Beyond Intervals: Phase Transitions Lead to More General Ranges

Karen Villaverde[1,*], Gilbert Ornelas[2]

[1]*Department of Computer Science, New Mexico State University, Las Cruces, NM 88003, USA*

[2]*The Computing Sciences Corporation, El Paso, TX 79912, USA*

Abstract

One of the main tasks of science and engineering is to use the current values of the physical quantities for predicting the future values of the desired quantities. Due to the (inevitable) measurement inaccuracy, we usually know the current values of the physical quantities with interval uncertainty. Traditionally, it is assumed that all the processes are continuous; as a result, the range of possible values of the future quantities is also known with interval uncertainty. However, in many practical situations (such as phase transitions), the dependence of the future values on the current ones becomes discontinuous. We show that in such cases, initial interval uncertainties can lead to arbitrary bounded closed ranges of possible values of the future quantities. We also show that the possibility of such a discontinuity may drastically increase the computational complexity of the corresponding range prediction problem.
ⓒ2010 World Academic Press, UK. All rights reserved.

Keywords: interval uncertainty, phase transitions

1 Formulation of the Problem

1.1 Objectives of Science and Engineering

One of the main tasks of science and engineering is to use the current values of the physical quantities $x_1, \ldots, x_n$ to predict the future values y of the desired quantities.

To be able to perform this prediction, we must know how y depends on x_i, i.e., we must know the algorithm $y = f(x_1, \ldots, x_n)$ which transforms the current values $x_1, \ldots, x_n$ into the desired prediction y. Once we know this algorithm, and we know the values of the physical quantities $x_1, \ldots, x_n$, we can then predict y as $y = f(x_1, \ldots, x_n)$.

Comment. In reality, often, the algorithm f represents the actual physical dependence only approximately. For example, in quantum physics, only probabilistic predictions are possible, so any deterministic prediction algorithm is approximate.

In many practical situations, however, the real-life dynamics is known reasonably accurately. In such situations, we can safely assume that the algorithm f describes the exact dependence. This is the assumption that we make in this paper.

1.2 Measurement Inaccuracy

In the above description, we assumed that we know the exact current values of the quantities $x_1, \ldots, x_n$. In practice, however, these values usually come from measurements, and measurements are never 100% accurate. As a result, the measured value $\widetilde{x}_i$ of the i-th quantity is, in general, different from its (unknown) actual value x_i.

Usually, the manufacturer of the corresponding measuring instrument provides us with a guaranteed upper bound Δ_i on the (absolute value) of the measurement error $\Delta x_i \stackrel{\text{def}}{=} \widetilde{x}_i - x_i$ of the i-th quantity.

In this case, after we measure x_i and get the measurement result $\widetilde{x}_i$, we can conclude that the actual value of x_i belongs to the interval $\mathbf{x}_i = [\widetilde{x}_i - \Delta_i, \widetilde{x}_i + \Delta_i]$. In other words, due to the (inevitable) measurement inaccuracy, we usually know the current values of the physical quantities with interval uncertainty.

Comment. Often, in addition to the range $\mathbf{x}_i$ of possible values of x_i, we also know the probabilities of different values $x_i \in \mathbf{x}_i$; see, e.g., [12]. In this paper, however, we only consider the range information.

---

*Corresponding author. Email: kvillave@cs.nmsu.edu (K. Villaverde).

1.3 The Effect of Measurement Inaccuracy on Prediction

In this paper, we assume that we know the exact algorithm $f(x_1, \ldots, x_n)$ which transforms the current values $x_1, \ldots, x_n$ into the desired future value y. Under this assumption, in the idealized situation in which we know the exact values x_i of the current quantities, we could compute the exact value $y = f(x_1, \ldots, x_n)$ of the desired future quantity.

In practice, for each i, we only know the interval $\mathbf{x}_i$ of possible values of x_i. In this case, the only thing that we can conclude about the quantity y is that y belongs to the set

$$\mathbf{y} = f(\mathbf{x}_1, \ldots, \mathbf{x}_n) \stackrel{\mathrm{def}}{=} \{f(x_1, \ldots, x_n) : x_1 \in \mathbf{x}_1, \ldots, x_n \in \mathbf{x}_n\}.$$

1.4 Case of Fuzzy Uncertainty: From the Computational Viewpoint, It Can be Reduced to the Crisp Case

In the fuzzy case, for each possible value of $x_i \in \mathbf{x}_i$, the experts describe the degree $\mu_i(x_i)$ to which this value is possible.

For each degree of certainty α, we can determine the set of values of x_i that are possible with at least this degree of certainty – the α-*cut* $\{x \mid \mu(x) \geq \alpha\}$ of the original fuzzy set. Vice versa, if we know α-cuts for every α, then, for each object x, we can determine the degree of possibility that x belongs to the original fuzzy set [1, 3, 5, 6, 7]. A fuzzy set can be thus viewed as a nested family of its α-cuts.

If instead of a (crisp) interval $\mathbf{x}_i$ of possible values of the i-th quantity, we have a fuzzy set $\mu_i(x)$ of possible values, then we can view this information as a family of nested intervals $\mathbf{x}_i(\alpha)$ – α-cuts of the given fuzzy sets.

Our objective is then to compute the fuzzy number corresponding to the desired value $y = f(x_1, \ldots, x_n)$. In this case, for each level α, to compute the α-cut of this fuzzy number, we can apply the interval algorithm to the α-cuts $\mathbf{x}_i(\alpha)$ of the corresponding fuzzy sets. The resulting nested sets form the desired fuzzy set for y.

1.5 Traditional Assumption: All Physical Dependencies are Continuous

Traditionally, it is assumed that all the processes are continuous; in particular, that the function $y = f(x_1, \ldots, x_n)$ computed by the algorithm f is continuous. It is well known that the range of a continuous function on a bounded connected set, e.g., on the box

$$\mathbf{x}_1 \times \ldots \times \mathbf{x}_n,$$

is an interval. Thus, for continuous functions f, the range $\mathbf{y}$ of possible values of the future quantity y is an interval.

Thus, due to inevitable measurement inaccuracy, we can only make predictions with interval uncertainty. Computing such intervals is one of the main tasks of *interval computations*; see, e.g., [2].

1.6 Discontinuous Dependencies: A Physical Possibility

Some physical processes are discontinuous: e.g., phase transitions. When a water is heated and boils, its density abruptly changes from the density of water to the (orders of magnitude) smaller density of steam.

Of course, all the molecules in water move continuously. So, strictly speaking, the density cannot change abruptly: theoretically, it does continuously change from the density of water to the density of steam. However, for all practical purposes, this transition is so fast that from the prediction viewpoint, we can safely assume that:

- the future density can be equal to the density of water,

- the future density can be equal to the density of steam, but

- the future density cannot be equal to any intermediate value.

1.7   Formulation of the Problem

How does the possibility of discontinuous dependencies change the class of possible ranges? How does it affect the computational complexity of computing these ranges?

These are the questions that we will handle in this paper.

2   How Discontinuities Affect the Class of Possible Ranges?

Let us first describe how the possibility of discontinuous dependencies changes the class $\mathcal{S}$ of possible ranges S. Before we describe this problem in precise terms, let us make some preliminary comments.

Comment. From the mathematical viewpoint, the following result is similar to the results from [9, 10].

2.1   It is Sufficient to Consider Closed Ranges

Let $S \in \mathcal{S}$ be a possible range, i.e., a possible set of values of some physically relevant quantity y. Let us also assume that for this range S, the values $s_1, s_2, \ldots, s_k, \ldots$ are all possible (i.e., $s_k \in S$), and that the sequence s_k converges to a certain number s. In this case, no matter how accurately we compute s, we will always find a number s_k that is indistinguishable from s (and possible). Therefore, it is natural to assume that this limit value s is also possible.

In other words, it is natural to assume that every set $S \in \mathcal{S}$ contains all its limit points, i.e., that it is a *closed* set.

2.2   It is Sufficient to Consider Closed Classes of Sets

A similar requirement can be formulated for different sets $S \in \mathcal{S}$.

Indeed, on the class of all bounded closed sets, there is a natural metric – Hausdorff distance $d_H(S, S')$. This distance is defined as the smallest $\varepsilon > 0$ for which S is contained in the ε-neighborhood of S' and S' is contained in the ε-neighborhood of S. In more precise terms, the Hausdorff distance is the smallest number ε for which

$$\forall s \in S \, \exists s' \in S' \, (d(s, s') \leq \varepsilon)$$

and

$$\forall s' \in S' \, \exists s \in S \, (d(s, s') \leq \varepsilon),$$

where $d(s, s') = |s - s'|$ is the standard distance between the points on the real line.

Informally, it means that if $d_H(S, S') \leq \varepsilon$, and we only know the values $s \in S$ and $s' \in S'$ with accuracy ε, then we cannot distinguish between the sets S and S'.

So, if the sets $S_1, S_2, \ldots, S_k, \ldots$ are all possible (i.e., $S_i \in \mathcal{S}$), and the sequence of sets S_k converges to a certain set S (i.e., $d_H(S_k, S) \to 0$), then no matter how accurately we compute the values, we will always find a set S_k that is indistinguishable from the set S (and possible). Therefore, it is natural to assume that this limit set S is also possible.

In other words, it is natural to assume that the class $\mathcal{S}$ contains all its limit points, i.e., that it is a *closed* class under the Hausdorff metric.

2.3   Towards Formalization of the Problem

We know that continuous dynamics functions are physically possible.

We assume that at least one function describing physical dynamics is discontinuous. In the simplest case, we have a monotonic function of one variable that has a "jump": it continuously grows until some threshold value, then makes a jump, and then again continuously grows. In this case, the range of this variable is not a single interval, it is a union of two intervals.

Thus, we arrive at the following definition.

Definition 1. *A class $\mathcal{S}$ of closed bounded non-empty subsets of the real line is called a* class of ranges *if it satisfies the following conditions:*

   (i) the class $\mathcal{S}$ contains an interval;

(ii) the class $\mathcal{S}$ is closed under arbitrary continuous transformations, i.e., if $S \in \mathcal{S}$ and $f(x)$ is a continuous function, then $f(S) \in \mathcal{S}$;

(ii) there exist a value x_0 and a function $f_0(x)$ which is continuously increasing for $x < x_0$ and for $x > x_0$ and which has a "jump" at x_0 ($f_0(x_0-) < f_0(x_0+)$) such that the class $\mathcal{S}$ is closed under f_0, i.e., if $S \in \mathcal{S}$ then $\overline{f_0(S)} \in \mathcal{S}$; and

(iv) the class $\mathcal{S}$ is closed under Hausdorff metric.

2.4   Main Result

Theorem 1. *The class of ranges coincides with the class of all bounded closed sets.*

2.5   Proof

First, every two intervals can be obtained from each other by a continuous transformation – e.g., by a linear function. Since the class $\mathcal{S}$ contains an interval and it is closed under arbitrary continuous transformations, we can thus conclude that this class contains all possible intervals.

Let us take an arbitrary interval $I = [a_1, a_2]$ that contains a point x_0 inside. We have already shown that this interval belongs to the class $\mathcal{S}$. By the construction of a discontinuous function $f_0(x)$ as monotonic, for this interval, the image $\overline{f_0(I)}$ is the union of two disjoint intervals: $[f_0(a_1), f(x_0-)] \cup [f(x_0+), f_0(a_2)]$.

Now, let us consider unions of two disjoint intervals, i.e., sets of the type $[a_1, a_2] \cup [a_3, a_4]$ with $a_2 < a_3$. Every two sets $[a_1, a_2] \cup [a_3, a_4]$ and $[a'_1, a'_2] \cup [a'_3, a'_4]$ of this type can be obtained from each other by a continuous transformation – e.g., we can take a piece-wise linear transformation $f(x)$ which maps:

- $[a_1, a_2]$ into $[a'_1, a'_2]$,

- $[a_2, a_3]$ into $[a'_2, a'_3]$, and

- $[a_3, a_4]$ into $[a'_3, a'_4]$.

Since the class $\mathcal{S}$ contains one such set and it is closed under arbitrary continuous transformations, we can thus conclude that this class contains all possible two-interval sets.

Let us take an arbitrary two-interval set

$$S = [a_1, a_2] \cup [a_3, a_4]$$

for which the second interval $[a_3, a_4]$ contains a point x_0 inside. We have already shown that this two-interval set belongs to the class $\mathcal{S}$. By the construction of a discontinuous function $f_0(x)$ as monotonic, for this two-interval set, the image $\overline{f_0(S)}$ is the union of three disjoint intervals:

$$[f_0(a_1), f_0(a_2)] \cup [f_0(a_3), f_0(x_0-)] \cup [f_0(x_0+), f_0(a_4)].$$

Now, let us consider unions of three disjoint intervals, i.e., sets of the type

$$[a_1, a_2] \cup [a_3, a_4] \cup [a_5, a_6]$$

with $a_2 < a_3$ and $a_4 < a_5$. Every two sets

$$[a_1, a_2] \cup [a_3, a_4] \cup [a_5, a_6]$$

and

$$[a'_1, a'_2] \cup [a'_3, a'_4] \cup [a'_5, a'_6]$$

of this type can be obtained from each other by a continuous transformation – e.g., we can take a piece-wise linear transformation $f(x)$ which maps:

- $[a_1, a_2]$ into $[a'_1, a'_2]$,

- $[a_2, a_3]$ into $[a'_2, a'_3]$,

- $[a_3, a_4]$ into $[a'_3, a'_4]$,

- $[a_4, a_5]$ into $[a'_4, a'_5]$, and

- $[a_5, a_6]$ into $[a'_5, a'_6]$.

Since the class $\mathcal{S}$ contains one such set and it is closed under arbitrary continuous transformations, we can thus conclude that this class contains all possible three-interval sets.

By applying f_0, we can now conclude that the class $\mathcal{S}$ contains all 4-interval sets, etc., and any finite unions of intervals.

Let us now prove that the class $\mathcal{S}$ contains an arbitrary bounded closed set S.

Indeed, for every ε, we can consider an interval-based approximation S_ε to the set S, by taking the union S_ε of all the grid intervals $[k \cdot \varepsilon, (k+1) \cdot \varepsilon]$ (with integer k) for which $[k \cdot \varepsilon, (k+1) \cdot \varepsilon] \cap S \neq \emptyset$. One can easily check that in the limit $\varepsilon \to 0$, we have $S_\varepsilon \to S$. Thus, from the fact that the class $\mathcal{S}$ contains all finite unions of intervals S_ε, we conclude that the class $\mathcal{S}$ must also contain their limit S.

The theorem is proven.

3   How Discontinuities Affect the Computational Complexity of the Prediction Problem

3.1   Computational Complexity of the Prediction Problem: Interval Uncertainty, Linear Functions

Before we discuss how discontinuities affect the computational complexity of the prediction problem, let us recall the computational complexity of the prediction problem in the continuous cases, i.e., under interval uncertainty. Let us start with the simplest case of a linear function

$$y = f(x_1, \ldots, x_n) = a_0 + \sum_{i=1}^{n} a_i \cdot x_i.$$

In this case, substituting the (approximate) measured values $\widetilde{x}_i$, we get the approximate value

$$\widetilde{y} = a_0 + \sum_{i=1}^{n} a_i \cdot \widetilde{x}_i$$

for y.

The approximation error $\Delta y = \widetilde{y} - y$ of this approximation can be described as

$$\Delta y = \sum_{i=1}^{n} a_i \cdot \Delta x_i,$$

where each input error Δx_i can take any value from $-\Delta_i$ to Δ_i.

The sum $\sum_{i=1}^{n} a_i \cdot \Delta x_i$ attains its largest possible value if each term $a_i \cdot \Delta x_i$ in this sum attains the largest possible value:

- If $a_i \geq 0$, then this term is a monotonically non-decreasing function of Δx_i, so it attains its largest value at the largest possible value $\Delta x_i = \Delta_i$; the corresponding largest value of this term is $a_i \cdot \Delta_i$.

- If $a_i < 0$, then this term is a decreasing function of Δx_i, so it attains its largest value at the smallest possible value $\Delta x_i = -\Delta_i$; the corresponding largest value of this term is $-a_i \cdot \Delta_i = |a_i| \cdot \Delta_i$.

In both cases, the largest possible value of this term is $|a_i| \cdot \Delta_i$, so, the largest possible value of the sum Δy is

$$\Delta = |a_1| \cdot \Delta_1 + \ldots + |a_n| \cdot \Delta_n.$$

Similarly, the smallest possible value of Δy is $-\Delta$.

Hence, the interval of possible values of Δy is $[-\Delta, \Delta]$, and the interval of possible values of the actual value y is $[\widetilde{y} - \Delta, \widetilde{y} + \Delta]$.

The corresponding range can be computed in linear time, i.e., efficiently.

3.2 Computational Complexity of the Prediction Problem: Interval Uncertainty, Quadratic (and Higher Order) Functions

Already for quadratic functions

$$y = f(x_1, \ldots, x_n) = a_0 + \sum_{i=1}^{n} a_i \cdot x_i + \sum_{i=1}^{n} \sum_{j=1}^{n} a_{ij} \cdot x_i \cdot x_j,$$

the problem of computing the exact range

$$\mathbf{y} = f(\mathbf{x}_1, \ldots, \mathbf{x}_n) = \{f(x_1, \ldots, x_n) : x_1 \in \mathbf{x}_1, \ldots, x_n \in \mathbf{x}_n\}$$

over interval inputs $x_i \in \mathbf{x}_i = [\widetilde{x}_i - \Delta_i, \widetilde{x}_i + \Delta_i]$ is, in general, NP-hard; see, e.g., [4, 13].

3.3 What is NP-Hard? A Brief Description

NP-hard means, crudely speaking, that no feasible (polynomial time) algorithm can compute the exact end-points of the range $\mathbf{y}$ for all possible intervals $\mathbf{x}_1, \ldots, \mathbf{x}_n$. (Strictly speaking, this interpretation is only true under the widely believed but still unproven hypothesis that P$\neq$NP.

3.4 What is NP-Hard? A More Detailed Description

To be more precise, it is well known that some algorithms are practically useful, while some other algorithms are computationally useless: even for reasonable size inputs, they require time which exceeds the number of particles in the Universe. This distinction is very different to formalize. Usually:

- algorithms for which the computation time $t_A(x)$ is bounded by some polynomial $P(n)$ of the length $n = \text{len}(x)$ of the input (e.g., linear-time, quadratic-time, etc.) are practically useful, while

- for practically useless algorithms, the computation time grows with the size of the input much faster than a polynomial.

In view of this empirical fact, in theoretical computer science, algorithms are usually considered *feasible* if their running time is bounded by a polynomial of n. The class of problems which can be solved in polynomial time is usually denoted by P; see, e.g., [11].

Not all practically useful problems can be solved in polynomial time. To describe such problems, researchers have defined several more general classes of problems. One of the most well known classes is the class NP. By definition, this class consists of all the problems which can be solved in *non-deterministic* polynomial time – meaning that if we have a guess, we can check, in polynomial time, whether this guess is a solution to our problem.

Most computer scientists believe that NP$\neq$P, i.e., that some problems from the class NP cannot be solved in polynomial time. However, this inequality has not been proven, it is still an open problem. What *is* known is that some problems are *NP-hard*, i.e., any problem from the class NP can be reduced to each of these problems in polynomial time. One of such NP-hard problems is the problem SAT of propositional satisfiability: given a propositional formula F, i.e., a formula obtained from Boolean (yes-no) variables $x_1, \ldots, x_n$ by using &, $\vee$, and $\neg$, check whether there exist values $x_1, \ldots, x_n$ which make this formula true.

NP-hardness of SAT means that if NP$\neq$P (i.e., if at least one problem from the class NP cannot be solved in polynomial time), then SAT also cannot be solved in polynomial time. In other words, SAT is the hardest of the problems from this class.

It is known that all the problems from the class NP can be solved in exponential time. Indeed, for a problem of size n, there are $\leq a^n$ possible guesses, where a is the size of the corresponding alphabet, so we can simply try all these guesses one by one.

3.5 Computational Complexity of the Prediction Problem: General Uncertainty, Linear Functions

We have already mentioned that due to possible discontinuities, the range of possible values of each input x_i is, in general, different from the interval; there may be gaps – specifically, it can be equal to an arbitrary

bounded closed set. In particular, when each gap is the largest possible, this range can be equal to the 2-point set $\{\underline{x}_i, \overline{x}_i\}$.

Let us show that for such 2-point inputs, the problem of computing the range becomes NP-hard already for linear functions $f(x_1, \ldots, x_n)$.

The proof is typical proof of NP-hardness: we reduce a known NP-hard problem to our problem. Specifically, we take the partition problem [11]. In this problem, we are given n positive integers $s_1, \ldots, s_n$, and we must check whether there exist values $\varepsilon_i \in \{-1, 1\}$ for which $\sum_{i=1}^{n} \varepsilon_i \cdot s_i = 0$. We will reduce each particular case of this problem to the following particular case of our problem: $a_0 = 0$, $a_i = s_i$, $\underline{x}_i = -1$, and $\overline{x}_i = 1$ for all i. For the resulting linear function

$$y = f(x_1, \ldots, x_n) = \sum_{i=1}^{n} s_i \cdot x_i,$$

0 belongs to the range

$$f(\{\underline{x}_1, \overline{x}_1\}, \ldots, \{\underline{x}_n, \overline{x}_n\}) = \{f(x_1, \ldots, x_n) : x_1 \in \{\underline{x}_1, \overline{x}_1\}, \ldots, x_n \in \{\underline{x}_n, \overline{x}_n\}\}$$

if and only if the original problem has a solution. The reduction is proven, hence our problem is indeed NP-hard.

Comment. This result was, in effect, proven in [4, 8]. The difference is that in [4, 8], this NP-hardness was proven to justify the use of intervals, while we already know that we have to go beyond intervals, so our NP-hardness is the (inevitable) complexity of an important practical problem.

4   Conclusions

One of the main tasks of science and engineering is to use the current values of the physical quantities for predicting the future values of the desired quantities. Due to the measurement inaccuracy, we usually know the current values of the physical quantities with interval uncertainty. Traditionally, it is assumed that all the processes are continuous; as a result, the range of possible values of the future quantities is also known with interval uncertainty.

However, in many practical situations (such as phase transitions), the dependence of the future values on the current ones becomes discontinuous. In this paper, we have shown that in such cases, initial interval uncertainties can lead to arbitrary bounded closed range of possible values of the future quantities. We have also shown that the possibility of such a discontinuity may drastically increase the computational complexity of the corresponding range prediction problem: e.g., for linear functions, the complexity increases from linear time to NP-hard.

Acknowledgments

References

[1] Dubois, D., and H. Prade, Operations on fuzzy numbers, *International Journal of Systems Science*, vol.9, pp.613–626, 1978.

[2] Jaulin, L., M. Kieffer, O. Didrit, and E. Walter, *Applied Interval Analysis: With Examples in Parameter and State Estimation, Robust Control and Robotics*, Springer Verlag, London, 2001.

[3] Klir, G., and B. Yuan, *Fuzzy Sets and Fuzzy Logic*, Prentice Hall, New Jersey, 1995.

[4] Kreinovich, V., A. Lakeyev, J. Rohn, and P. Kahl, *Computational Complexity and Feasibility of Data Processing and Interval Computations*, Kluwer, Dordrecht, 1998.

[5] Moore, R.E., and W. Lodwick, Interval analysis and fuzzy set theory, *Fuzzy Sets and Systems*, vol.135, no.1, pp.5–9, 2003.

[6] Nguyen, H.T., and V. Kreinovich, Nested intervals and sets: Concepts, relations to fuzzy sets, and applications, In: R. B. Kearfott and V. Kreinovich (eds.), *Applications of Interval Computations*, Kluwer, Dordrecht, pp.245–290, 1996.

[7] Nguyen, H.T., and E.A. Walker, *A First Course in Fuzzy Logic*, CRC Press, Boca Raton, Florida, 2006.

[8] Nogueira, M., and A. Nandigam, Why intervals? Because if we allow other sets, tractable problems become intractable, *Reliable Computing*, vol.4, no.4, pp.389–394, 1998.

[9] Ornelas, G., *Set-Valued Extensions of Fuzzy Logic: Classification Theorems*, Master's Thesis, University of Texas at El Paso, Department of Computer Science, 2007.

[10] Ornelas, G., and V. Kreinovich, Set-valued extensions of fuzzy logic: Classification theorems, In: M. Reformat and M.R. Berthold (eds.), *Proceedings of the 26th International Conference of the North American Fuzzy Information Processing Society NAFIPS'2007*, San Diego, California, June 24–27, pp. 549–553, 2007.

[11] Papadimitriou, C.H., *Computational Complexity*, Addison Wesley, San Diego, 1994.

[12] Rabinovich, S., *Measurement Errors and Uncertainties: Theory and Practice*, Springer-Verlag, New York, 2005.

[13] Vavasis, S.A., *Nonlinear Optimization: Complexity Issues*, Oxford University Press, New York, 1991.