

***COMPUTATION
AND COMPLEXITY
IN ECONOMIC
BEHAVIOR AND
ORGANIZATION***

***Kenneth R. Mount
Stanley Reiter***

Computation and Complexity in Economic Behavior and Organization

KENNETH R. MOUNT

Northwestern University

STANLEY REITER

Northwestern University



CAMBRIDGE
UNIVERSITY PRESS

CAMBRIDGE UNIVERSITY PRESS

Cambridge, New York, Melbourne, Madrid, Cape Town, Singapore, São Paulo

Cambridge University Press

The Edinburgh Building, Cambridge CB2 8RU, UK

Published in the United States of America by Cambridge University Press, New York

www.cambridge.org

Information on this title: www.cambridge.org/9780521800563

© Kenneth R. Mount, Stanley Reiter 2002

This publication is in copyright. Subject to statutory exception and to the provisions of relevant collective licensing agreements, no reproduction of any part may take place without the written permission of Cambridge University Press.

First published 2002

This digitally printed version 2007

A catalogue record for this publication is available from the British Library

Library of Congress Cataloguing in Publication data

Mount, Kenneth R.

Computation and complexity in economic behavior and organization / Kenneth R. Mount, Stanley Reiter.

p. cm.

Includes bibliographical references and index.

ISBN 0-521-80056-0

1. Economics, Mathematical. 2. Organizational behavior. 3. Computational complexity. I. Reiter, Stanley. II. Title.

HB135 .M744 2002

330'.01'51 – dc21

2001035644

ISBN 978-0-521-80056-3 hardback

ISBN 978-0-521-03789-1 paperback

Computation and Complexity in Economic Behavior and Organization

This book presents a model of computing and a measure of computational complexity that are intended to facilitate the analysis of computations performed by people, machines, or a mixed system of people and machines. The model is designed to apply directly to models of economic theory, which typically involve continuous variables and smooth functions, without requiring an analysis of approximations. The model permits an analysis of the feasibility and complexity of the calculations required of economic agents in order for them to arrive at their decisions. The treatment contains applications of the model to game theory and economics, including a comparison of the complexities of different solution concepts in certain bargaining games, to the trade-off between communication and computation in an example of an Edgeworth Box economy, and to problems of economic organization.

Kenneth R. Mount is Emeritus Professor of Mathematics at Northwestern University, where he has spent 40 years on the faculty. Professor Mount received the E. Leroy Hall Award for excellence in teaching at Northwestern, and he also served as a visiting professor and researcher in France and Argentina. Professor Mount's articles and coauthored articles have appeared in leading refereed publications such as the *Journal of Mathematical Economics*, *Advances in Mathematics*, *Proceedings of the American Mathematical Society*, *Econometrica*, *Journal of Complexity*, and *Economic Theory*. His professional research has been supported by the National Science Foundation, NATO, and UNESCO.

Stanley Reiter is Morrison Professor of Economics and Mathematics in the Weinberg College of Arts and Sciences and Morrison Professor of Managerial Economics and Decisions Sciences in the Kellogg School of Management, Northwestern University, where he directs the Center for Mathematical Studies in Economics and Management Science. He previously served as Krannert Professor of Economics and Mathematics at Purdue University. A Fellow of the American Academy for the Advancement of Science, the American Academy of Arts and Sciences, the Guggenheim Foundation, and the Econometric Society, Professor Reiter is the editor of *Studies in Mathematical Economics* (1987), coeditor of *Information, Incentives, and Economic Mechanisms* (1987), and associate editor of the journals *Economic Design* and *Complex Systems*.

Additional Praise for Computation and Complexity in Economic Behavior and Organization

“This book summarizes the research done over the past two decades by these two pioneers in the theory of bounded rationality in organizations. Anyone who is trying to model economic agents in an organization, and especially anyone who is concerned with the processing of information by organization, will find this an important reference. The models in this book, where agents are information processors within a network, are significantly richer than the conventional model of a single boundedly rational agent as a finite automaton. This approach offers a fresh perspective and tools for modeling computational complexity in an organization, tools that will be very valuable in capturing within a model the limited computational capabilities of both individuals and organizations. The treatment is both insightful and rigorous, making the book particularly suitable to advanced graduate students and researchers.”

– In-Koo Cho, *University of Illinois*

“This book opens a challenging new path in the theory of organization. An organization’s task is to compute a function of certain external variables. A well-designed organization does so quickly. It breaks the task into subtasks, each requiring a unit of time to complete, with the result becoming an input for a higher subtask. Some of the subtasks can be performed simultaneously. The challenge is to arrange the subtasks in a network so as to minimize the total elapsed time until the full task is finished. This is a novel and fruitful way to look at efficient organizations and to compare the difficulty of the tasks they undertake. Some general results are obtained and they are illustrated in a rich assortment of examples, including resource-allocating organizations and games. Contemporary work in the economic theory of organization has many motives and many approaches. Those who seek to move it in new directions ought to make a serious study of this book.”

– Thomas Marschak, *University of California, Berkeley*

“Mount and Reiter overcome the idiosyncratic, problem specific nature of previous models of computation and complexity by developing an approach based around the most common building blocks of economic models: real numbers and smooth functions. On the technical side this powerful innovation opens the way for the use of classical analysis and algebra in analyzing complexity of decision-making. At the same time the use of real numbers and smooth functions makes Mount and Reiter’s approach immediately applicable to standard models in game theory and organizational economics. The detailed examples in the text allow the applied theorist to see this new approach at work in familiar problems without having to master all the theoretical underpinning of this powerful new theory.”

– Kieron Meagher, *University of New South Wales, Australia*

To Bertha, Cynthia, John, Lisa and Greg.

K.R.M.

To Nina, Carla, Frank, Carrol and Miles.

S.R.

Contents

<i>Acknowledgments</i>	<i>page ix</i>
1. Introduction	1
1.1. The Modeling of Computing and Economic Agents	1
1.2. Complexity, Mathematics, and Human Capacities	7
1.2.1. Complexity and Computability	8
1.3. Computing and Economic Organization	9
1.4. Chapter Summaries	15
2. \mathcal{F} Networks	18
2.1. Graphs and Trees	18
2.1.1. The Network Model	25
2.1.2. Conditional Branching	41
2.1.3. Symmetrical Computation	42
3. Networks of Real-Valued Functions	46
3.1. The Leontief Theorem	46
3.1.1. Necessary Conditions	48
3.1.2. An Example	50
3.1.3. Sufficient Conditions	53
3.2. Local Conditions	57
3.3. Computability in Excess Time	65
4. Applications to Economics	73
4.1. Computation with Human Agents	73
4.1.1. Example 1: Reading Handwriting	74
4.1.2. Example 2: Chernoff Faces	77
4.2. Decentralized Mechanisms	79
4.3. The Edgeworth Box Economy	83
4.3.1. Linear Coordinate Changes in the Message Space	89
4.3.2. Linear Coordinate Changes in Parameter Spaces	90
4.4. The Efficient Frontier	91

5. Applications to Games	95
5.1. Bargaining Games	95
5.1.1. Bargaining Games with Quadratic Boundaries	95
5.1.2. The Kalai–Smorodinsky Solution for Quadratic Boundaries	97
5.1.3. The Nash Solution for Quadratic Boundaries	99
5.1.4. Bargaining Games with Cubic Boundaries	103
5.2. Computational Power and Strategic Advantage	105
6. Lower Bounds and Approximations	111
6.1. Revelation Mechanisms	111
6.1.1. Constructions	117
6.2. Finite Approximations	128
6.2.1. Lattice Decomposition of \mathfrak{R}^n	128
6.2.2. A Limit Theorem	130
7. Organizations	137
7.1. Coordination Problems	137
7.1.1. Costs of Information Processing and Efficient Assignments	140
7.2. Two Examples	145
7.2.1. Example 1	145
7.2.2. Example 2	157
7.3. A Formal Model	162
7.3.1. Technology and Production	162
7.3.2. Efficient Production	164
7.3.3. Information, Communication, and Coordination	164
7.4. Structure of Organizations	170
7.4.1. Larger Organizations	170
7.4.2. Revised Cost Model	170
7.4.3. Example 3	171
A Appendix to Chapter 2: Graph Theory	179
B Appendix to Chapter 3: Real-Valued Functions	196
B.1. Uniqueness Results	197
B.1.1. An Example	202
B.2. Leontief’s Theorem	207
B.2.1. An Example	209
B.2.2. Example of the General Leontief Theorem in a Low-Dimensional Case	214
C Appendix to Chapter 5: Application to Games	225
<i>Bibliography</i>	229
<i>Index</i>	235

Acknowledgments

We are indebted to Tom Marschak for extensive comments on an earlier version of this book, a debt that cannot adequately be recognized in the references. We are also grateful to Tim Van Zandt and to Leo Hurwicz for their interest and for helpful discussions. We thank Fran Walker for her help in preparing the manuscript, including some graphics, and for her cheerful patience in dealing with many changes in its evolution. Thanks also to Sondra Fargo for a very helpful editorial contribution.

Portions of Section 6.1 and Appendix B appeared in Mount, K. R. and S. Reiter, “A lower bound on computational complexity given by revelation mechanisms,” *Economic Theory*, 7, 237–266. Copyright Springer-Verlag 1996.

Example 1 in Section 7.1 appeared in Reiter, S. “Coordination of Economic Activity: An Example,” *Review of Economic Design*, 2001. Copyright Springer-Verlag 2001.

Research reported in this book was supported by National Science Foundation Grants IST-8314504, IST-8509678, and IRI-9020270/01.

1 Introduction

1.1. THE MODELING OF COMPUTING AND ECONOMIC AGENTS

This book presents a model of computing, called the modular network model, and a measure of complexity, intended to apply to computations performed by a mixed system of people and machines. Applications of the model to problems in game theory and economics are also presented.

The model has two primitives: a set, usually denoted \mathcal{F} of elementary functions or elementary operations, and a set of directed graphs. The modeler can choose the set of elementary operations to fit the problem at hand. It is assumed that an elementary operation is carried out in one unit of time. Every computation is described as a superposition of elementary operations. Choice of the set of elementary operations permits limitations on computing capabilities to be expressed formally. It also gives the modeler control over the level of reduction of analysis. The topology of the directed graph can also be restricted by the modeler, though in this book we do not do so; instead we assume that any directed graph is available.

These features facilitate the application of the model to human agents and to economic models. Parallel and distributed computing and the dispersion of information among agents are naturally expressed in this model. In each application the modeler can choose the set of elementary functions to suit the problem. The class of elementary operations may include functions of real variables, as well as functions whose domains are discrete sets, as, for instance, functions defined on strings from a finite alphabet. When the alphabet is finite and the elementary functions are Boolean, the model is equivalent to the finite-state automaton model.

Computing with real numbers offers some important advantages in the context of scientific computing (see Blum et al., 1998). It is also relevant to applications in economic theory. Economic models typically use real variables and functions of them. A model of computing in which the elementary operations are functions of real variables allows that model to be directly applied to

standard economic models, without requiring an analysis of approximations in each application. In cases in which the analysis in the economic model is itself numerical, then, as is the case with numerical analysis generally, computation is necessarily finite and typically is carried out by persons who use a finite-state machine. This raises a question about the relationship between the analysis of complexity in our model when real variables are involved and the complexity of the corresponding computation performed by a finite-state automaton that uses a finite alphabet. Instead of analyzing this question case by case, we show (in Chapter 6, Theorem 6.2.1) that the measure of complexity of a function obtained from the continuum model (real variables and smooth functions) is a limit of a sequence of measures of complexity obtained from finite networks, equivalent to sequential machines, computing approximations to the smooth function in question. The limit theorem presented in Chapter 6 shows that when regularity assumptions are satisfied, our model of computation, and the measure of complexity obtained in it, is an idealization of finite computing in the same sense in which measurement with real numbers is an idealization of measurement with rational numbers.

Real computing opens connections to classical mathematics. When computing is done over a finite alphabet, the technical machinery of analysis available is combinatorial mathematics, which is difficult and, in the setting of standard economic models, awkward. In contrast, when the alphabet is the real numbers, or real vectors, the apparatus of classical analysis is available. In this book the class of elementary operations is often taken to be the class of functions of at most r variables, each a d -dimensional real vector, whose value is also a vector in a d -dimensional Euclidean space. In terms of machines, the parameter d can be thought of as the *size* of a register, and the parameter r as the number of registers. When a human being is doing the computing, the parameter d may refer to the number of modalities of perception via the senses. Smoothness conditions are sometimes imposed. A person typically receives visual, auditory, and other sensory inputs simultaneously. In our model, the number of these is d . Further, a person can perceive more than one input at a time, but not very many. In our model the number of these is r . According to a classic paper (Miller, 1956) for a human being the number r is approximately seven. A modular network whose elementary functions satisfy the restrictions imposed by r and d is called an (r, d) network. Much of the analysis in this book concentrates on analysis of (r, d) networks.

How can an (r, d) network represent computations performed by a system consisting of human beings and machines? When the class of elementary functions consists of functions between Euclidean spaces (or between smooth manifolds), it is not obvious that the (r, d) -network model can represent computations performed by human beings, or by a combination of people and machines. However, we can extend the model so that more abstract computations can be reduced to computations with real quantities. For this purpose




Figure 1.1.1.

we introduce the idea of an *encoded version of a function*, and its computation. The formal definition appears in the first section of Chapter 4. Here we make do with an informal description sufficient to understand the example that follows.

Human beings are good at detecting patterns in low dimensions. For instance, people have little trouble recognizing the pattern shown in Figure 1.1.1 as a triangle.

We continue to recognize a triangle even if the corners, and perhaps other pieces of the perimeter, are missing, or if the sides are slightly curved. Recognizing a pattern can be thought of as computing a function that expresses the relation between a subset of the plane, and the act of saying that it is a particular pattern, in our example a triangle. Thus, recognizing a pattern is represented as evaluating a function, ρ , whose value at the subset shown in Figure 1.1.1 is the word *triangle*. The domain of ρ is the set of subsets of the plane, not a Euclidean space, and the range of ρ is the set of English words, or some suitable subset of it, also not a Euclidean space. In the case of pattern recognition by a human being, it is natural to consider ρ to be an elementary function. When it is possible to encode the more abstract domain and range of the function in terms of elements of Euclidean spaces, then, as Figure 1.1.2 shows, evaluating ρ becomes equivalent to evaluating a function, h , whose domain and range are Euclidean spaces. Although h may be complex if evaluated by a machine,




Figure 1.1.2.

when ρ is elementary for a human agent, and the (fixed) encoding and decoding functions are elementary, the function h becomes an elementary function in a modular network model that applies to the system consisting of human agent and machine. Figure 1.1.2 shows the scheme.

More broadly, the standard models of computer science are not convenient for expressing the characteristics of human beings, in particular, economic agents, as information processors. A machine can easily do some things that people find difficult. Other things are easy and natural for people and difficult for machines. Machines have no special difficulty in handling high-dimensional data, whereas human beings do. Human beings are good at recognizing patterns in low-dimensional data, whereas that task can be quite difficult for a machine. Reading handwriting is relatively easy for literate human beings, and quite complex for machines. The widespread and heavy use of computers to generate two-dimensional visual representations of high-dimensional or complex data so that a human being can study them graphically, rather than look at a printout of the underlying function in the form of numerical data, testifies to the special ability of humans to see patterns in low dimensions; it also testifies to the limitations of humans to comprehend the same information presented in another form. The standard models of computing represent these tasks independently of who is to perform them, and the complexity of a task is the same whether the computation is done by a human or a machine. A model in which reading handwriting can be a simple task if done by a human being and a complex one if done by a digital computer would come closer to capturing computation done by persons equipped with machines.

Computers are sometimes used directly by human beings, who determine and control their inputs, and receive and use their outputs. (Some computational devices are internal to other machinery, such as the computers used to control automobile engines and the like. In these cases the interaction with humans is more remote.) In some cases when a human being is an active participant in the course of a computation, with the result that the whole computation depends on both the actions of the machine and of the person, it is nevertheless possible to focus entirely on the analysis of the computation performed by the machine.

But when, as in some economic models, the object of the analysis includes the person, it is not appropriate to separate the person from the machine. When these situations are important, there is need for a model of computation that facilitates analysis of computational problems to be solved by a combination of human and machine resources. The modular network model allows operations that can be carried out by a person to be included in the set of elementary functions. When an encoded version of such an elementary function exists, it can be included in the set of elementary functions of an (r, d) network. Then the analysis of the entire computational task of human and machine can be modeled and analyzed seamlessly. (Analysis of some examples of computations performed by a combination of humans and machines are presented in the first section of Chapter 4. Also see Mount and Reiter, 1998.) Of course, the entire computation might also be represented in one of the classical models of computing, a Turing machine, or a finite-state automaton. Some work along this line is in the literature. The finite-state automaton model and the Turing machine model have been used in game theory (Neyman, 1985), where computational issues also arise. In game theory, as in economic theory generally, it is assumed that players are fully rational and have unlimited computational abilities and resources. These assumptions provide a basis for deep analysis of situations in which the interests of agents may run together, and also conflict to some extent. As in economic theory, there are attempts to weaken the assumptions of full rationality and unlimited computational capabilities. The finite-state automaton model has been used to analyze the complexity of strategies, and the Turing machine model has been used to study the complexity of games. The questions addressed are, for a given game: "How difficult is it for a player to carry out a given strategy in that game?" and: "How difficult is it to solve the game?"¹ Regarding the first question, Aumann (1981) suggested using finite-state automata to measure the complexity of strategies. The measure is the number of states of the smallest automaton that can execute the strategy. (see Gilboa and Zemel, 1989). This is called *strategic complexity*. Infinitely repeated games have been studied in this way (Rubenstein, 1979, 1986; Abreu and Rubenstein, 1988; and Cho, 1995). There is substantial literature in which finite-state automata or Turing machines are used to restrict strategic complexity (Kalai and Stanford, 1988; Kalai, 1996). With respect to the second question, the standard model of computational complexity in computer science, namely the classification of computations based on the asymptotic theory of complexity classes, has been used to analyze the complexity of solving games. (Also see Ben-Porath, 1989; Kalai et al., 1993; Papadimitriou, 1992.) Ben-Porath (1986) showed that there is an advantage to a player in a repeated two-person zero-sum game with patient players from having a larger automaton. The

¹ Kalai (1996) surveys the literature relevant to these questions and provides a bibliography listing the basic papers.

advantage can be considerable, but the automaton must have an exponentially larger number of states. We address these questions in Chapter 5 in the context of the modular network model of computing.

The idea that human cognition is modeled well by Turing machines, or finite-state versions of them, is widely asserted, but it remains controversial.²

There is another issue that deserves a brief comment here. An economic agent experiences his environment directly. How that agent represents his environment in his own mind is usually not observable by others. The economic analyst considering the agent's behavior in his environment "constructs" her own model of the situation in which she can deduce the optimal action of the agent. We apply the (r, d) -network model to the analyst's model in order to analyze the computational complexity of the decisions of the agent. Because the analyst's choice of how to model the agent's situation may have arbitrary elements in it, the measure of complexity might reflect those arbitrary elements, and consequently might be misleading. It is therefore important that the model of computation and the complexity measure it defines do not depend on a particular parameterization of the problem. We want the measure of complexity to be invariant with respect to transformations of the problem that preserve its essential elements. Specifically, we want the measure of complexity to be the same under coordinate changes in the spaces of the variables that define the computation. The methods of analysis in Chapter 3 are coordinate free; in Chapter 4, where the model is applied to analyzing the trade-off between communication and computation in finding the equilibrium of a decentralized message process, we show explicitly that the result is invariant under appropriate coordinate transformations of the underlying spaces.³ There is a second reason why invariance of the measure of complexity under coordinate transformations is important. Changing coordinate systems can implicitly perform computations without our taking explicit notice of them. For instance, to solve a system of linear equations without doing any work, just assume that the coordinate system is one that makes the matrix of the system diagonal. There are also other ways of "smuggling" computation, but these are ruled out in our analyses by regularity conditions.

² Put briefly, the assertion is that the human brain (mind) is a Turing machine. Among others, Roger Penrose does not subscribe to this assertion. His book (Penrose, 1994) is an excellent guide to the issues and to the literature. It is not necessary for our purpose here to take a position on this question. Even if it were the case that the brain is a Turing machine, it would not necessarily be useful in applications to economic agents or economic organizations to model them as executing algorithms in which the elementary steps are evaluations of Boolean functions over a finite alphabet.

³ This is done in steps, beginning with linear coordinate transformations, and ending with general nonlinear coordinate transformations. We show the invariance under linear coordinate transformations explicitly. The proof for nonlinear transformations is tedious and does not lead to any new insight. We therefore refer the reader to Mount and Reiter (1990), where it is presented in full.

1.2. COMPLEXITY, MATHEMATICS, AND HUMAN CAPACITIES

There is a direct connection between the (r, d) -network model and certain classical problems in mathematics. For instance, when we restrict attention to analytic functions, computation of a function F by an \mathcal{F} network, where \mathcal{F} is taken to be the class of analytic functions of two variables, is related to Hilbert's 13th problem. That problem asks whether an analytic function of several variables can be expressed as a superposition of analytic functions of two variables. This is the same as asking whether an analytic function of several variables can be computed by a $(2, 1)$ network whose modules are analytic functions. There is literature stemming from Hilbert's 13th problem that includes contributions by Kolmogorov and others, such as Kolmogorov (1961a, 1961b) and Arnol'd (1963); also see Lorentz (1966). Kolmogorov first showed that each continuous function of n real variables could be written as a superposition of functions of three variables. Arnol'd showed that only functions of two variables are required. Kolmogorov refined this result and showed that each continuous function of n variables could be written as a superposition of continuous functions of one variable and the binary function of addition. In general, the functions required for superposition, besides addition, are not differentiable. The situation is more complicated when the functions in the superposition are required to be smooth. It is known that there are s times differentiable functions of n variables that cannot be written as a finite superposition of s times differentiable functions of fewer variables (see Lorentz, 1966 or Vitushkin, 1961). In this book we work mostly with elementary functions that are twice continuously differentiable, d -vector-valued functions of r variables, each a d -dimensional real vector. Sometimes real analytic functions are used as elementary functions, as in the paper by Mount and Reiter (1998). That paper presents some of the ideas of our model in a less technical setting, and it also presents some applications of the model to human computing, specifically to Chernoff faces in pattern-recognition problems (Chernoff, 1973).

In our model, *computability* and *complexity* are *relative* concepts. The complexity of a given computation can be different depending on the class \mathcal{F} of elementary functions (and, if relevant, the class of graphs permitted). Consider a polynomial of degree ten in one variable, and consider the function that associates the array of its roots with the vector of its coefficients. A person who knows the coefficients and must compute the roots can be in one of at least three situations. She may not have a computer available, or she may have a computer but not have a program for the task, or she may have a computer equipped with a program for computing the roots of polynomials from the coefficients, for example, the program Gauss, or Mathematica. A person without access to a computer, or one using a personal computer that lacks a program for that purpose, would find this a time-consuming task – a complex task. However, that same person using a computer equipped with Gauss or Mathematica could accomplish the task with a few keystrokes. In that case it would be sensible in

many situations to regard the function that associates the roots to the coefficients as an elementary operation, and not pursue the analysis to more detailed levels of the program doing the work. To require that every computation be reduced to a fixed given level of elementary operations, such as Boolean functions over a finite alphabet, results in a model that is awkward to apply to computations done in the context of economic problems by economic agents or by economic theorists, whether they do or do not have access to computing machines.

The idea that complexity is a relative concept is in keeping with practice in mathematics. There the notion of solution to a problem is also relative. For example, what does it mean to solve a differential equation? Among other possibilities, it can mean to find an integral, perhaps expressed in some abstract form, or it can mean to find the solution trajectories numerically. The complexities of these two tasks can be quite different. In mathematics a problem can be considered as solved in cases that are quite different. For instance, a problem might be considered solved if it is shown to be equivalent to another problem for which a solution is known, or to one that is known to have a solution, or to one for which there is an algorithm.

1.2.1. Complexity and Computability

In the \mathcal{F} -network model, the complexity of a function F relative to the class \mathcal{F} (\mathcal{F} complexity) is the minimum over all \mathcal{F} networks \mathcal{N} of the number of sequential steps in the longest path in the network \mathcal{N} . We also refer to this as the *time* it takes \mathcal{N} to compute F . When we take account of the resources used to evaluate elementary functions, the time can vary depending on the assignment of elementary operations to agents (see Chapter 7). If the \mathcal{F} complexity of F is infinite, then F is not \mathcal{F} computable; that is, it is not computable by networks with modules in the class \mathcal{F} . We will sometimes refer to the complexity of F , omitting reference to the class \mathcal{F} when it is clear which class is being used.

Note that the complexity of a function depends only on the class of elementary functions, and not on a particular algorithm that might be used to compute it. In some cases, in which the functions being computed are smooth and the set of elementary functions is appropriately specified, we are able to give lower bounds, sometimes exact, on the complexity of a function F in terms of properties of F alone, that is, independently of the algorithms used. This means that the complexities of, say, different polynomial or rational functions (functions in the standard complexity class \mathcal{P}) can be compared without having to count the steps of the algorithms being used to compute them. In the case of smooth functions between Euclidean spaces (or, more generally, smooth manifolds) the lower bound on computational complexity is determined by the number of variables that the function *actually* depends on (as distinct from the number of its nominal arguments). The calculation of the lower bound does not require