Journal of Official Statistics, Vol. 15, No. 3, 1999, pp. 373–393

Towards Systems of Social Statistics – Some Principles and Their Application in Statistics Canada

Ivan Fellegi and Michael Wolfson[1]

The main proposition of the article is that the further development of social statistics should be based on a careful review of the fundamental needs of social policy. In turn, these needs should be based on the character of social programs. They are: outcome oriented; should be based on an understanding of causal relationships; involve long lead times between causes and effects; are sectoral; most often focus on distributional issues; involve more than government. The article then explores the implications of these features of *social policy* for the *design of supporting data systems*. It emphasizes the need for conceptual frameworks, feeder data systems including traditional cross-sectional surveys as well as hybrid data systems, longitudinal surveys, the importance of clean micro data bases, and microsimulation modelling. It also outlines an implementation strategy which stresses collaboration among government statisticians, departmental policy analysts, and academic personnel.

Key words: Social policy; conceptual frameworks; data systems; longitudinal surveys; micro data bases; microsimulation modelling.

1. Background – A Long Search

There is a rich but generally unsuccessful history of efforts to develop a unified and coherent system of social statistics for a nation. Ideally, we seek an integrated system of social statistics:

- whose measures shed light on the state of society
- which is based on a defensible understanding of dynamic relationships among the main factors or variables, and
- therefore provides a framework for the study of causes, effects and outcomes, and the roles and possible impacts of policies.

Such a system would parallel, for social statistics, the success of the System of National Accounts for economic statistics – both in terms of its utility for users as an integrated set of information, and for producers of statistics as an organizing framework for the development of comprehensive economic statistics.

[1] Statistics Canada, 26A R.H. Coats Bldg, Ottawa, K1A 0T6, Canada. E-mail: fellegi@statcan.ca
Acknowledgments: A version of this article was presented at the 51st Session of the International Statistical Institute, Istanbul, August 1997.
We are indebted to an editor and several referees for valuable comments. Otherwise, we remain responsible for any errors or infelicities.

© Statistics Sweden

In the absence of this ideal, users and producers of social statistics have sought a more modest goal: the achievement of consensus on a set of social indicators that could serve as a focus for public policy debates. As a minimum, these statistical indicators would provide at least basic descriptive and monitoring information spanning the range of major social concerns. While they would likely not reflect the dynamic relationships among the variables involved, nor the impacts of possible government policies, they would nevertheless be reasonable and understandable both individually and collectively. Unfortunately, not even this approach has been successful. Many national and international organizations tried unsuccessfully to identify a small number of commonly agreed indicators, which jointly would illuminate the state of social issues and could become a focus for enlightened public debate.

Yet current interest in social issues, hence in social statistics, is increasing again. One reason is continuing stagnation in household incomes. Another driving force has been pressures on governments to cut expenditures in order to achieve more balanced fiscal positions. Many of these cutbacks are made in social programs; but to the extent that empirical foundations to judge their impacts are lacking, the decisions must be made largely in the dark. Given the potential long-term impacts of these cuts for individuals and, indeed, for the fabric of society, there are increasing demands for ''evidence-based decision making'' – policy decisions that take the fullest account of their implications.

The theme of this article is that the time has come to consider renewed efforts directed toward developing coherent systems of social statistics. Such efforts are most likely to be successful if they build on the following premises:

- Development of social statistics must be openly policy relevant;
- Since social policy is itself sectoral, systems of social statistics should also be constructed along sectoral lines;
- Statistical offices must support such development with a critical mass of analytical capacity – both for substantive reasons, and as gateways to external policy and academic analysts;
- Longitudinal surveys should be a central component of the supporting data development strategy, but at the same time such surveys should provide a continuing stream of more immediate benefits in order to maintain funding support;
- Data developments associated with new statistical systems should exploit the benefits of explicit micro data foundations;
- Statistical systems should build on a range of data sources, and consider the use of models such as microsimulation models as a key methodology for integrating and creating coherent structures for such data; and finally
- Productive collaboration is needed among national statistical offices, not least to seek means of building consensus on measures of key social outcomes.

The balance of this article provides the underpinnings for these conclusions. In the next section we explore some basic features of social policy which should have a determining impact on the evolution of social statistics. We then propose a broad strategy for action. The last section explores, through examples drawn from Statistics Canada, some specific implications for the data systems that we need to develop.

2. The Character of Social Policy

Social policy is more than the design or analysis of programs by officials in government ministries. It is fundamentally concerned with activities in which citizens of a nation engage collectively. It therefore involves a range of interested parties, including the general public, researchers, public and private sector managers, and politicians.

For social policies to be effective, it is clearly preferable that they be grounded on facts and information – a large part of which must necessarily be supplied by social statistics (Fellegi 1995). Conversely, if they are to be relevant, the character of social statistics must derive from the character of the social policies and programs which they are designed to illuminate. This section reviews those basic characteristics of social policy which have, or should have, a major impact on the nature of social statistics.

2.1. Interest in outcomes

To begin, social policy is about achieving desired outcomes – adequate incomes, a sufficient number of jobs, high levels of health status in the population, improved levels of educational attainment, reduced crime rates, preservation of cultural identities. There is, therefore, a clear interest in monitoring the achievement of these objectives, hence a continued need for the standard type of monitoring information on the levels and trends for key outcome variables.

However, one can only monitor outcomes if there is at least an implicit social consensus on just what are the desirable outcomes. For the examples just given, there is indeed such a social consensus: policies should aim at reducing unemployment and crime rates, increasing average incomes, and avoiding an extreme polarization of incomes. Subject to the availability of resources, statistical systems by and large do a reasonable job of *monitoring* these outcomes – even if we do considerably less well in *explaining* them.

Greater difficulties face us in domains where there is no social consensus on the operational meaning of the desired outcomes. While we all want a better education system, there is no agreement as to whether and to what extent this should relate to any of the following: the acquisition of basic knowledge, an acquired ability for learning, skills needed by the labour market, or social attitudes that we accept as desirable or essential for citizenship. There are equal difficulties in the field of health. While we know that ''better health'' is the general objective, should the emphasis be on reducing health problems defined in terms of clinical disease, or functional limitations, or the handicaps society may or may not impose on those with a given disability (WHO 1980), or should health be viewed more positively in terms of resilience and as a resource for living? There is no consensus on which of these definitions should be given priority, let alone how each should be operationalized.

2.2. The need for information on causal relationships

To the extent that we succeed in monitoring outcomes, the data measure, *post facto*, the extent of progress towards the objectives of social programs. But they provide no insight about whether a given outcome is attributable to a social program or to a multiplicity of other possible causes. Nor can this kind of data provide guidance about the possible effects

of alternative policies as they might affect various population groups, nor the effects of other factors external to the policies in question. The reason, simply, is that any discussion of impacts depends fundamentally on an ability to assess alternative scenarios – what the outcome would have been in the absence of the given policy. In other words, discussing impacts presupposes some empirically validated theory linking causes and outcomes. Thus, to the extent that intelligent debates about policy options must inevitably focus on their intended and likely effects, the usefulness of statistical information depends directly on its ability to shed light on causal relationships.

Unfortunately, causal pathways between policies and desirable outcomes are seldom direct. For example, it is unlikely that any single government policy can directly and permanently move female-headed single parent families out of poverty. But government programs can provide income support, subsidised childcare, and job-related training. Statistical information can facilitate the search for effective policies if it can point the way to those policies and interventions that have a high probability of leading to long term improvement, for example of incomes.

There is increasing recognition, therefore, that social statistics should go beyond description and monitoring. They should be constructed as a system, knit together by a conceptual framework that reflects the best research knowledge and evidence on causal relationships, and the linkages between policies and outcomes. We shall come back to this point.

2.3. Long lead times between causes and effects

A significant difficulty in linking policies with their intended outcomes is that in social domains, cause-effect relationships often take a long time to manifest themselves. The delays can range from a few months (e.g., between the time of training and a new job) to decades (e.g., the time between exposure to carcinogens and the diagnosis of cancer). The connections can also be complex and cumulative (e.g., unfortunate circumstances in early childhood leading to poor school performance, and eventually to poor labour market experience; e.g., Power, Manor, and Fox 1991). The elapsed time between causes and effects, and the likelihood of intervening factors, pose major challenges for the specification of conceptual frameworks, for the validation of underlying hypotheses, and for the planning of data systems.

2.4. Policy is sectoral

In spite of the increasing recognition of the inter-relatedness of social phenomena, policies remain largely sectoral. There are separate ministries and other organizations dealing with policies in the fields of human resources and labour markets, welfare, health, education, justice and crime, culture, etc. While there is some history of governments creating ''super-ministries'' in an attempt to reflect inter-relatedness in the development of policies and programs, they tend to be umbrellas for sector-specific decision making.

And this is reasonable: social phenomena are simply too wide-ranging, so that comprehensive approaches attempting to take account of all factors might lead to paralysis. This sectoral orientation within government is reinforced by the fact that each of the major domains of social policy has its specific constituencies – who provide primary triggers,

audiences, and spokespersons in public policy debates within their respective areas of interest. The sectoral character of most social policy development in government is further reinforced by the general absence of substantial inter-disciplinary academic work.

For example, when a health department is considering policy options, it is principally interested in evaluating the health effects of alternative policies, and the reactions of key constituencies such as health care providers and disease-related charities, and not their effects on employment or education. At the same time, the intellectual base also tends to be sectoral – the health department has a far wider clinical, bio-medical, and health services research literature to draw upon in thinking about its policies, than it does with regard to the inter-relationships among health, employment, and educational factors, even though there is high quality evidence that all are importantly related (e.g., Evans, Barer, and Marmor 1994). Analogous considerations apply to departments of education, social security, and employment, etc.

2.5. Social policy concerns are distributional

Much of social policy concerns one or other of two broad kinds of government function. One is providing services, or in-kind benefits, for example public education, publicly funded or managed health care, and the legal and protective services of the courts and police. The other function is providing income transfers, often as selective benefits to specific needy or deserving groups, for example the unemployed, seniors, and children in poor families. Income transfer policies and programs clearly require distributional information in order to identify the relevant groups to target, and subsequently to monitor program impacts. For example statistics on the distribution of income, wealth and consumption are needed to identify the ''poor'' (however defined), and to assess the effect of income support programs in reducing poverty among particular target groups (children, the elderly, single parent families, etc.).

But effective management of the first group of ''in kind'' or social service programs also requires distributional statistics. In their efforts to reduce deficits, many governments are cutting back the funds available for both the education and the health care systems, and are seeking ways to redirect funds in order to attain objectives more effectively. There are numerous ways of accomplishing this and choices must be made. Should the number of hospital beds be reduced? If so, how much should community care resources be increased not only to compensate but also to improve outcomes? How should the government alter the mix of public and private provision of prescription drugs, both to contain overall costs, and to assure equitable access? Every combination of choices has different effects on the health, out of pocket expenses, and informal care burdens of families according to their incomes, types of health problems, family situation, and community circumstances.

In both the health and the education sectors, current statistical information in most countries relates primarily to financial and operational characteristics – expenditure data related to various segments of these sectors (e.g., elementary-secondary versus post-secondary), and operating characteristics such as student-teacher ratios, graduation numbers, hospital discharges, etc. Such data are fine as far as they go, but they are insufficient for current policy needs. Statistical information can only be truly relevant if it helps to illuminate the likely effects of alternative choices *on different population groups*. In the

case of health, for example, this will require data on the joint distribution of individuals by health status, health service utilization, geographic neighbourhood, and socio-economic status. Indeed, the data increasingly seen as essential to understanding the effects of the health care system are individuals' trajectories through this system, linked both to data on their family and community circumstances, and to their health status. In summary, the needed outcome-oriented statistics are essentially distributional; they will typically involve micro-data that track a sample of individuals over time; and they require a sufficient range of information to assess how individuals in different circumstances are interacting with the programs in question.

2.6. Social policy is more than government

Before the advent of the welfare state, individuals had to rely exclusively on their own networks of support: first their immediate and extended families, then their employers, friends, and immediate community. An initial objective of the welfare state was to provide one more source of support when all else failed. Over time, this premise has evolved in both depth and breadth. A number of government programs have become major, if not primary, sources of support for an increasing number of people and families; and in some cases the idea of benefits of last resort evolved into programs with universal benefits. However, partly as a result of the budgetary squeeze faced by most governments, the orientation is changing, with targeting and selectivity becoming increasingly attractive in a range of instances.

This broad change in policy direction has rekindled interest in the roles played by more traditional sources of social support other than the state: members of extended families, friends, nonprofit voluntary organizations, and the like. But by and large, statisticians are not well placed to support such a changing policy orientation. In making a new start (or in trying to renovate) social statistics, we must begin the difficult task of reflecting on the role of the support networks available to individuals. Families are particularly important, not only as potential sources of support to individual members, but also because their internal dynamics play a major role in the socio-economic welfare for family members (e.g., a divorce often leads to less income for a spouse, adult children leaving home leads to more disposable income but possibly less informal support). Moreover, developing this kind of statistical information has become much more difficult with the increasing complexity of family relationships: higher rates of marriage breakdown and reformulation leading to blended families and joint custody arrangements, more common-law unions, and more informal care for frail elderly parents provided by adult children not living in the same household. These significant relationships often extend well beyond the group of people living together in a dwelling, the conventional unit for data collection.

2.7. Summary

The paragraphs above have indicated several broad requirements for new or renewed approaches to social statistics. Specifically, they have highlighted the importance of:

- Developing statistical measures of outcomes relevant to major areas of social policy.

- Pursuing a sectoral approach that reflects the institutional structure (current or emerging) of social policy.
- Capturing the impact of factors that shape outcomes, including, but not limited to those which are amenable to policy interventions.
- Recognizing the special importance of long lead times between causes and their effects.
- Giving pride of place to information which can shed light on dynamics, distributions, and the complexities of modern institutional forms like the family.

3. Implications for Social Statistics – Strategy

We now turn to outline some possible strategies. These are illustrated with reference to health and education, two of the largest sectors where progress in redeveloping systems of social statistics is particularly challenging.

3.1. The sectoral approach: A non-fatal compromise

Taking the System of National Accounts as a model, the notion of a single system of social statistics, such as that proposed by Sir Richard Stone (Stone 1973; UN 1975), appears ideal. In comparison to this ideal, a sectoral approach represents a compromise. However, a sectoral approach to the development of social statistics is practical for a national statistical agency, since it would match the external intellectual and organizational landscape which, as argued earlier, is itself sectoral. Thus, the medium-term objective should not be a single over-arching system of social statistics, but rather a collection of sector-specific statistical systems.

Sectoral systems, while simpler than some grand over-arching model, need not represent a fatal compromise from the viewpoint of *inter*-sectoral relationships, provided they are conceived broadly enough. Given consensus regarding the outcome variables (the *explicanda* or ''dependent'' variables) within a given sector, one can and should take account of all the explanatory (or ''independent'') variables that are likely important, from whatever sector, and whether or not they are amenable to policy. To illustrate at a very general level, when thinking about health outcomes, at least in the medium term, education might be considered an exogenous variable. Conversely, when the focus is on a system of education-related statistics with educational attainment as an outcome, child health might be an exogenous variable (which, however, should it turn out to be an important *independent* variable in the education model, could be targeted as a policy variable that can be influenced, e.g., through school-provided meal programs).

Since a sectoral approach need not lead to fatal conceptual compromises, its *relative* simplicity becomes a major consideration. Even so, the conceptual and data difficulties remain quite overwhelming.

3.2. Relevant outcome measures (plural!)

As discussed earlier, the primary interest of *social policy* is to modify outcomes. It follows that the dual priority for systems of *social statistics* should be: (a) the development of measures aimed at monitoring key outcomes, and (b) illuminating their most important correlates, particularly those with causal significance.

Agreement at this general level is far easier, however, than at the level of detail required to develop specific statistical indicators. At the same time, we are frankly sceptical of the ''holy grail'' of a single all-encompassing outcome indicator. The wide attention given to the United Nations ''Human Development Index'' (perhaps only in Canada because of our continued number one ranking) indicates the hunger, at least in the media, for such overall measures. Such summary measures are useful as rallying points for public discussion and for mobilizing support for public policies. However, these simple indices are not needed by serious analysts. Indeed, what serious economist looks to the System of National Accounts for only a single overall figure to describe recent economic events?

We do not foresee the development of useful indices of overall well-being which would bear logical and defensible relationships to their constituent parts – GDP (with its unifying thread of a common numeraire) bears to components such as consumer spending, exports and investment. It is truly daunting even to think of meaningfully combining measures of such disparate aspects of well-being as income, health, and educational attainment (the three main components of the Human Development Index; UNDP 1997), as well as other obviously relevant factors such as risks of victimization and availability of leisure time.

Instead, sector-specific outcome measures are a more appealing objective, one that might be approached developmentally. For example in the health domain, an early example of an overall health status indicator is ''disability-free life expectancy'' (DFLE), i.e., life expectancy modified to remove the years that might be spent in disability. This measure has the advantage of relative simplicity, and modest data requirements. But it suffers from several limitations: it recognizes only a binary health outcome at the individual level – healthy or disabled, rather than a continuum; it strongly depends on a threshold definition of disability; and it implicitly attaches zero value to time spent ''disabled.''

As a result of these limitations, more sophisticated alternatives have been proposed (e.g., disability-adjusted life years or DALYs – World Bank, 1993; health-adjusted life expectancy or HALE – Roberge, Berthelot, and Wolfson 1997). These summary outcome measures, which take account of a number of dimensions of health status, aggregate information over different kinds of health states and over individuals. But such aggregation necessarily employs some kind of weighting. How should one weigh a year of life spent with cancer versus AIDS, or being blind versus missing a limb? One response is to base the weights on a representative sample of people who are asked a structured series of questions to elicit their rankings or preferences over a carefully specified set of health states. But this is also problematic: the kinds of questions required are cognitively complex for respondents, and the results can be sensitive to subtle aspects of the way the questions are framed (Nord, Richardson, and Macarounas-Kirchmann 1993).

Apart from these empirical difficulties, there are also questions of public perception. If a summary health index were successful, it would be used as a measure of progress of the health system. In turn, rational decision-making would tend to allocate resources to those interventions most likely to improve population health *as measured by the index*. Such an allocation process might, for example, point to providing relatively fewer resources for care of heart disease patients, and more for arthritis or dementia, based on the respective burdens of these diseases. While policy decisions are never made on the

basis of a single indicator, even the possibility would naturally draw attention to the weighting scheme or judgements implicit in the index. This could become quite controversial and could affect the perceived objectivity of statistical offices associated with such attempts. This is already the case with ''poverty lines.''

The main challenges are to recognize that for social statistics there are, unavoidably, at least *several* relevant major outcome indicators in each sector, and that there is a need to build sufficient public consensus to support statistical measures of these key social outcomes. One strategy, the one we feel is most likely to be productive, is to begin producing, perhaps experimentally, a range of possible outcome indicators for each social sector. This would serve a dual purpose; it would draw increased attention to outcomes, even if imperfectly conceptualized and measured; and it could trigger a public debate about what *should be* the socially accepted outcome measures.

As countries seek increasingly to compare themselves to one another in the social domain, international comparability of social statistics has become more important. As a result, it would certainly be desirable that the conceptual development and specification of new summary outcome indicators in social domains like health and education take place under the auspices of international organizations.

3.3. Connecting outcomes and policies

High quality outcome indicators are an essential first step towards the development of useful systems of social statistics. However, a major goal of observers and analysts is to discern the relationship between outcomes and public policy interventions, the so-called ''policy levers'' available to government. Useful statistical data systems must therefore attempt to shed light on these.

Our current body of social statistics tends to reflect traditional policy levers: expenditures, human resource inputs, and fixed capital investment. However, at a time of cost containment and doing more with less, simply monitoring expenditures and other inputs is no longer sufficient. As a minimum, we are asked to shed light on the relatively simple issues of efficiency: i.e., to explore the *direct impact* of marginal changes and reallocations, e.g., within the formal health and education systems. To illustrate, the unit costs of a student-year of elementary or secondary schooling need to be regularly estimated. Moreover, these costs should be available by type of resource input and by individual school, so that productivity ''benchmarking'' comparisons across discrete organizational entities are feasible. This allows high and low unit cost outliers to be identified, which then provides the impetus for further exploration of the determinants of these per pupil costs. Similarly, in the health domain, cost differentials among hospitals, adjusted for the number of patients, the kinds of care or services provided, and the severity of their conditions, have begun to serve as important starting points for a better understanding of ''cost drivers.''

But there are weaknesses in the usually assumed connections between the traditional inputs and the desired outputs. For example, there is growing recognition that the factors determining rates of successful surgery or high rates of student achievement depend on more than the physical or financial volumes of resource inputs. Class size or pupil-teacher ratios, for example, turn out not to have a great influence (Hanushek 1981), while ''school

climate'' (Willms 1992) appears to be an important factor for student outcomes. Similarly, choice of ''appropriate'' surgical procedures and follow-on community services is recognized as central both to cost control and to improved patient outcomes in health care. Social statistics relevant to contemporary policy must include information about the *indirect and longer-term* effects of these additional attributes, not least because this is where the focus of policy has turned; these are important contemporary policy levers.

This leads us to our second set of problems. High levels of population health and educational outcomes (to continue with our pair of examples) depend on far more than the activities of hospitals and schools, even when we broaden our conceptions to include school climate and community services. For example, data from new statistical initiatives at Statistics Canada show that entry into nursing homes depends not only on the severity of specific clinical conditions, but also on social factors such as availability of kin support, and wealth (Tomiak, Berthelot, and Mustard 1997). In the area of education, early results from a new longitudinal survey of children (see below) enable us to explore the importance of non-school factors such as ''home curriculum'' as a determinant of educational outcomes.

The implication is that systems of social statistics should be based on broad views of the relevant causal factors and the ways they relate to one another – in other words, on a conceptual framework.

3.4. Conceptual frameworks

We believe that the requirement for connecting outcomes to policies leads directly to the need to develop conceptual frameworks. There is no widely shared agreement about what comprises an adequate conceptual framework for a statistical system. As a general definition, we regard such a framework as:

- An evolving schematic representation,
- of significant empirical interactions,
- together with the description of a set of feasible data bases,
- capable of illuminating these interactions.

A conceptual framework regarded as a schematic representation of significant empirical interactions is not, in and of itself, a theory. But it should be grounded in some theory.

The System of National Accounts (SNA), the best-articulated statistical framework, is not a theory of how national economies function. It is, rather, a complete and carefully structured accounting of monetized transactions. While the origins of the SNA are deeply rooted in Keynesian theories, the SNA itself is not uniquely wedded to them. It also supports analyses motivated from monetarist and other theoretical perspectives.

Similarly, a conceptual framework for a system of health or education statistics should draw on, or be motivated by, relevant theories – in the sense of reflecting the main interactions that, according to accepted theories, connect the variables that are of interest for the given social domain. Indeed, it would be ideal if the framework could accommodate most of the major theoretical perspectives in the given domain. In this way, the framework could avoid the criticism of being captive to one particular perspective; it could offer

benefits to a range of theoretical constituencies; and ideally it could begin to offer the empirical basis for judging amongst competing theories.

One illustration in the health domain of alternative perspectives, if not competing theories, is the significant difference in emphasis between clinicians with a bio-medical perspective, and public health practitioners whose concerns include individuals' functional limitations and the availability of community supports. A good conceptual framework will accommodate the full range of valid perspectives. An example of such a framework is the Health Information Template (Wolfson 1992) developed as part of Canada's National Task Force on Health Information (Wilk 1991).

A conceptual framework is not just a schema – for example a ''box and arrow'' diagram reflecting one or another theory – showing how a network of empirical factors (each in its own ''box'') interact (i.e., the network of arrows showing linkages among the various ''boxes''). It should also specify a tightly coupled measurement system. At the least, it should sketch a feasible set of data collection processes. Furthermore, these data systems ought to be rich enough to reflect all important or competing perspectives. For example, student-level data should encompass both standardized achievement results and formal qualifications. Similarly, individual-level health data should ideally include information on both clinical diseases and functional limitations.

The conceptual framework then assumes its full role in an evolving interaction with realizations of these data systems. On the one hand, conceptual frameworks can guide the evolution of data systems (e.g., the difficult choices of what groups of questions to place on a given survey), as they endeavor to quantify the interactions displayed by the framework. On the other hand, the data systems should also have a profound influence on the evolving conceptual framework: they should lead to the elimination of insignificant or irrelevant relationships, and to the possible further elaboration of the most important ones. In exploring and trying to represent relationships, a key role, data analysis might hint at altogether new relationships, i.e., at possible new theoretical insights.

In order to maintain policy relevance, the conceptual framework for one or other system of social statistics should be explicit about those relationships that connect the social outcomes of interest to a range of causal factors: those that qualify as direct policy levers, those amenable to policy intervention (even if not currently used as such), as well as other important causal factors beyond the ken of public policy. At the least, this requires explicit notions of social processes, along with their ''inputs'' and ''outputs.''

A conceptual framework might well need to recognize a hierarchy of levels at which outcomes of individuals are significantly affected. For example, in the education domain, key decision-makers or influences include teachers, principals, school boards, and provincial education ministries, as well as parents and community organizations. Correspondingly, the data collections associated with the framework should ideally include data at all these levels. The higher level data, however, should not stand in isolation from data at other levels. Rather, they should explicitly recognize the relevant connections (e.g., children to classrooms, classrooms to schools, etc.), to the extent that these connections are important for understanding.

The distributional focus of social policy implies yet another feature of a desirable conceptual framework. Since social policies are, ultimately, about individuals, it follows

that conceptual frameworks for domains of social statistics like health, work, income, and education should be grounded in individual-level data (micro-data).

Finally, the very notion of process in social domains like health, work and education, as mentioned before, often involves long lead-lag relationships. Thus, the relevant conceptual framework must be designed to absorb data on, and reflect back as summary statistical measures, information on processes like health care, learning, and development over time. These lines of argument lead to the following general *desiderata* for conceptual frameworks for systems of social statistics:

- Put individuals (e.g., as students and learners) at the centre;
- Encompass multiple levels (e.g., the student, classroom, school, community, etc.);
- Include as a core idea that inputs lead to outputs via social or institutional processes; and
- Include explicit micro level dynamics, such as the ''trajectories'' of individuals over time and through various social programs and institutions (e.g., school, health care).

Again, let us emphasize once more that we are talking about sectoral frameworks, and sector-specific systems of social statistics. It should be clear that, even with this simplification, the task remains almost overwhelmingly complex.

3.5. Need for collaboration

The type of statistical evolution outlined here clearly represents a very ambitious undertaking. Public policy needs, typically sectoral, justify the effort and government expenditure that is involved. To the extent that these statistical efforts are successful, public policy benefits through access to explanatory models validated (at least eventually) by rigorous analysis of solid statistical information, rather than being limited to intuition and ideology.

The major effort required cannot succeed without the intensive tripartite collaboration of the sectoral policy department concerned, the statistical office, and the social science community. The policy department's support must take at least three forms: (1) moral support, (2) direct financial support (or alternatively support for the financial requirements of the statistical office), and (3) most importantly collaboration in articulating policy requirements, using the resulting statistical data and related policy models, and providing feedback regarding their further desirable evolution.

Social science is fundamental to identifying the prevailing theories, in helping to design the measurement instruments that are trying to capture the relevant phenomena, and in testing, validating, and modifying theories. At the same time, social science benefits from availability of statistical information which can be used to sort out idle speculation from empirically validated hypotheses.

The implications of such a strategy for statisticians are significant. They must take the lead in convincing decision-makers of the critical importance *for them* of launching such developmental processes. They must also collaborate with policy analysts and social scientists in identifying the data systems that need to be designed. These data systems are typically longitudinal surveys that track outcomes in domains such as health, education, and labour market and income (as the case may be), together with the appropriate

explanatory variables. In addition, the data to be collected through different vehicles should be standardized, to the extent possible, so as to facilitate their joint utilization. And excellent documentation of the concepts and definitions used should be maintained.

4. Implications for Social Statistics – Data Systems

We have tried in Statistics Canada to apply the principles described above in giving a new impetus to the further evolution of our social statistics. Our efforts are intended to serve as illustrations. We summarize these below under four headings: longitudinal surveys; hybrid data systems (exploiting, where possible, the strengths of administrative records); micro data bases; and microsimulation modelling.

4.1. Longitudinal surveys

The longitudinal surveys briefly described below give rather full expression to our notions:

- Each is essentially *sectoral* in orientation.
- We have, to varying degrees, grounded them in *conceptual frameworks* that explicitly recognize a multiplicity of possible *outcome measures*.
- They all seek to illuminate the *connections between outcomes and policies*.
- They are all oriented towards the *microdata analysis of distributional impacts*.
- The frameworks and the corresponding data systems were developed in close *collaboration* with both the relevant policy departments and the academic community.

The first illustration is a longitudinal survey which explores the interactions of the social safety net, the labour market, and family circumstances – the Survey of Labour and Income Dynamics. It is designed to shed light on factors that account for some families successfully escaping poverty while others do not; factors accounting for successful and unsuccessful transitions (from school to work, between job and unemployment, from work to retirement); and the impact of changes in family circumstances (marriage, separation, having a child, death of family member). Not only the variables included, but the basic character of the survey is determined by these conceptual objectives. To illustrate, the period of follow-up (six years) was determined based on the particular dynamic relationships of interest (e.g., labour market and poverty transitions); and the survey unit was chosen to be all individuals in a household (rather than one selected individual) so they all could be followed, however their relationships to the original household might evolve.

The conceptual framework guiding the subject content of the survey is implicit, but it is well-grounded in prevailing theories about the social safety net and the labour market. A high-level (mostly academic) advisory committee is helping both in content development and analysis. Finally, funding for the package of statistical initiatives which included this survey, was obtained with the strong support of the major social policy departments. While Statistics Canada has explicit and ultimate responsibility, both our academic and policy partners are major stakeholders. Finally, the most productive exploitation of the resulting data must take the form of microdata analysis; it is impossible to do justice to the rich character of the data base through the dissemination of partially aggregated tabulations. This poses particular difficulties in terms of confidentiality, but the exploration of that issue goes beyond the scope of the present article.

Our second example is drawn from the field of education and child development. As in the case of the previous survey, we tried to adhere to all of the principles described in this article. In this case we started out with a less complete understanding of even the key factors of interest (i.e., those that lead to the development of productive and fulfilled members of society). Consequently, we did not have an articulated conceptual framework, though we built on earlier work such as the UK (e.g., Power et al. 1991) and more recent Canadian research. Indeed, we hope that the survey will play a key role in developing a framework by exploring the effects over time of the widest possible range of variables.

In compiling a list of possibly relevant variables, our first step was to consult a variety of experts about those factors suspected (or, indeed, known) to have a material impact on the development of children. These potential explanatory variables include mother's health during pregnancy, socio-economic conditions of the family, parenting styles and other attributes of the ''home curriculum,'' early signs of emotional or learning problems, neighbourhood characteristics, teacher's assessment of the child, and principal's assessment of the school. All of these variables (and others) were incorporated into a longitudinal survey of children, initially 0 to 11 years of age, which we are collecting at two-year intervals. The age cohort covered will be extended by two years in each round at least until the leading age cohorts reach early adult life. Indeed, because of clear indications from other longitudinal surveys that causal factors in this domain operate over very long time periods, the objective is to follow the same sample of children for up to 20 years. Of course, the survey is designed also to offer opportunities for important analytical and policy-relevant results on an immediate and continuing basis. In particular, we hope and expect that a considerable refinement and elaboration of the currently weak conceptual framework will emerge as an important by-product.

The third major longitudinal survey is the National Population Health Survey. As in the case of the longitudinal survey of children, its content is informed by broad consultation, but in this case we also invested significant effort into the development of an explicit conceptual framework (Wolfson 1992) which we, of course, expect to refine over time. The data system is conceived as a hybrid, involving both a household survey and administrative records. The survey itself contacts the sampled individuals every two years. Each survey cycle contains a set of core questions on topics including health status, disability, health care utilization, health problems, family situation, and labour market participation or other major activity. Each survey cycle also contains ''focus content,'' a series of questions that delve into a specific topic for that cycle only. For example, the initial cycle focused on stress and mental health. The survey is designed as a hybrid because virtually all respondents consented to have their health care administrative records (e.g., provincial data on hospital stays and physician visits) linked to their survey responses.

4.2. Hybrid data feeder systems

Longitudinal surveys are essential components of the data development strategies needed to support contemporary systems of social statistics. But, of course, they have limitations. They are costly, and their full exploitation may require a long time period. Like most social surveys, they are limited by the weight one can impose on respondents' knowledge

and recall. And, of course, they are subject to sampling error. A variety of supplementary strategies are therefore needed.

First, wherever feasible, consideration should be given to speeding up the accumulation of longitudinal information through the use of recall at the beginning of the survey. Indeed, for major events where good memory is expected, recall surveys can at least in part substitute for longitudinal ones. For example, our family history recall surveys provided a relatively inexpensive initial picture of the covariates of family formation and breakdown, including the strong interactions with educational attainment, fertility history, and labour force participation (Rowe 1989). A module of this sort is now included at the beginning of each six-year cycle of the Survey of Labour and Income Dynamics mentioned above.

Second, we can gain important insights from a given administrative record system if it can be linked longitudinally to cover the needed time span. In Canada we have been exploiting the fact that the social insurance number must be reported by all taxpayers on their income tax forms, enabling the creation of longitudinal tax files. We have also succeeded in linking individual tax forms (for a large sample of individuals) into a good approximation of families. The tax files contain information about individuals' employment and incomes, as well as their age and some indication of their occupation, but of course, few other explanatory variables. While considerably cheaper than longitudinal data, statisticians have much less power to include the explanatory variables that might be relevant to explore causal relationships between *outcomes* and *policy levers*. Also, they may or may not link neatly with our *conceptual frameworks*.

Third, in some circumstances important causal relationships can be studied by creating data sets from the linkage of two files. For example, we have a long history of providing a record linkage service to help study the detailed mortality effects of different occupational health exposures. It is based on our machine readable file on deaths dating from the early 1950s, where a file of persons known to have been exposed some years earlier to some specified health risk is linked to this mortality file. A typical analysis involves the exploration of *causal relationships* (in this case between exposure and health outcomes) by testing for high differential mortality. As another example, and in a prototype mode, very detailed high quality health care data for a sample of one province's residents have been linked to census data. Such linked data are very valuable in exploring the socio-economic correlates of health outcomes and in developing small area health care ''needs'' indicators that can take some account of socio-economic status (Mustard, Derkson, Berthelot, and Wolfson 1997). Such indicators are being used increasingly in some Canadian provinces as the basis for funding new regional health authorities. (Note that this is a linkage between very sensitive data sets. It was therefore preceded by extensive consultations, including the policy ministries concerned, their Ministers, and the Federal Privacy Commissioners. Analysis of the resulting linked data is only conducted on Statistics Canada premises by individuals under the authority of the Statistics Act.)

Sometimes, the best of both worlds can be achieved with hybrid data systems – combinations of administrative and survey data. We have already mentioned the hybrid collection of data in our National Population Health Survey. Another example is provided by our school leaver surveys. Here the sample frame is the administrative data on

enrolled students. The survey contacts (a sample of) the students sometime after leaving school, and collects data on how they are faring. Such hybrid data are potentially more powerful than either source alone because they combine the advantages of each.

It is also worth emphasizing that virtually all of the surveys and innovative data collections just described are the products of strong partnerships both with the relevant government policy departments, and the social science research community.

4.3. Clean micro data bases

The heart of any statistical system is the data on which it is based. As recognized almost 20 years ago (UN 1979), any possible combination of statistics and indicators can be derived from clean micro data bases. For example, enrollments by province, by type of post-secondary institution, and by course of study represent a sequence of tables, all of which could be generated by a single micro data file of students containing information on province, type of institution, course of study, etc. This micro data file is ultimately far more compact and flexible than even a large combination of pre-specified tables. The derived statistics and indicators are then just that, derived. We can add to this traditional and well-known argument for maintaining clean micro data bases our additional consideration that micro data provide unique support for the exploration of *distributional issues*.

The availability of powerful computing and richly detailed microdata sets also changes the nature of the discussion of social indicators. In the 1970s, at the height of the (so-called) social indicators movement, a major focus of discussion was defining and seeking consensus on a specific set of indicators. However, such a pre-occupation would be misplaced to-day. Rather, the underlying question should be the desired *range* of statistical indicators. The focus should be on the content and coverage of the data feeder systems, and the format of the resulting clean micro data bases. Then, software can be readily developed to produce the desired range of indicators. As a consequence, the ''stakes'' involved in agreeing in advance on specific indicators are much lower. If for some reason a change is required, the underlying data can be inexpensively reprocessed (e.g., tabulated) to produce the revised indicator.

It is undoubtedly true that, from the perspective of the statistical office, the most effective form of storing its information holdings is through micro data. This also represents the most convenient access to data, at least for serious researchers. This does not mean, of course, that we should abandon the production of summary information and simply make available clean micro data to our users. We continue to need to produce summary tables and analyses in part to meet the needs of a wide range of clients, and because of the constraints of confidentiality. But we do advocate a shift of perspective where the production, documentation, and preservation of micro data is seen as a central function.

In practice, confidentiality constraints are very real, and are significantly more serious with longitudinal data, where the amount of information accumulated over time can almost certainly uniquely identify almost any individual in the sample. A full discussion of this issue would clearly take us too far afield. We would like to note, however, that Statistics Canada is proceeding on a triple track. First, we *are* disseminating micro data from all of our social surveys, subject to suppressing geographic detail, collapsing codes, etc. Second, we are also experimenting with providing ''softly coupled'' access to

researchers to the full micro data. This means the provision to researchers of a great deal of detail about both the logical and physical structure of the full internal clean microdata file, including a file with exactly the same format and structure as our internal file but containing dummy data. Researchers can use the publicly available but partially suppressed micro data file for exploratory research, formulate (in some high level language) more detailed retrieval requests, test these on the dummy file, and submit them to Statistics Canada electronically. We run the fully specified data analysis requests against the full internal version of the file, check the results for confidentiality disclosure, and send the retrieved information back to them also electronically. Thirdly, we are exploring with the social science community an even more direct form of access, via designated regional centers of Statistics Canada, under specified legal and scientific conditions which are not as yet fully defined. We will explore the experience of the United States Bureau of the Census with respect to this third track.

A specific example of exploiting the power of micro data bases is the development in Statistics Canada of the so-called Total Work Accounts System (Stone and Chicha 1996), a network comprised of a micro data file (a detailed time use survey in the present instant), a conceptual framework encompassing both paid and unpaid work, statistical tables linked to the micro data files, and some pre-analysed statistical indicators. The TWAS is designed to explore such questions as: the prevalence of child care activities across key sub-groups of men and women; volume and value of informal health care activities; gender differences in the pattern of balancing work and family obligations; inter-group differences in the volume of informal supports provided and received; etc.

4.4. Microsimulation modeling

Finally, it is important to highlight a technique both for forming statistical systems, and for making them more useful for public policy, namely microsimulation models. Such models in the social sciences have their longest history in Ministries of Finance, where they are used to analyse prospective income tax changes. As such, they are clearly driven by high priority policy needs. At the same time, their construction often involves innovative and very intensive use of micro data on income distributions (Citro and Hanushek 1991). Indeed, some of the earliest examples of hybrid data involved the linkage of income distribution surveys and income tax records. The *de facto* result is an estimate of the household distribution of income that is both fully based on explicit micro data foundations, and an amalgam of the best possible data from a range of sources (Adler and Wolfson 1989; Bordt et al. 1990).

This latter aspect, drawing on a range of data sources, is a hallmark of the success of the SNA, but one that has generally not been followed in social statistics. But policy needs, even at a sectoral level, usually involve a range of factors that are not, and typically could not, all be included in a single omnibus survey. Thus, some sort of synthetic apparatus or methodology is required to knit together data from several sources. Moreover, to maintain the desired microanalytic character, new methods are required, compared to the semi-aggregate ''balancing'' and ''macro-editing'' methods used in constructing the SNA.

Microsimulation models and related techniques provide one means to achieve these statistical objectives. For example, dynamic microsimulation models can be used to study

the long-term evolution of a population cohort under different policy scenarios. The cohort may consist of hypothetical persons, or a sample of persons initially drawn from a cross-sectional survey. The longitudinal dynamics are injected by simulating the passage of time through the use of transition probabilities. For example, all persons are exposed to the probability that they marry during the (modelled) year, have a child, divorce, get a job or lose one, earn a certain income, start smoking, have a heart attack, become disabled, etc.

Such models are good vehicles for the integration of our existing knowledge about the phenomena under study, including information on dynamics and causal factors represented in the form of transition probabilities (or better yet, transition probability densities, expressed as functions of relevant covariates). Subject to our level of knowledge, these kinds of models provide a foundation for coherent *families* of summary outcome indicators.

One example is a family of variants of life expectancy that generalize the notion of disability-free life expectancy (DFLE) discussed earlier by allowing a richer description of health states. The resulting summary measure can then be disaggregated in various standard ways (e.g., by age group or province or educational attainment). Moreover, it can be extended by means of ''what if'' questions, such as ''how much would health-adjusted life expectancy change if health problem X did not exist?'' (Roberge et al. 1997; Wolfson 1996a). The answer to this kind of question is analogous to, and a generalization of, cause-deleted life expectancy which, in turn, is often used in policy discussions to assess the relative importance of different groups of health problems (e.g., heart disease versus cancer).

Analogous social statistical indicators, also generalizations of life expectancy, are feasible for socio-economic factors, provided individuals are the focus, and microsimulation is used to weave together both cross-sectional and longitudinal data from a variety of sources. For example, estimates can be produced of the average number of jobs an individual can expect over his or her lifetime, assuming current job turnover rates; or the number of years a typical child can expect to spend growing up with a lone parent or without siblings, assuming current divorce and fertility rates (Wolfson 1989; Wolfson 1996b).

These same kinds of microsimulation models also support the exploration of alternative policy scenarios. The alternatives may be represented directly as changes in program parameters (e.g., public pension entitlements, health care treatment protocols), or implicitly in the form of modified transition probabilities (e.g., the effects of alternative student loan programs on participation in post-secondary education). Such policy changes can then be connected to policy-relevant outcomes such as proportions of the future elderly likely to be counted as having low income (Murphy and Wolfson 1991; Wolfson and Murphy 1997), the prospective government cash flow and student indebtedness effects of changes to student loans (Statistics Canada 1997), or the effects of alternative chemotherapies on lung cancer progression (Evans et al. 1997).

Finally, the process of building and using such models highlights existing gaps in our knowledge, and in data availability. Moreover, via appropriate ''guesstimates'' combined with sensitivity analysis, they can help in assessing their relative importance, thereby providing strategic guidance for statistical priorities (e.g., Wolfson and Rowe 1996).

No single survey, longitudinal or otherwise, can possibly encompass the full range of potentially relevant information needed to study the kinds of issues just mentioned, and those typical of the social policy domain. The necessary information will sometimes only be possible if assembled from a number of sources in the form of models that seek to capture the likely causal pathways, and provide a capacity to simulate the impacts of alternative policies.

Such complex models will, however, never be complete. As a result, they should only be contemplated in an evolutionary mode, based on effective collaboration among decision makers, social scientists, and statisticians. These individuals will have to work together on a repeating cycle of conceptual articulation of models, collection of relevant data, causal inferences from these data, construction of empirical models that embody these inferences, derivation of model results, assessment and testing against other statistical data collection and analysis, and then back again through an iterative loop to refine the data and the model.

5. Concluding Comment

We close by asking the reader to recall the conclusions summarized at the outset. The confluence of renewed social policy concerns, dramatic enabling improvements in computing and hence statistical methods, and the possibilities arising from more explicit tri-partite relationships among social policy departments, the academic community, and national statistical offices signal exciting prospects for the development of more coherent and forward-looking systems of social statistics.

6. References

Adler, H. and Wolfson, M. (1988). A Prototype Micro-Macro Link for the Canadian Household Sector. Review of Income and Wealth, 34, 4.

Bordt, M., Cameron, G., Gribble, S., Murphy, B., Rowe, G., and Wolfson, M. (1990). The Social Policy Simulation Database and Model: An Integrated Tool for Tax/Transfer Policy Analysis. Canadian Tax Journal, 38, 48–65.

Bottani, N. and Tuijnman, A. (1992). International Education Indicators: Framework, Development and Interpretation. Making Education Count, Developing and Using International Indicators. OECD, Centre for Educational Research and Innovation, Paris.

Citro, C.F. and Hanushek, E.A. (1991). Improving Information for Social Policy Decisions: The Uses of Microsimulation Modeling. Vol 1, Review and Recommendations. Panel to Evaluate Microsimulation Models for Social Welfare Programs, National Research Council, National Academy Press, Washington DC.

Evans, R.G., Barer, M.L. and Marmor, T.R. (1994). Why Are Some People Healthy and Others Not? The Determinants of the Health of Populations. Aldine de Gruyter, New York.

Evans, W.K., Will, B.P., Berthelot, J-M., and Earle, C. (1997). The Cost of Combined Modality Interventions for Stage III Non-Small Cell Lung Cancer. Journal of Clinical Oncology, 15, 3038–3048.

Fellegi (1995). Social Science and Social Policy. Commencement Speech at Simon Fraser University, October 1995.

Hanushek, E.A. (1981). Throwing Money at Schools. Journal of Policy Analysis and Management, Vol. 1, No. 1, Fall.

Krieger, N. (1994). Epidemiology and the Web of Causation: Has Anyone Seen the Spider?. Social Science and Medicine, Vol. 39, No. 7, 887–903.

Murphy, B.B. and Wolfson, M. (1991). When the Baby Boom Grows Old. Statistical Journal of the United Nations Economic Commission for Europe, Vol. 8, No. 1.

Mustard, C.A., Derkson, S., Berthelot, J-M., and Wolfson, M.C. (1997). A Comparison of Household and Neighbourhood Level Measures of Socio-Economic Status in the Study of Population Health Status.

Nord, E., Richardson, J., and Macarounas-Kirchmann, K. (1993). Social Evaluation of Health Care versus Personal Evaluation of Health States. International Journal of Technology Assessment in Health Care, 9:4, 463–478.

OECD (1976). Measuring Social Well-Being, Paris.

OECD (1998). Social Indicators: An Illustrative Compilation. Paper prepared for the June, 1998 Meeting of the Employment, Labour and Social Affairs Committee at Ministerial Level, Directorate for Education, Employment, Labour and Social Affairs.

Power, C., Manor, O., and Fox, J. (1991). Health and Class: The Early Years, Chapman and Hall.

Roberge, R., Berthelot, J-M., and Wolfson, M.C. (1997). Adjusting Life Expectancy to Account for Morbidity in a National Population. Quality of Life Newsletter, MAPI Research Institute (instdoc@mapi.fr), Lyon, France, No. 17.

Rowe, G. (1989). Union Dissolution in a Changing Social Context. J. Legare, T.R. Balakrishnan, and R.P. Beaujot (eds.), The Family in Crisis: A Population Crisis?, Royal Society of Canada, Ottawa.

Statistics Canada (1997). LifePaths: Student Loans, Software Users' Guide. Social and Economic Studies Division, Ottawa.

Stone, L.O. and Chicha, M.-T. (1996). The Statistics Canada Total Work Accounts System. Statistics Canada, Catalogue No. 89-540-XPE.

Stone, R. (1973). A System of Social Matrices. Review of Income and Wealth, Series 19, No. 2, 143–166.

Tomiak, M., Berthelot, J-M., and Mustard, C.A. (1997). Factors Associated with Nursing Home Entry for Elders in Manitoba, Canada. Analytical Studies Research Papers, Statistics Canada, forthcoming.

United Nations (1975). Towards a System of Social and Demographic Statistics (SSDS). Studies in Methods, Series F, No. 18, ST/ESA/STAT/SER F/18, New York.

United Nations (1979). The Development of Integrated Data Bases for Social, Economic, and Demographic Statistics. Studies in Methods, Series F, No. 27, ST/ESA/STAT/SER F/27, New York.

United Nations Development Programme (1997). Human Development Report, 1997, Oxford University Press.

Wilk, M.B. (1991). Health Information for Canada. Report of the National Task Force on Health Information. Ottawa.

Willms (1992). Monitoring School Performance, A Guide for Educators, The Falmer Press.

Wolfson, M.C. (1989). Divorce, Homemaker Pensions, and Lifecycle Analysis. Population Research and Policy Review, 8, 25–54.

Wolfson, M. (1992). A Template for Health Information. World Health Statistics Quarterly, Vol. 45, No. 1, 109–113 plus diskette.

Wolfson, M. (1996a). Health Adjusted Life Expectancy. Health Reports, Vol. 8, No. 1, Summer, 1996.

Wolfson, M.C. (1996b). LifePaths: A New Framework for Socio-Economic Accounts. 24th General Conference of the International Association for Income and Wealth, Lillehammer, August, 1996.

Wolfson, M.C. and Rowe, G. (1996). Perspectives on Working Time Over the Life Cycle. Canadian Economic Research Forum conference on Changes in Working Time in Canada and the United States, June, Ottawa.

Wolfson, M.C. and Murphy, B.B. (1997). Aging and Canada's Public Sector: Retrospect and Prospect. In K.G. Banting and R. Boadway (eds.). Reform of Retirement Income Policy, International and Canadian Perspectives, School of Policy Studies, Queens University Press, Kingston, Ontario.

WHO (1980). International Classification of Impairments, Disabilities, and Handicaps: A Manual of Classification Relating to the Consequences of Disease (ICIDH). Geneva.

World Bank (1993). World Development Report 1993, Investing in Health. Oxford University Press.

Received October 1997
Revised December 1998