tag{17}$$

Then $H(t, x, y)$ is convex in y and, for any (t, x) such that $L(t, x, v)$ is lsc, proper and convex with respect to v, it obeys the reciprocal formula

$$L(t, x, v) = \sup_y \{\langle v, y\rangle - H(t, x, y)\}. \tag{18}$$

On the other hand, when $L(t, x, \cdot)$ lacks convexity, but at least is proper and majorizes an affine function, the expression in v given by the right side of (18) is the "lsc convex hull" of $L(t, x, \cdot)$ and is proper. In that case we encounter, therefore, convexification of the kind already viewed in the classical context in connection with the Weierstrass condition.

This approach to what the Hamiltonian for problem $(\mathcal{P})$ should be was first proposed in [11], in the days when only convex analysis offered a means by which optimality conditions could be stated without resorting to gradients, which of course for functions like L and l might well not exist. It was essential then to assume $L(t,x,v)$ to be convex not only with respect to v but with respect to (x,v). In Hamiltonian terms, that corresponds to having $H(t,x,y)$ be concave in x in addition to being, as always, convex in y.

It was for convex functions that robust substitutes for gradients were initially developed, so we focus now on that case here. A vector $y \in I\!R^n$ is called a *subgradient* of a convex function $\varphi : I\!R^n \to \overline{I\!R}$ at a point v if

$$\varphi(v') \;\geq\; \varphi(v) + \langle y,\, v' - v\rangle \;\text{ for all }\; v' \in I\!R^n. \tag{19}$$

The set of such subgradients y is denoted by $\partial\varphi(v)$ and is always closed and convex, but perhaps empty. Interestingly $\partial\varphi(v)$ consists of a single y if and only if φ is finite and differentiable at v, in which case $y = \nabla\varphi(v)$. Furthermore, when φ is also lsc and proper, the set-valued subgradient mapping $v \mapsto \partial\varphi(v)$ and the corresponding mapping $y \mapsto \partial\varphi^*(y)$ for the lsc, proper, convex function φ^* conjugate to φ, are the inverses of each other:

$$y \in \partial\varphi(v) \iff v \in \partial\varphi^*(y).$$

An extensive calculus is available for determining subgradients of convex functions in taking advantage of the formulas that may be used to construct such functions. Subgradients can similarly be defined for a function φ that is concave instead of convex by reversing the inequality in (19); the notation $\tilde{\partial}\varphi$ then supplants $\partial\varphi$.

To save words, let us now speak formally of the *fully convex* case of problem $(\mathcal{P})$ as the one in which $L(t,\cdot,\cdot)$ and l are lsc, proper and convex on $I\!R^n \times I\!R^n$, with L being a normal integrand that in addition satisfies the following minor technical condition: $L(t,x,v) \geq a(t,x,v)$ for an expression $a(t,x,v)$ that is affine in (x,v) and summable in t. (The latter is assured for instance under the growth condition on L in (15).) Then in particular, J is a well defined *convex* functional from $\mathcal{A}^1_n[\tau_0,\tau_1]$ to $\overline{I\!R}$ that nowhere has the value $-\infty$.

Consider further in this setting the subsets C_l and C_L of $I\!R^n \times I\!R^n$ that express the two kinds of endpoint constraints implicit in $(\mathcal{P})$; specifically, let C_l consist of all (x_0,x_1) such that $l(x_0,x_0) < \infty$ (i.e., $C_l = \operatorname{dom} l$) and let C_L consist of all (x_0,x_1) such that

$$\exists\, x \in \mathcal{A}^1_n[\tau_0,\tau_1] \;\text{ with }\; x(\tau_0) = x_0,\; x(\tau_1) = x_1,\; \int_{\tau_0}^{\tau_1} L(t,x(t),\dot{x}(t))dt < \infty.$$

Note that these sets C_l and C_L are convex, and that

$$J[x] < \infty \quad \Longrightarrow \quad (x(\tau_0), x(\tau_1)) \in C_l \cap C_L \neq \emptyset. \tag{20}$$

We will want to look also at the *relative interiors* of C_l and C_L, which are their interiors relative to their affine hulls (cf. [8] or [14]). In a constraint qualification introduced below, the issue will be whether not only C_l and C_L themselves have nonempty intersection, as in (20), but also their relative interiors have nonempty intersection.

Theorem 5. *In the fully convex case of problem $(\mathcal{P})$, placed in the arc space $\mathcal{A}^1_n[\tau_0, \tau_1]$, let $x \in \mathcal{A}^1_n[\tau_0, \tau_1]$ be an arc that is feasible in the sense of having $J[x] < \infty$. If there is an arc $y \in \mathcal{A}^1_n[\tau_0, \tau_1]$ satisfying the generalized Euler-Lagrange condition*

$$(\dot{y}(t), y(t)) \in \partial_{x,v} L(t, x(t), \dot{x}(t)) \tag{21}$$

for almost every $t \in [\tau_0, \tau_1]$ and also the generalized transversality condition

$$(y(\tau_0), -y(\tau_1)) \in \partial l(x(\tau_0), x(\tau_1)), \tag{22}$$

then x is optimal. Conversely, the existence of such an arc y is necessary for the optimality of x if the relative interiors of C_l and C_L have nonempty intersection.

Here the Euler-Lagrange condition (21) can be written equivalently as the generalized Hamiltonian condition

$$\dot{x}(t) \in \partial_y H(t, x(t), y(t)), \qquad -\dot{y}(t) \in \tilde{\partial}_x H(t, x(t), y(t)). \tag{23}$$

The sufficiency of the generalized Euler-Lagrange and transversality conditions in this theorem was brought out in [10], and the necessity (much harder to prove) in [12]. The equivalence between the Euler-Lagrange condition and the Hamiltonian condition (23) was demonstrated in [11]. These works further reveal that the adjoint arc y solves a certain *dual* problem, which developed in a certain way out of the convex functions conjugate to $L(t, \cdot, \cdot)$ and l. This dual problem again fits the generalized Bolza format. (One of the many virtues of the notion of normal integrand, incidentally, is that in passing to the conjugate of $L(t, \cdot, \cdot)$ for each t one gets another normal integrand.)

The optimality conditions in Theorem 5 can be elaborated in fine detail when particular structures are given for L and l. A remarkable feature is the way that a very wide range of endpoint formulations and corresponding transversality conditions can all be combined in the single subgradient relation (22). But Theorem 5 only covers problem $(\mathcal{P})$ in the fully convex case.

Can subgradients be defined in for a much broader class of functions in such a manner that the subgradients of convex analysis, as just described, are recovered when the functions are convex, but ordinary gradients are obtained when the functions are smooth? And is it possible that way to derive necessary conditions for optimality resembling those in Theorem 5 even for Bolza problems $(\mathcal{P})$ that are not fully convex? Yes.

A breakthrough in that direction was made in Clarke's thesis [1], as already mentioned. Subsequently the topic were extensively developed further by him and many other researchers, with modifications here and there in the concepts, the most important being the recognition of the need to relinquish certain convexifications that seemed altogether natural at the start of the theory, but later got in the way. In that respect key contributions were made by Mordukhovich; cf. for instance [7].

This larger subject and its history are too much to explain here, but some references can be given. The theory of subgradients in current form is presented comprehensively in the book [14]. Recent achievements in characterizing optimality in terms of subgradient conditions that extend the ones in Theorem 5 can be found for example in [13], [5], [4], [6], [3].

Even today, though, there continue to be new developments in the setting of full convexity, too. In this vein can be mentioned the recent results in the Hamilton-Jacobi theory associated with generalized Bolza problems $(\mathcal{P})$, cf. [15], [16], where the duality theory in [12] has strongly been utilized.

References

[1] F.H. Clarke (1973), *Necessary Conditions for Nonsmooth Problems in Optimal Control and the Calculus of Variations*, thesis, Univ. Washington.

[2] F.H. Clarke (1983), *Optimization and Nonsmooth Analysis*, Wiley.

[3] A.D. Ioffe (1997), "Euler-Lagrange and Hamiltonian formalisms in dynamic optimization," *Trans. Amer. Math. Soc.* 349, 2871–2900.

[4] A.D. Ioffe and R.T. Rockafellar (1996), "The Euler and Weierstrass conditions for nonsmooth variational problems," *Calculus of Variations & PDE]* 4, 59–87.

[5] P.D. Loewen and R.T. Rockafellar (1996) "New necessary conditions for the generalized problem of Bolza," *SIAM J. Control Opt.* 34, 1496–1511.

[6] P.D. Loewen and R.T. Rockafellar (1997) "Bolza problems with general time constraints," *SIAM J. Control Opt.* 35, 2050–2069.

[7] B.S. Mordukhovich (1988), *Approximation Methods in Problems of Optimization and Control*, Nauka, Moscow (in Russian).

[8] R.T. Rockafellar (1970), *Convex Analysis*, Princeton University Press, Princeton, New Jersey (available from 1997 also in paperback in the series Princeton Landmarks in Mathematics).

[9] R.T. Rockafellar (1975), "Existence theorems for general control problems of Bolza and Lagrange," *Advances in Math.* 15, 312–333.

[10] R.T. Rockafellar (1970), "Conjugate convex functions in optimal control and the calculus of variations," *J. Math. Analysis Appl.* 32, 174–222.

[11] R.T. Rockafellar (1970), "Generalized Hamiltonian equations for convex problems of Lagrange," *Pacific J. Math.* 33, 411–428.

[12] R.T. Rockafellar (1971), "Existence and duality theorems for convex problems of Bolza," *Trans. Amer. Math. Soc.* 159, 1–40.

[13] R.T. Rockafellar (1996), "Equivalent subgradient versions of Hamiltonian and Euler-Lagrange equations in variational analysis," *SIAM J. Control Opt.* 34, 1300–1315.

[14] R.T. Rockafellar and R.J-B Wets (1998), *Variational Analysis*, Springer Verlag, Berlin.

[15] R.T. Rockafellar and P.R. Wolenski (2000), "Convexity in Hamilton-Jacobi theory I: dynamics and duality," *SIAM J. Control Opt.*

[16] R.T. Rockafellar and P.R. Wolenski (2001), "Convexity in Hamilton-Jacobi theory II: envelope representations," *SIAM J. Control Opt.*