INFORMATION PROCESSING, CONTEXT and PRIVACY

Alan Dix[†]

Human-Computer Interaction Group and Department of Computer Science,
University of York, YORK, YO1 5DD
0904 432778, alan@uk.ac.york.minster

This paper is about an old concept, data processing, but one that has taken on new meaning with the increasing complexity and interconnection of systems and the burgeoning of expert systems and connectionism. Classical information theory has been found to be inadequate even in the relatively formal context of security, but this inadequacy is intensified when we consider more human issues like privacy. Further, writers like Suchman and Winograd & Flores emphasise context in understanding communication and information. Relating these issues to a simple information life-cycle, this paper questions how we can retain an understanding of human issues when interacting with such complex systems.

Keywords: Information theory, information processing, privacy, formal analysis

1. Interacting with computers

As I write this I am sitting at my desk at home using a personal computer. Later, I shall take the disk into the University and read it into a compatible computer. On my desk there I have a workstation which I will use to format and print this paper. I am also networked to various other computers in the department, which I occasionally use more or less transparently. These are simple face-to-screen interactions, and are obvious examples of personal interaction with computers, but the majority of our interactions with computers are less direct.

There are, of course, the large range of embedded microprocessors in day to day equipment such as washing-machines, electric drills, hi-fi, and cars, but as these are not information systems, they are not relevant to this paper. Moving outside our homes, we might use automatic bank teller machines, pay for goods using electronic funds transfer at a supermarket, or stand behind counters in various institutions while those we talk to enter information about us and consult terminals to national networks. However, there are many more systems that we never even see at all, with which our interactions are totally indirect. Electricity, gas and telephone companies all hold records and have largely automatic billing. Both the university which employs me, and SERC which pays my grant, use a mixture of manual and automatic mechanisms, as do insurers, credit card companies, and mail order firms. Some systems are so diffuse that it is hard even to be aware of their importance: corporate decision support systems, simulations for local government road planning, party political and media planning. All of these are supported to some extent by automated information systems, all clearly affect us by their decisions but, less clearly, all take input from us: market research and sales figures, monitoring of traffic flows, opinion polls and audience viewing figures. Finally, there are those systems which are most discussed and of which (perhaps) least is known, that is, governmental, police and security service files, including social security, tax and census information.

We are also, of course, the recipients and users of many information systems. In addition to direct services such as bank teller machines, we have television, radio, newspapers, tele-text services and direct mailings. Further, we are affected by the diffuse decision-making processes mentioned above.

Privacy and bugs

When considering large corporate computer systems, both governmental and private, the emphasis is frequently on the deliberate misuse of personal data. It is to counter such misuse that privacy legislation such as the Data Protection Act in the UK is aimed. To make an assessment of the acceptability of various uses of personal data and of the adequacy of such legislation, we need to understand the nature of information and the ways it is transformed and used.

A further problem is bugs. In private or embedded computer systems these are relatively well understood, and we can notice and deal with them at a local level. Large numbers of breakdowns can of course have very damaging results, just as large numbers of engine breakdowns would clog up the rail system, or misbehaving telephones damage communications. However, from just such analogies we have some idea of how to deal with local breakdowns. Even though it can be very distressing to lose a file or deal with a bad interface, it is comparable to a typewriter with sticking keys and probably less damaging than a leaking gas appliance. Thimbleby has noted that the attitude to bugs in the software industry is far from professional, and this bodes ill for the correspondingly more obscure errors of design and interpretation possible in more complex systems.

When we look at large scale information systems, in addition to the familiar bugs in the access mechanisms and their interfaces, there are a whole new set of problems connected with the access to and interpretation of the large information bases involved. In the early days of data processing such bugs usually involved multi-million pound electricity bills or alternatively sudden massive bank balances. The current equivalent would be credit

† The author is funded by a SERC Advanced Fellowship in Information Technology

blacklists, where people find themselves mistakenly refused credit because someone with a similar name or a previous occupier of the same address once defaulted. Less clearly mistaken is the use of neighbourhood as a basis for assessing creditworthiness, or to demand additional deposits. For example, I was recently told by one of my banks that they could not post cheque books to my home address as I had a "high risk postal code". One wonders seriously about the quality of information used for such decisions, especially as the frequent response to queries is that the "system says so". Bugs in information systems may also be just as life threatening as bugs in aircraft control systems; Vallee [1] quotes an example where French police shot a motorist due to mistaken information in the computerised record system.

From a management perspective, poor, mistaken or misunderstood information could easily lead to corporate ruin. For instance, credit companies do not really want to refuse credit to a good potential customer (or allow it to a bad one) any more than the electricity companies wanted to send out silly bills. In the same way as we had to understand the nature of information to assess the deliberate misuse of information, we need a similar understanding to detect and prevent its accidental misuse.

In some ways, it makes little difference to the above examples whether the relevant systems are computerised or not. Large corporate information systems are by their nature complex and can easily use, or misuse, poor quality information. Similarly, they are often hidebound by arbitrary rules which differ little from those executed in a computer system. In fact, the rules may be inferior if they deliberately ignore information to make the human processing task simpler. It is the massive quantitative increase in the complexity of information *processing* that gives a different qualitative nature to computer centred information systems.

2. Formal models of information

As we have noted, in order to assess issues of privacy and misuse of information we need some understanding of the nature of information. One simple requirement is to measure how much information we have, and where the information is going. The simplest measure of information is length. For instance, this paper is aiming to be about six pages long, and at various stages I will do a word count to see how I am getting along. However, when thinking of formal treatment and measures of information, one turns to classical Shannon and Weaver information theory [2]. Its concepts colour the understanding of information even among those who have little grasp of its mathematics and range of application. Indeed, one would surmise that the popular idea of information lies somewhere between classical information theory and word counts! In fact, such a model turns out to be inappropriate when considering human issues like privacy.

Classical information theory

The origins of classical information theory lie in the analysis and design of data channels and optimal data transmission strategies. The information content of a message is given a precise value in the context of the ensemble of possible messages and their likelihoods. Information is measured using entropy:

$$I(\textit{message channel}) = -\sum p_i \log_2 p_i$$
$$I(\textit{particular message}) = -\log_2 p_i$$

where the p_i are the probabilities of the various alternative messages. The information conveyed by a message of fixed length is maximised when all alternatives are equally possible.

Of course, this is a more subtle measure of information content than the simple expedient of measuring the length of a message. However it has no idea of the content of the message except in so far as this is captured by relative probabilities. This means, for instance, that in a nuclear reactor control room a light indicating that the kettle has boiled for tea carries more information content than an indicator of potential meltdown! It also takes no account of the context of interpretation so that translating a message into Swahili would leave its information content unchanged (reasonably) but of course would not be equally acceptable in a British power station. (In fact the opposite is far more likely to occur in practice.)

The vital importance of context in understanding language, and, by extension, computer data, has been emphasised by authors such as Winograd and Flores [3]. To be fair, however, the Shannon and Weaver definition was not intended for such purposes. The problem is that its conceptual impact is pervasive and is likely to be translated to inappropriate areas.

One of the nice things about an entropy-based definition is its pleasant mathematical properties; for instance if A and B are two independent messages then

$$I(A+B) = I(A) + I(B)$$

In general there is correlation between messages (for instance after reading the letter "q" the successive letter "u" has little additional information). In this case we have the inequality:

$$max(I(A), I(B)) \leq I(A+B) \leq I(A) + I(B)$$

the left hand side equality being obtained when the correlation is complete, and the right when the messages are independent as above. We can summarise these properties as "linear" behaviour.

Security

Interest in security comes from two main groups, commercial and military. The desire here is to minimise the "information" obtained by an adversary, either the enemy or the competitor. Some work in these fields has concentrated on classical concepts of information, reducing the bandwidth of information loss, ideally to zero, but more practically merely to insignificant levels (eg. a bit a day) [4].

Even in this relatively formal field, however, the classical approach has been found inappropriate [5]. One reason is the relative importance of different messages, alluded to above. For instance, if an enemy spy opening and shutting a blind means "D-day tomorrow", then it is clearly a more important "bit" of information than "fish and chips in the mess again".

In addition, measures of security which include importance (related to military or commercial damage) do not follow the simple linear patterns of composition. For instance, enemy knowledge of *either* a ship's latitude *or* longitude may not be important, whereas the combined information is highly secret. That is, security measures exhibit "super-linear" behaviour:

$$\exists\ X, Y \quad s.t. \quad S(X+Y) > S(X) + S(Y)$$

It is crucial to be aware of this behaviour as it denies the intuition obtained from classical information theory. The fallacy of linearity (albeit essential when appropriate) is as dangerous when applied to information as it is to environmental damage or straws on camels backs.

Privacy

Privacy, whether personal or corporate, is subtly different from secrecy. Things are private because the very fact of another person *knowing* them is as important as what they might do with that knowledge. Private things may be secret (because we do not want others to know) but not necessarily vice versa. For instance, the formula for a paint would be secret, because one would not want other manufacturers to copy it. The dumping of residues from its manufacture in the North Sea would be private as it would damage the corporate image. Similarly troop strengths would be secret from an enemy, to prevent them using the information in planning their attacks, but private from ones own population to guard morale.

The properties of privacy are more complex again than those of secrecy. It is usually the case that the more pieces of secret information someone has the less happy we are. This is not true of privacy. It is frequently the case that you do not want someone to know something, but, if they do find it out, you want them to know something else as well. For instance, a business man may not want his wife to know that when he said he was working late he was actually with another woman. However, if this fact comes to light he may reveal to his wife that the other woman was a silver-smith with whom he was agreeing the details of a bracelet for their wedding anniversary.

We could call this behaviour "sub-linear", that is, given a measure of privacy P:

$$\exists\ X, Y \quad s.t. \quad P(X+Y) < P(Y) < P(X)$$

There are parallels with this behaviour in the security world, like flooding an adversary with useless information to hide the important bits. However, these trade on the inability of the adversary to access the relevant information; in the case of privacy we may not mind other parties *knowing* something so long as they know the context as well.

This sub-linear property of privacy is related to the need for context that Winograd and Flores [3] demonstrate in the understanding of language. One could characterise problems of privacy like the example above in terms of context/understanding. Essentially, there is some information that has some meaning in its true context with which we are happy. However, if someone attempts to understand it without that contextual background their understanding will inevitably be based on some assumed context yielding a damaging interpretation.

Sub-linear behaviour reminds us that we must be as worried about too little information being stored as too much. Automated systems are particularly worrying as, if there is one thing that computers do better than remembering, it is irrevocably forgetting. This concern about loss of context is important both for individual and official information. Most government and corporate misinformation is true.

Resumé - formal properties of information and privacy

It is obvious that, on the one hand, bringing together seemingly innocuous pieces of data can produce something valuable, and, on the other hand, ignorance of relevant data or context can significantly alter the interpretation of other data. However, these obvious facts violate the two inequalities in the linearity equation implied by classical information theory, the very measures that colour our perception of modern information systems.

3. Information life cycle

The foregoing discussion concentrates on "information" as a passive item. However, any understanding must look at the *process* by which data is gleaned, analysed and used rather than just the stored entities. A simple model would be to divide the information life cycle into three parts:

- Collection
- Processing
- Use

This is clearly an over-simplification. For instance, there will be various phases of analysis before data is collected, the use of one piece of data will affect subsequent data collection, and there are likely to be feedback loops. However, since the purpose is to show the complexity of the process, considering such a simplification will be sufficient.

To each phase in this life cycle we can add various attributes:

Collection

Who (or what) is the subject of the information and in which of its attributes are we interested? Who is collecting the information, for what purpose and in what context? When the information is gathered in, who owns it? This last question probably has different answers depending on whether one looks at it from an ethical, legal or de facto point of view.

Processing

Once data has been codified it will be processed in some way. It is important to know the intentions and purpose behind the algorithms used for processing the data. There is no reason to assume that this purpose is related to the purpose for which the data is collected or for which it is eventually used. By whom is this processing done? There is a wealth of difference between an application for postponement of tax payments being considered by the district inspector of taxes and an application for housing benefit being considered by a programmer in Bletchley.

Use

Finally, how is the processed information to be used? Someone (or something) will use it for some purpose (not necessarily related to the reason for collection). The actions that arise are likely to affect the original subject, amongst others. Of course, even if these actions are acceptable and "correct", there are many issues to consider in the way this process interfaces to the user, and whether the manner of collection and use are consistent with the dignity and privacy of the subject. Automation may help in this respect: it is far less humiliating to be refused money by a bank telling machine than to be refused cash over the counter.

One could expand upon each of these phases but it is the processing which I consider central. It is interesting that over the years the phrase "data processing" has become rather out-dated,

even "information processing" sounds a little too technical and we are all now engaged in "information technology" (note my funding!). The processing aspect is played down, and yet it is precisely this processing that gives computer based information systems their power and danger.

Examples of the information life cycle

In the simplest human information systems, the collection, processing, and use would all be carried out by the same individual, perhaps at the same time. The context is taken into account when analysing and acting on the data and the purpose for which the data is to be used is apparent when it is collected. There is, of course, plenty of room for misinterpretation and misunderstanding but it is *relatively* easy to focus on these and correct them. An example of this would be if the businessman in §" 2 met his wife while with the silver-smith.

An example of a more complicated human information system would be a small building society. A borrower is behind with her payments partly because of a delay in receiving her salary after a job change and partly because of delays due to a postal strike. She phones up the building society to apologise and explain and a note is made in the margin of the *paper* records. Later, she goes to request an extension of her mortgage to cover some new central heating. The manager looks up her payment record and notices the period of arrears. At first he is reticent but then, perhaps after discussion with the borrower, he sees the marginal notes and agrees the loan.

Imagine now that the building society computerises its records. There is no room for the marginal note and the information is lost. This time, if the borrower submits her (streamlined and efficient) postal application she is refused. Perhaps she follows it up and the matter is cleared up, perhaps she gets finance elsewhere, or perhaps she sticks with her coal fire.

Finally, let us consider a totally automated system. The effects of the need for automatic processing are felt even at the collection stage. Usually only predefined information can be entered, so that additional marginal notes would never have existed in the first place. Also the information collected will have to fit into categories that may not be appropriate for the actual data subject. For instance, I was once computerising the personnel records system for an education authority. The teachers' titles were given code digits, and I was specifically requested *only* to allocate codes for Mr, Mrs and Miss. Clearly the teaching profession contains no Doctors, no Nobility and no Feminists!

In this last example the failure of the information system is apparent even before processing begins. What hope is there that sensible decisions can be made by the eventual user of the information after it has been changed in form by the processing stage?

4. The nature of processing - alienation

From the increasingly automated examples above, we see that the role of automatic processing in the middle of the information life cycle separates the user of the information from its source. This means that valuable contextual and hard to codify information is not available. Furthermore, the processing itself may make it hard to trace back to the original data, such as it is. We shall now look at the way that different types of processing relate to the properties of information discussed in section 2.

Again, we will grossly simplify and consider processing under three headings:

Selection

This is direct access to data stored in the system, by some sort of query. Information is of itself useless without an effective access mechanism and it is the ability of modern data base systems to access data quickly and by various mechanisms which has revolutionised this aspect of information systems. This is probably the type of processing envisaged when people consider privacy issues: can someone get at information about me? The Data Protection Act recognises this and makes allowance for personal information that can be accessed by keys not actually held upon computer. A better use of computer facilities is perhaps the indexing of data stored manually, the prime example being library indices. As far as I am aware the UK legislation makes no special provision for computer indices to manual data. However, in terms of accidental misuse this form of processing is probably the most benign, because it at least presents the user with the data as originally collected. (One could even make provision for free format fields on all records and thus approximate the paper form.)

Collation

This is the gathering together of related information and is a crucial requirement of any data base query facility. This bringing together of disparate data brings into play the super-linear behaviour discussed under security. From the point of view of privacy, items of personal information gathered separately which individually I am willing to divulge may be brought together, and thus compromise my privacy. On the other hand it is exactly this super-linear behaviour that gives collation its power. The query yields information of greater quality than the sum of the data brought together. On the whole the major problem here is from malicious misuse of information. However, the fallacy of linearly additive information might lead someone to use collated information in a way which infringes one's privacy without considering that there are any human or ethical questions involved.

Filtering

This is where we deliberately throw away unwanted information. It is similar to selection, but there we were considering selecting entire data records, whereas here we consider the extraction of parts of a single record for some purpose. Formally, the two are pretty much equivalent, but if we consider a record to be the total of the information collected about a single subject we see that filtering has a quite different effect. If we recall the sub-linear behaviour of information with respect to privacy, we see that filtering of data can invade privacy by losing relevant information. In fact, we have already seen that all data collection involves some loss of context and this is especially so when the data is intended to be entered into a computer system. Thus filtering really begins at the collection stage. The collector may well be aware of the paucity of the information gathered, whereas the eventual user may be blissfully ignorant. It is well known that the choice of indicators to be taken into account when describing people can be used as a subtle means of discrimination against specific groups. Thus deliberate misuse under this category is possible. However, it is probably one of the most frequent ways in which information is unintentionally misused and it is important principally for this reason. In fact, most of the modern "bugs" in information systems seem

to stem from the ignoring of pertinent information.

Statistical analysis

Statistical measures of large data bases may be seen as protectors of privacy, by hiding the particular in the whole. For instance, UK census data is deliberately only supplied in aggregate form. Clever use of multiple statistical queries may in fact compromise individual information [6], but this problem is well understood and techniques are available to protect against it. It is unlikely that such a circumstance would arise unintentionally. There is however again the problem of loss of context. A parent may drive his child 100 metres down the road to school because the road is unsafe to cross. On the way, he passes an observer measuring road usage. Because the road is used such a lot it is widened, attracting more traffic and thus making it more dangerous.

Mathematical and symbolic analysis

We are thinking here about applying complex algorithms to derive information about a specific subject. This relates to statistical analysis in the same way that selection relates to filtering. It comprises aspects of both collation and filtering, for instance, $a+b$ brings together two pieces of information while at the same time losing their individuality. Further, the very complexity of the operations may make it almost impossible to relate the derived information to the originally collected data.

We notice that once we take into account the super-linear and sub-linear nature of human information, each of the modes of processing has the possibility of violating privacy or misusing information due to loss of context. If the information system has been designed as a whole, these effects may be taken into account, but it is more likely that users of information will come to an existing information system. The data with which they work may have been collected for a totally different purpose. Important facts may be missing, and even where data is complete the way in which the real world has been disambiguated in order to codify it may not be pertinent for the new purpose. Further, the user may not have the raw data available, but will instead access it in an already processed form. Again, the purposes of this processing will affect the way information is collected together or ignored. With the best of intentions the user is forced to take action on the basis of this incomplete and possibly irrelevant data. Any statistician knows the problem of being asked to produce some sort of analysis from derived statistics, where perhaps the wrong type of data was gathered and the original data lost. One could compare this situation with the industrial revolution. At that time, workers became alienated from the source and purpose of the material artifacts they produced. Today, information workers are similarly alienated from the subject and context of the data which they process and use.

Connectionism and expert systems

Until quite recently, the majority of processing has been of a fairly simple algorithmic nature. In *principle*, one could analyse the method of collection and processing to decide whether the derived information was suitable for the purposes required. For some time now, expert systems have been used as part of the processing element in information systems. The complex nature of their decision-making can increase the distance between the user and the original data. To some extent, the problems of alienation have been recognised; users are unwilling to trust the conclusions of such systems without access to the decision-making process, and explanation facilities are often included.

The advent of connectionism into commercial systems poses more severe problems again. Amongst others, credit companies are considering the use of neural networks in assessing creditworthiness. The problem with such systems is that the processing they perform is so diffuse and unstructured (and deliberately so) that it may be theoretically impossible to obtain a similar explanation. That is, the processing may contain elements of collation and filtering, but it will be impossible to know what information is being ignored and what brought together. Information users may violate the personal rights of a subject without being aware of it, and without the ability to detect that they are so misusing the data. Although research is being done to address the "explanation" of neural net decisions it is likely that the use of such systems will fast outstrip the ability to understand them.

5. Addressing the problem

The picture painted is a trifle gloomy. One reaction might be to eschew all complex processing because of the inherent dangers. For instance, I used to have a book dating from the late 1960s [7] that suggested, amongst other anti-computer measures, magnetically "wiping" your cheques to force them to be processed by hand. Luddism has a bad name, but faced with technologies that benefit the powerful at the expense of the ordinary man and woman such a response may well be defensible [8]. A second reaction is exemplified by the Data Protection Act in the UK, which, continuing the analogy to the industrial revolution, we may see as following in the tradition of reformers by statute such as Wilberforce and Shaftesbury. Finally, and most relevant to the reader, there are the awareness and professional standards of information technologists. But what measures, if any, can we take to make effective use of information systems and to protect the integrity of their subjects?

The value of processing

No one would bother with the costly business of processing information unless it yielded some value. In this the super- and sub-linear behaviour of information is specifically used. We collate disparate data or discard unnecessary data because that enhances the value of the derived data. If the derived data is of more value to the data user it is not surprising that it may also be of more value to the subject.

This leaves us in somewhat of a dilemma. It appears that it is fundamentally impossible to process data effectively without necessarily also causing problems of ownership and privacy. Of course, some types of processing may maximise the commercial value of derived information whilst causing few personal problems, and vice versa. However, the connection between power and privacy does imply that we cannot simply bar classes of processing as unacceptably violating civil liberties and allow others wholesale; any type of processing, if it is useful, is a potential privacy problem.

On the one hand, the value of correctly processed information can be a strong argument for avoiding misuse, as this misuse is just as harmful to the data user as to the data subject. On the other hand, data collection, storage and processing is costly in itself. One reason why bugs in information systems (such as the problems of credit blacklisting) persist is that the costs of

proper processing and use outweigh the losses due to misuse. Unfortunately, as in many spheres, the costs to the individual are not included in this formula.

People-friendly processing

Is there such a thing? As with all areas the solutions to the problems cited will be as varied as the systems designed. One can however give a few suggestions.

To begin with, knowing that the problem exists is the key to solving it. In particular, the most potent and most easily overlooked cause of information misuse is due to the sub-linear nature of information, ignoring pertinent knowledge which could lead to better systems both for users and subjects. A "people-friendly" company would include these factors in their assessment of appropriate information processing strategies.

Processing as the source of alienation is the key problem. We can therefore design systems that attempt to bridge the gap between collection (and subject, context etc.) and use. Traditional data files mirrored the constrained fields of punch cards and the fixed field approach has been uncritically extended into more advanced data base architectures. One could imagine data base designs that include space for free format comments or perhaps have several different answers to the same question related to some (coded?) indication of the context of use. Even more unconventionally, we might imagine doubtful or special items of information being active so that if they are used in any calculation they could signal to the user their specialness. On a more traditional level, we could just choose that wherever we can, we present as much of the original data as possible even when this is done in conjunction with derived data.

Auditing may well be at the heart of a people-friendly company's information policies. At a per transaction level this would imply that as far as possible the user was made aware of the source of information used, perhaps in the ways described above. More importantly, this should be a professional activity in its own right. In the same way as we might assess the energy efficiency of a factory or check for financial irregularities in a company's accounts, the professional information auditor would examine the information systems for the way they use or misuse personal information. They might well be able to tell the company how better to use their information as well as giving it a "people clean" bill of health.

Even more radically, one could attempt where possible to replace 'information up, decisions down' processing with one where strategic information was passed down for local, contextual, decision making. This has its technical problems, as the local structure of use may differ from the structure of the data it depends on. However, the main problem with such an information structure is not technical, but that it would conflict with the corporate power structures.

6. Conclusions

We have seen how classical information theory, the basis of much of our intuition about information flows, fails to provide appropriate measures of privacy and personal importance. Processing, at the heart of the information life cycle, may thus violate personal rights subtly and perhaps unintentionally. Moreover it causes a separation between the user of information and the subject of that information which makes effective and proper use impossible. Some strategies have been suggested for tackling this alienation, but whether complex systems can be tamed in this way is debatable. Of particular concern is the prospect of connectionist approaches to the processing of personal information, which make it impossible to tell whether or not privacy is threatened.

References

1. J. Vallee, *The Network Revolution*, Penguin Books (1982).
2. C.E. Shannon and W. Weaver, *A Mathematical Theory of Communication*, Univ. of Illinois Press (1949).
3. T. Winograd and F. Flores, *Understanding computers and cognition : a new foundation for design*, Addison-Wesley Publishing Company, Inc., New York (1986).
4. DoD, ''Trusted Computer System Evaluation Criteria'', (CSC-STD-001-83), US Department of Defence (15 Aug 1983).
5. J E Dobson and J A McDermid, ''Security Models and Enterprise Models or Information Flow Models considered a Denial of Service Attack on Computer Security'', *Proceedings of The 1988 Workshop on Database Security*, Kinsgton, Ontario,Canada (Oct 5-7, 1988).
6. Wiebren do Jonge, ''Compromising statistical databases responding to queries about means'', *ACM Trans. on Database Systems* **8**(1), pp. 60-80, ACM (March 1983).
7. *The Beast of Business*, author and publisher unknown.
8. E.P. Thompson, *The Making of the English Working Class*, Pelican Books (1968).