Decision, Contract and Emotion: Some Economics for a Complex and Confusing World

W. Bentley MacLeod*
Boston College and C.R.D.E., Université de Montréal

May 31, 1996
This version June 16, 1996

Abstract

This essay illustrates that if Savage's small world assumption is relaxed, one can construct a theory of bounded rationality that incorporates some of the insights from recent work in cognitive psychology. The theory can be used to explain why contracts are incomplete and the existence of endowment effects in exchange.

JEL Classification: D0, J3.

Keywords: bounded rationality, decision theory, complexity

*This paper was prepared for the Harold Innis Memorial Lecture, Canadian Economics Association Meetings, St. Catherines, Ontario, May 31, 1996. I would like to thank Susan Atheny, Robin Boadway, Janet Currie, Lorne Carmicheal, Ross Finnie, Peter Gottschalk, Edi Karni, Jim Malcomson, Eric Maskin, Vernon Smith and Jean Tirole for stimulating discussions on bounded rationality. I am also grateful for the comments from seminar participants at the Boston College Brown Bag Seminar, the Harvard-MIT theory workshop, and the University of Toulouse. The financial support of the S.S.H.R.C. of Canada for this research is gratefully acknowledged.

"..the notion of utility is raised above the status of a tautology by economic theories as make use of it and the results of which can be compared with experience or at least with common sense."

From Theory of Games and Economic Behavior by John Von Neuman and Oskar Morgenstern.

1 Introduction

One of Adam Smith's important insights is that the market economy is a system that coordinates activities among a large number of self-interested individuals. Since the publication of the *Wealth of Nations* this version of *homo economus* has been greatly refined to include the fact that individuals make decisions based on models of the way that other members of *homo economus* response to economic events (see for example Lucas (1976)'s famous critique of public policy). Yet as the model of rational choice is further refined it is also increasingly rejected by the data, both in laboratory experiments (Kagel and Roth, eds (1995)) and with economic data (Deaton (1992)). This suggests a need to take up again the program on bounded rationality begun by Simon (1982). The purpose of this essay is to illustrate one way that the economist's model of rational choice can be modified to incorporate some of the exciting advances that have occurred in our understanding of the human mind[1], and to apply this new model to problems of contract formation and bargaining.

The standard model of rational choice begins with the expected utility model of Von Neuman and Morgenstern (1944) and Savage (1972). The conventional approach to resolving observed violations of expected utility maximization is to explore modifications of Savage's axioms, especially the sure-thing principle (see Camerer (1995)). The theme of this essay is that the focus of our attention should be on what Savage calls the "small world assumption". This assumption supposes that the decision maker understands all the consequences of her actions, and is able to assign a probability to each state of nature. Dropping this assumption opens the door to a new set of potentially testable economic models.

The importance of the small world assumption has recently been highlighted in the work of Tirole (1994) and Maskin and Tirole (1995)[2]. They highlight the tension that exists between the model of *homo economus* in Savage and the

[1] See Kosslyn and Koenig (1995) for an excellent review of recent advances in cognitive neurosciences.

[2] An essay of this breadth must necessarily be selective, and so I do not discuss some interesting work on incomplete contracts that accepts the small world assumption, including Bernheim and Whinston (1995) and Segal (1995). See Anderlini and Felli (1994) for an approach that explores the consequence of assuming the contract is computable. Important work on bounded rationality that I do not discuss here includes Osborne and Rubinstein (1994) and Kalai and Stanford (1988) with applications to game theory, and Gilboa and Schmeidler (1995) with an application of case-based decision making. Sargent (1993) reviews models of bounded rationality that have been applied to macro economics.

assumptions underlying the recent work on incomplete contracts that began with Grossman and Hart (1986) (see Hart (1995) for a review of this literature). This is a serious criticism given that this literature has provided new and important insights into the theory of the firm. In section 3 I show, using a modification of the multi-tasking model of Holmström and Milgrom (1991), that contracting over goods and services with multi-dimensional attributes is impossible due to the inherent environmental complexity. The expected value of the relationship cannot be computed, and dynamic programming, as required by Maskin and Tirole (1995), cannot be carried out. Hence bounded rationality arising from environmental complexity may provide a basis for models that begin with the hypothesis that complete contracts are impossible. I also illustrate that this approach can be used to provide some testable implications for the form of compensation contracts.

In the 'large world' considered in this paper it is not possible for individuals to be completely rational in the sense that they are able to think about all the events that are relevant to a decision. The question then is how do sensible people make decisions? Section 4 introduces a model of choice based on the recent work of Damasio (1995). He reviews a set of cases studies based on individuals with damage to the frontal lobe of the brain, and concludes that emotions are an essential part of good decision making. I formalize this idea using a model of heuristic choice taken from the literature on artificial intelligence (see Nilsson (1980) and Pearl (1984)). The result is a model of decision making that is a compromise between search (rational thought) and intuitive (or emotional) judgement. The paper finishes with an application of this emotional decision making approach to the problem of bargaining and property rights based on Carmichael and MacLeod (1996). We show that fairness, endowment effects, and intransitive preferences are natural outcomes of an evolutionary model in which the choice between thought and emotional commitments is endogenous.

2 Choice Theory

2.1 The Model of Expected Utility (Von Neuman and Morgenstern (1944), Savage (1972))

The basic postulates of Savage's theory of expected utility are found inside the front cover of his book "*The Foundations of Statistics*". Consider the following summary of the base model:

1. There are a finite number of states of the world, S, with subsets $A, B, C...$ describing possible events[3].

2. A set of consequences, F, with elements $f, g, h...$.

[3] The assumption of a finite number of elements is actually not consistent with postulate P6 of Savage's model, which is needed to ensure that the subjective probability measure is unique. However, here I am concerned with the weaker property of existence of a measure that in principle is more easily satisfied with a finite state space.

3. Acts, $\mathbf{D} = \{f | f : S \to F\}$, consisting of arbitrary functions from the states of nature to consequences.

4. Complete, transitive preferences $\succsim$ over $\mathbf{D}$.

The notion of a state in 1 is very general. It includes, for example, the set of natural events that may occur at any date in the future, and hence the model is rich enough to deal with dynamic decision processes. By assuming that time is discrete and horizons are finite, we may suppose that the set of states is finite. The set of consequences are the things that individuals care about such as money or goods consumed. For example it might be the clothes one wears each day.

Item 3 is the set of acts: functions from the set of states to consequences. Individuals are assumed to make choices from the set of acts. For example in the case of clothing an act might be the function describing the clothes to wear at each date as a function of the weather in the morning. In choosing an act an individual is selecting a plan for future consumption as a function of events. It is in this way that Savage's model explicitly assumes foresight on the part of agents. Preferences are then defined over the set of acts.

Savage outlines a set of postulates for preferences that imply that individuals form beliefs on the likelihood of each state and assign a value or utility to each consequence. Preferences over acts are then derived from the expected utility that is computed for each act. If a decision today has an implication for consequences tomorrow, then the agent must take these implications into account when computing her expected utility.

This model is the basis for much of modern economics[4]. Though many modifications of the model have been studied in recent years, one assumption that is rarely questioned in economics is the ability of agents to think about all the different possible acts[5]. Letting n denote the number of states and m the number of consequences, then the number of possible acts is m^n. As the tables later in the essay illustrate, exponential functions of this form grow so quickly that for even moderately sized problems thinking about all the different acts would take centuries.

Savage was well aware of this limitation of the model, and related it to the "Look before you leap principle". He states that:

> "...the look before you leap principle is preposterous if carried to extremes."[6]

In particular his objective was the construction of a theory for "small worlds", that is situations for which the state space is relatively small. Section 5.5 of *The Foundations of Statistics* outlines some sufficient conditions under which the small world analysis is consistent in the "large world" of day to day life.

[4] See Epstein (1992) for a recent review of behavior under risk.

[5] An exception is Binmore (1993) who also points out that the 'small world' assumption is often too strong.

[6] Page 16, Savage (1972).

The crucial ingredient is that the decision maker is able to construct what Savage calls a microcosm in which each state corresponds to a large world event. In section 5.5 he illustrates how his postulates may be modified to ensure that decision making in the microcosm is consistent with rational choice in the large world.

One of the basic requirements is that the set of states in the microcosm form a partition of the large world state space[7]. Given the size of the large world, this is a condition that I shall argue is unlikely to be satisfied in many economic domains. In other words there exist states or events that may occur with positive probability, that are relevant to the decision at hand, but yet the decision maker has not constructed a plan of action for this state or event. I suggest that this is a common occurrence in day to day decision making that leads to some interesting implications.

It is instructive to recall Savage's discussion of the Allais paradox. Individuals are asked to choose between the lotteries given in Table 1[8]:

	Situation 1:
Choose between:	
Gamble 1.	\$500,000 with probability 1;
Gamble 2.	\$2,500,000 with probability 0.1,
	\$500,000 with probability 0.89,
	status quo with probability 0.01
	Situation 2:
Choose between:	
Gamble 3.	\$500,000 with probability .11,
	status quo with probability 0.89;
Gamble 4.	\$2,500,000 with probability 0.1,
	status quo with probability 0.9.

Table 1: Choices for the Allais Paradox

When first faced with these choices Savage chose gamble 1 over gamble 2 in situation 1, and gamble 4 over gamble 3 in situation 2. However these choices are not consistent with expected utility maximization or more precisely the sure-thing principle because situation 1 is transformed into situation 2 by substituting the prize of \$500,000 with probability .89 with the status quo outcome for each gamble. Given similar laboratory experiments, a great deal of research has studied ways to relax the axioms of expected utility so that observed behavior is consistent with utility maximization[9].

What is interesting is Savage's response upon realizing that his choices did not satisfy the sure-thing principle: he changed his choices and decided upon gamble 3 in situation 2. A cynic might suggest that he was only trying to protect his theory. An alternative view is that even these apparently simple choices

[7] Ibid, page 85.

[8] Ibid, section 5.6.

[9] See Camerer (1995) for a nice review of the experimental literature on choice. Machina (1987) and Epstein (1992) review the literature on non-expected utility theory.

involve subtleties that may lead even a great statistician like Savage astray. In other words, bounded rationality is likely to be ubiquitous, and responses to apparently simple choices in the laboratory do not necessarily reflect underlying preferences. The next section considers some evidence on how individuals think about problems, and why they make mistakes.

2.2 Hard simple problems

An issue that has attracted a great deal of attention is whether logic or reason is innate to the thought process. Johnson-Laird (1983) presents results of laboratory experiments in which individuals are asked to state all the logical implications of a syllogism. Bounded rationality would imply that they may not find all of the true implications, however if they are natural deductive reasoners then all the conclusions they arrive at would be correct. For example consider the following syllogism:

- None of the authors are burglars.
- Some of the chefs are burglars.

Possible conclusions are:

- None of the authors are chefs.
- None of the chefs are authors.

In a laboratory experiment 6 out of 20 subjects reported one of these two incorrect conclusions. Seven of the twenty subjects arrived at the only correct implication:

- Some of the chefs are not authors.

From this and other evidence Johnson-Laird concludes that logical reasoning is not an innate ability. He proposes a theory based on the idea that individuals use mental models to think about the world. In the case of syllogisms one can think in terms of Venn diagrams, as is illustrated in Figure 1. Johnson-Laird suggests that one of the reasons individuals make mistakes is that they consider only one or two of the possible models. In particular the first Venn diagram illustrated is consistent with the incorrect conclusions drawn above.

To test this idea Johnson-Laird studies the likelihood of an error as a function of the number of possible true models consistent with a given syllogism. The results for one of these experiments are illustrated in Table 2. Observe that the error rate increases with the complexity of problem as measured by the number of models needed to represent the syllogism.

This suggests that even in relatively simple problems individuals commit many errors. Furthermore, the error rate increases with the complexity of the mental representation of the problem. When thinking about a syllogism the individuals must consider the consistency of a number of different models. What

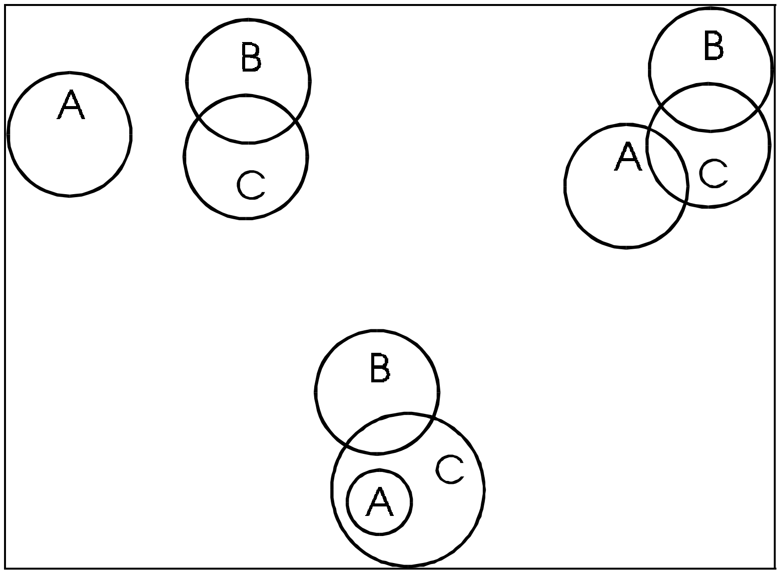

Figure 1: Possible Models for Syllogism: A = author, B = burgler, C = chef

	Premises requiring one model	Premises requiring two models	Premises requiring three models
Experiment 1	92%	46%	27%
Experiment 2	80%	20%	9%
Experiment 3	62%	20%	3%

Taken from Johnson-Laird (1983), page 117.

Table 2: Percentages of Correct Reasoning in Experiments on Syllogistic Reasoning

is happening is that individuals are *searching* over the set of possibilities. Those who are making errors have not completed a careful search over all possible models. More generally, as Newell (1990) states:

> "Search is fundamental for intelligent behavior. It is not just another method or cognitive mechanism, but a fundamental process."

Viewing the thought process as search provides a natural way to think about bounded rationality. If the space of possibilities is sufficiently large (as I shall argue is the norm) then search is incomplete, and hence individuals may not have discovered the correct answer or the best decision. The model of bounded rationality I consider is one that begins with a description of the potential state space, proposes a method for searching over this state space, and ends with some rules for decision making.

What the results of Johnson-Laird suggest is that even if we do not completely understand the decision process, we can make testable predictions about behavior. In his model we expect the error rate to increase as the complexity of the environment increases. I shall argue that even apparently simple decisions involve a level of complexity that makes exhaustive search impossible, and hence rational decision making is not based on a complete ranking of all alternatives.

To see this consider the following common decision problem: What should one wear? Suppose that we are rational in the sense of Savage. That implies that we begin with all the alternatives and rank them according to our preferences. Suppose that 1/2 of a second is used to assign a utility to each choice. The time needed to think about all the possible combinations is presented in Table 3 for different possible wardrobes. In the first case a completely rational choice would take at least 40 hours! Observe that this problem can be greatly simplified by the use of combinations, such as suits or complete outfits (a process called *chunking*).

Item	Total Number	Choice with suits	Full Outfits
Shirt	10	10	20
Pants	10	10	1
Jacket	10	1	1
Tie	10	10	1
Socks	6	6	1
Shoes	4	4	1
No. of Decisions	240,000	24,000	20
Time needed	40 hours	4 hours	12 seconds

Table 3: Dressing

This dressing algorithm involving search over all alternatives is not what we do in practice, yet this is what economists do in theory. In a very simple environment brute force search of this type is likely to be effective, yet for many situations such an approach is not practical. In the next section I discuss the implications of environmental complexity for contract formation.

3 Bounded Rationality and the Theory of Contract

The idea that it is not possible to search over all possible choices can be used to explain why contracts are incomplete. Consider a model of the employment relationship between a worker and a firm that combines the multi-tasking model of Holmström and Milgrom (1991) with the repeated agency model of MacLeod and Malcomson (1989). Each period the following sequence of moves is carried out:

1. The principal and agent agree on compensation and expectations for performance (which may include the continuation of a previous agreement).

2. The state of the world $\omega_t \in \Omega$ is revealed.

3. The agent divides a time endowment of Y among k different tasks: $\mathbf{y}_t \in \Re^K$.

4. The principal pays the agent W_t.

5. Both principal and agent decide whether to continue the relationship or not.

The date is denoted by the subscript t, and K is the number of possible tasks. The twist upon the previous literature concerns the interpretation of the state of nature. Suppose that both the cost and benefit of different actions are unknown *ex ante;* for example a fireman may not know which house will catch fire; how difficult it will be to put out the fire; nor is he able to anticipate the set of actions that will need to be carried out upon entering the burning house. Let us begin by defining a state space that incorporates uncertain costs and benefits for each of the possible tasks:

$$\Omega = \left\{\left\{\alpha^1, ..., \alpha^n\right\} \times \left\{\beta^1, ..., \beta^m\right\}\right\}^k, \tag{1}$$

where $\alpha_k \in \left\{\alpha^1, ..., \alpha^n\right\}$ denotes one of n levels of productivity for task k, while $\beta_k \in \left\{\beta^1, ..., \beta^m\right\}$ represents one of the m cost levels for task k. The total benefit from an effort choice $\mathbf{y}_t$ is defined by $\boldsymbol{\alpha}^T \mathbf{y}_t$ (boldface represents a vector), while the total cost to the worker of producing this effort is

$$C\left(\mathbf{y}_t, \boldsymbol{\beta}\right) \equiv \sum_{i=1}^{K} \left(\beta_i y_{it}^2 - \delta\left(y_{it}\right) f\right). \tag{2}$$

The function $\delta\left(y_{it}\right)$ is 1 if y_{it} is positive and zero otherwise.

The benefits and costs have been modelled as functions, however it is explicitly assumed that a measurement system does not exist. Consider a secretary who carries out a variety of tasks, including typing, answering the phone, filing, making travel reservations etc. The costs and benefits for these different activities vary with the day to day demands of the office. For example, several people in the office may need to go to the same conference, raising the productivity of allocating time to travel plans, and resulting in a cutback in typing throughput. On the cost side, if the conference is during a busy period (for example college convocation), then one may have to call several hotels to find accommodation. Not only do these costs and benefits vary in an independent way from day to day, it is not clear (to me at least) how one would construct a measurement system to directly compare the costs and benefits of the different actions.

This does not imply that measurement systems never exist. In many jobs, particularly sales, compensation is based on a commission. In this case the level of sales accurately aggregates the efforts of an individual. In the current example, although I assume that a measurement system aggregating the costs and benefits does not to exist, I do assume that in principle agents are able to write explicit contracts. Such a contract permits the worker and the firm to agree upon actions and payments as a function of the different states. The lack of a measurement system implies that each state must be described explicitly in the contract[10]. This is common in many contracts. For example the contract for a singer at a concert may explicitly list acceptable reasons, such as laryngitis, that excuse the individual from providing the contracted upon services.

Thus I make the extreme assumption that in principle the worker and firm can agree upon a state contingent contract that specifies what should be done for each state of nature. Formally the contract is a function $c : \Omega \rightarrow X = \Re \times \Re^k$, where for each state $\omega \in \Omega$, the $c(\omega) = (w(\omega), y(\omega)) \in X$ defines the wage payment and the output expected from the agent. This assumption differs from the incomplete contracts literature where it is assumed that such a contract is impossible, while maintaining the hypothesis that individuals understand all the possible outcomes and can recontract based on the *ex post* realization of the state.

For this model an efficient complete contract, $c^*(\omega) = (w(\omega), \mathbf{y}(\omega))$, is the solution to the following program:

$$\mathbf{y}(\omega) \in \arg\max_{\mathbf{y}'} \boldsymbol{\alpha}\mathbf{y}' - C(\mathbf{y}', \boldsymbol{\beta}), \text{ subject to:} \tag{3}$$

$$|\mathbf{y}| \equiv \sum_{i=1}^{k} y_i' = Y, \text{ and} \tag{4}$$

$$w(\omega) = \bar{U} + C(\mathbf{y}(\omega), \boldsymbol{\beta}). \tag{5}$$

where $\bar{U}$ is the one period alternative utility for the worker. Following Townsend (1979) and Dye (1985) suppose that there is a cost for including additional contract contingencies, given by γ per contingency. For this multi-tasking model one has the following result.

Proposition 1 *The cost of implementing the complete contract procedure when all states occur with positive probability is* $n^k m^k \gamma$.

What is important to observe is that the cost of the contract is an *exponential* function of the number of tasks. The literature on computational complexity emphasizes the impossibility of implementing algorithms whose costs are exponential in the size of the problem (see Garey and Johnson (1979))[11]. To see

[10] This assumption can be contrasted with the agency approach to compensation as outlined in Baker (1992) and Holmström and Milgrom (1991). This work examines the optimal way to incorporate imperfect signals of worker performance into the pay package.

[11] Notice that this approach is very different from Anderlini and Felli (1994) who use computability as a criterion. In this model complete contracts are assumed to be computable in the sense that a contract can be agreed upon in finite time. The important point, as emphasized in the computer science literature, is that an algorithm is only useful if a solution can be achieved in a "reasonable" amount of time. If the problem is of exponential complexity, then even moderately sized problems take several centuries to solve.

why this is the case suppose that $\gamma = 1$ *cent*, and that the number of cost and performance levels are the same ($n = m$). Table 4 presents the costs of the complete contract as a function of the number of tasks and effort levels.

Number of Cost and	Number of Tasks			
Performance Levels	2	5	10	15
2	\$0.16	\$10	\$10,000	\$10 million
3	\$0.81	\$600	\$35 million	\$2 trillion
4	\$2.56	\$10,000	\$11 billion	\$11,000 trillion
5	\$6.25	\$100,000	\$1000 billion	\$10 million trillion
Cost of a contract clause:	1 cent			

Table 4: Cost of a Complete State Contingent Contract

As one can see, the use of a complete contract when there are more than say 10 tasks is impossible. Furthermore, given that these costs reflect the number of underlying states, dynamic programing is impossible because one could not compute the expected value of the relationship. Observe that piece rate contracts correspond to basing compensation on one dimension of output. In this simple setup complete contracts are very inexpensive; hence they should be observed when the number of tasks to be measured is small. I now turn to the governance of the employment relationship when a complete contract is not possible.

A solution, outlined in MacLeod (1996), begins with Simon (1951)'s distinction between an employment relationship and a sales contract. A sales contract is one for which the terms and conditions are agreed upon before the state of nature is revealed, while an employment contract delays determination of the appropriate action until after the state is reveal. In the context of the multitasking model proposition 1 demonstrates that a state contingent sales contract has a cost that is exponential in the number of tasks, and hence it is necessarily incomplete when there are a moderate number of tasks. The computational complexity of the contracting relationship can be dramatically reduced by delaying decision making until after the state of nature is revealed; at this point, one has only to determine the appropriate action for a single state.

Notice that a fixed wage or salary contract provides exactly such a solution. A secretary is paid a fixed salary and is retained as long as her performance is satisfactory. The definition of satisfactory is never explicitly outlined before the fact, but is judged after the secretary's response to day to day events is observed. Another example is the academic tenure decision. At the time an assistant professor is hired some general criteria for promotion may be outlined, but rarely are a specific set of necessary and sufficient conditions for promotion agreed upon. Rather, at the end of a fixed period of time, the work of the individual is judged and compared to his or her peers, and then a decision is made to continue or terminate employment.

In the context of the current model, what is crucial is that the evaluation of performance given the state of nature is orders of magnitude less complex than trying to specify acceptable performance for every conceivable eventuality.

In this case compensation takes the form of a fixed wage payment, with the decision to keep the worker conditional upon an *ex post* evaluation of output. Also observe that in principle a complete contract could be written *ex ante*, however the complexity of the state space makes this impossible. However once the state of nature is revealed then both the worker and the firm have symmetric information. This is consistent with assumptions that are typical in the incomplete contracts literature discussed in Hart (1995). In this model individuals cannot compute the expected value of the relationship because it is not possible to have well defined beliefs over all the relevant states. Hence the model does not satisfy the condition that Maskin and Tirole (1995) identify as key to their demonstration of the equivalence of complete and incomplete contracts models.

Another solution to the incentive problem is the reallocation of property rights so that one individual is given residual control over non-contractibles. When the technology is such that property rights can be reallocated so that the output of each individual can be marketed, we are then back to the one dimensional model for which complete contracts are possible. However, if there is team production in the sense of Alchian and Demsetz (1972), then such reallocation of property rights is not possible, resulting in an employment contract with incentives provided after an *ex post* evaluation of worker performance.

To summarize, when a job involves a significant amount of multi-tasking then it may be more efficient to agree upon a fixed wage at the beginning of the period, but to delay judgement of the worker's performance until the end of the period. MacLeod and Malcomson (1989) shows the enforcement of a contract requires the existence of a rent that is related to the amount of contract incompleteness. The number of tasks provides a way to parameterize the level of contract incompleteness, and hence should be related to the level of rents in a relationship. From an empirical perspective this is problematic because rents are generated in a number of ways, including mobility costs, firm specific investments and reputations, all of which are notoriously difficult to measure.

However, MacLeod and Malcomson (1989) also show that generically there are two types of contracts[12]. The first type of contract is an efficiency wage that pays a worker more than her market alternative, and when combined with the threat of dismissal, provides an incentive to perform. A second type of contract is performance pay given at the end of the period in the event that the worker achieves an acceptable level of performance. The existence of bonus pay provides a potential test of the theory outlined above.

In standard principal-agent analysis there is no discretionary bonus pay. Moreover, the analyses of Baker (1992) and Holmström and Milgrom (1991) suggest that there would be less use of performance pay with increased complexity due to the incentives to game the reward system, while this model predicts the opposite, namely bonus pay should be correlated with the complexity of the job. With less complex jobs there is either an explicit contract such as a piece

[12] MacLeod and Malcomson (1996) show how the form of contract used can be explained by the form of technology and market forces.

rate, or the level of incompleteness is sufficiently small that turnover costs alone are sufficient to provide incentives.

Though labor economists often assume that individuals are paid a fixed hourly wage, bonus pay is in fact widespread, with 20% to 30% of workers receiving some form of non-contingent bonus pay[13]. Such bonuses would not be paid in the world of complete contracts, hence the complexity model provides some predictions on the *incidence* of bonus pay. Table 5 presents results from ongoing work with Daniel Parent for a simple probit on the incidence of piece rate pay and bonus pay as a function of different job characteristics (using the National Longitudinal Survey of Youth)[14].

One would expect the incidence of bonus pay to be positively correlated to job complexity, with the reverse holding for piece-rate contracts. The reason for the later result follows from the fact that complete contracts are feasible only when the number of tasks (or performance measures) is low. Hence increasing job-complexity would make it less efficient to implement an explicit piece-rate, while increasing the incentive to use bonus pay in some situations. Notice that the coefficients for jobs requiring autonomy are positive and significant for bonus payment and of the opposite sign for piece rate pay, as we would expect. Workers are also asked if their job requires completing a well defined task. The effect of this is reported in the second line, where the effect is negative for bonus pay incidence and positive for the incidence of piece rates. Finally, increasing job responsibility has a positive effect on the use of bonus pay, though its effect on piece rate incidence insignificant.

Goldin (1986) also provides some interesting evidence on the relationship between job design and compensation at the turn of the century. For some production processes, particularly for coats and cigars, firms have a choice between one of two forms of organization. In one case different persons produce different parts of the product, which are then assembled in a final step. In this case the output of each worker in the first step is simply the number of pieces produced, and hence it is efficient to use a piece rate contract for this group of workers (usually women).

The second possibility is to produce the product in a single step. This is the preferred method when a higher quality product is desired. What is observed is that compensation is based on time, rather than output. Goldin argues that the choice is based on relative costs of monitoring, with the one step processes requiring more supervisory input. That begs the questions why it is not possible to use a complete contract that pays by the piece condition on a minimum quality standard for output? Within the context of the current model I would suggest that high quality cigars and coats depend on several characteristics that make it difficult to contractually specify acceptable quality. In that case the efficient form of compensation is a time rate combined with some form of implicit reward at the end of the period based on a subjective

[13] See Brown (1990) for theory and evidence regarding piece rate pay. See Milgrom and Roberts (1992), Gibbons (1995) and Lazear (1995) for some excellent overviews of the literature on performance pay.

[14] MacLeod and Parent (1996)

evaluation of performance. This later prediction seems to be consistent with the observed faster increase in earnings for these jobs, that Goldin argues is a form of delayed compensation.

More work must be done before we fully understand the relationship between jobs design and compensation, however I find it encouraging that this approach can produces potentially testable hypotheses that seem to be borne out by the data.

Independent Variable	*Bonus*	*Piece Rate*
Autonomy	0.5989	-0.3256
	(0.1625)	(0.1843)
Complete Task	-0.3328	0.6998
	(0.1440)	(0.1787)
Variety	-0.2125	-0.0042
	(0.0999)	(0.1112)
Increase in Responsibility	0.0110	-0.0070
	(0.0084)	(0.0099)
Schooling	0.0237	0.0133
	(0.0091)	(0.0110)
Local Unemployment	-0.0924	-0.0530
	(0.0250)	(0.0250)

Standard erros are in parenthesis. Sample size is 8,045.
Other covariates include sex, race, rural/urban, industry and occupation dummies, labor market experience, experience squared, tenure, tenure squared, regional dummies, and an intercept.

Table 5: Method of Pay Probits

4 Emotion and Choice

The argument thus far is that the world in which day to day decisions are made is so complicated that an exhaustive search of all the possibilities cannot be carried out. In essence completely rational thought is impossible. How then do we make decisions? Damasio (1995) brings together some impressive evidence suggesting that emotions are an integral part of decision making. In one case, an individual he calls Elliot, had a brain tumor resulting in extensive damage to the frontal lobe of the brain. What is interesting about this case is that after the operation Elliot performed well on a set of standard tests of perceptual ability, past memory, short term memory, new learning, and arithmetic ability. Yet his life was a mess. After his operation he lost his job, and has a series of failed projects.

At issue is why an apparently able individual was unable to manage day to day life? Damasio observes that Elliot lost the ability to feel or experience emotion, and with that loss he also lost what we might call judgement. In *Descartes' Error* he argues that emotion is an essential part of decision making (an effect based on what Damasio calls the *somatic marker hypothesis*). This can be illustrated with the dressing example discussed above. In the morning we clearly do not consider all possible combinations, rather we quickly eliminate many possibilities because they do not "look right". Essentially we make decisions before we are aware of all the possibilities. If we did not, then dressing would be an impossible task that we would do once a week or once a year.

These observations are consistent with the models of decision making developed in the early years of artificial intelligence. For example both Newall and Simon (1995 (orginally published 1963)) and Minsky (1995 (orginally published 1963)) view the process of thought as a combination of search and pattern recognition. The process of pattern recognition can be equated with emotional responses or intuition. For example chess players carry out a certain amount of search, and at the point they stop searching they attempt to evaluate the strength of the position. No matter how much reasoning a person uses, at some point in a chess game intuition concerning the strength of the position must be used to evaluate the strength of a position (except at the end of the game when a clear victory can be established using backwards induction). What Damasio demonstrates is that the source of this high order thought originates in the frontal lobe, which is also the source of emotion. Hence he concludes that there does not exist a clear demarkation between pure reason and pure emotion.[15]

Rather one has a continuum between spontaneous decision making and the careful evaluation of as many options as possible. How then are we to model the process of decision making? One approach, outlined in MacLeod (1995), is based on heuristic search algorithms widely used in artificial intelligence (See Newall and Simon (1972), Nilsson (1980) and Pearl (1984)). In this model rather than begin with an *a priori* description of the full state space, as is done is standard decision and game theory, one begins with the rules for *constructing* the state space. In the dressing example discussed above the construction of the state space proceeds with sequential addition of articles of clothing, say pants, then shirt, then jacket. A heuristic algorithm adds an additional element: a function that assigns a value to search in an particular direction.

To see how the process of heuristic search might proceed, view the dressing example as a search problem on a graph, as illustrated in Figure 2. In this problem the individual first considers the choice of a shirt, indicated by node n_1 or n_2. A heuristic is a function h that assigns a number $h(n_i)$ to each node reflecting the value of search in that direction. In this example I suppose that the individual prefers dressing with a white shirt, and hence $h(n_2) > h(n_1)$. This is followed by considering which pair of pants to wear with the shirt. The process continues until fully dressed with a white shirt, grey pants and a blue

[15] See also the work by Baron-Cohen (1995) on autism. He argues that much of day to day interaction involves higher level modelling of our environment, including the intentions of others. Autistic individuals do not have this ability, and accordingly he calls them *mindblind*.

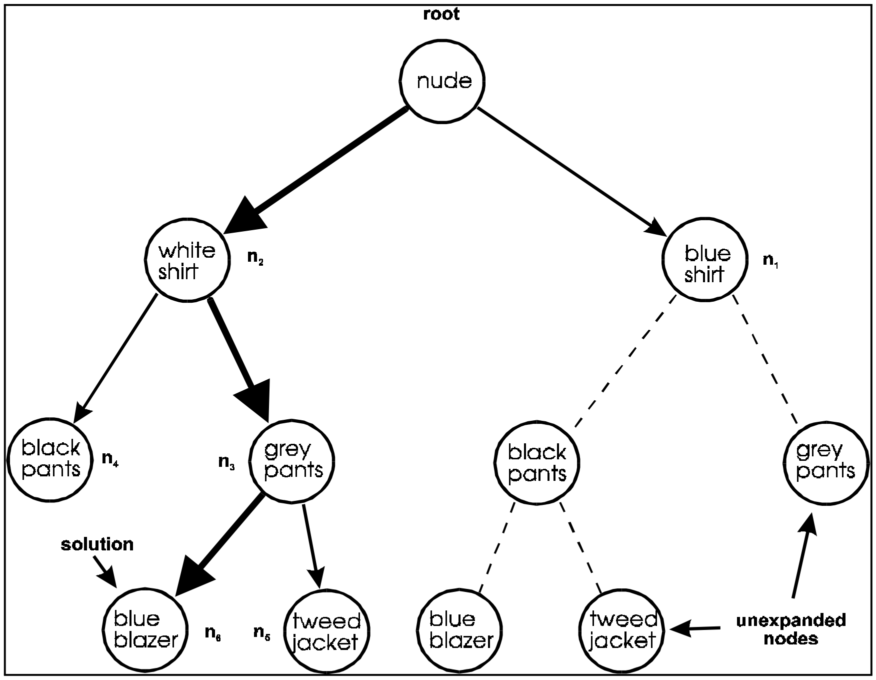

Figure 2: The Dressing Graph

blazer[16]. If after being fully dressed, the result is a payoff that is lower than $h(n_1)$, then the individual would begin considering the consequences of dressing with a blue shirt. Search continues until a terminal node is reached with a value that the higher than $h(n_i)$ for all nodes n_i that have been explored.

In practice the heuristic is an easily computable function that is hopefully correlated to the value of continued optimal search. One can show that the complexity of the heuristic search algorithm very much depends on the quality of the heuristic function. If agents were unboundedly rational then the heuristic would be the value function computed from choosing optimally after the expansion of a node. In that case backtracking to explore other parts of the graph would never occur because the heuristic would always accurately reflect the value of search in a particular direction. MacLeod (1995) shows that for any problem, there exists a heuristic algorithm that results in an optimal strategy. In many situations, such as in the game of chess, computing the value function on which the optimal heuristic would be based is impossible.

My understanding of the literature on heuristic search is that for many problems there exist heuristic functions that work well in practice, but the theory is not yet well developed. For my purposes the heuristic search model

[16] There are several possible algorithms for heuristic search of a graph. See Pearl (1984) for a review. MacLeod (1995) uses a version of best-first search called A*.

provides a useful paradigm model for boundedly rational decision making that generalizes the model of Simon (1955). An interesting aspect of the heuristic search model is that the process of search is carried out by searching *forward* over the graph, not backwards as is done with dynamic programing. This is consistent with experimental evidence from Camerer, Johnson, Rymon and Sen (1993) using a repeated bargaining game. The twist on the standard experimental setup is that the individuals examine the consequences of their decisions by opening windows on the computer screen corresponding to different possible moves. Thus it is possible for the experimenter to record some of the steps in the thinking process of an individual. Camerer, Johnson, Raymond and Sen find that typically individuals began their examination of the game tree by proceeding *forward* and studying the consequences of moves today, rather than beginning with an examination of the terminal nodes of the graph.

Johnson, Camerer, Sen and Rymon (1996) find additional evidence supporting the hypothesis that agents are boundedly rational in their analysis of the bargaining game. In addition they also find that some individuals who had been trained in dynamic programming use it to solve the game, and came closer to the subgame perfect equilibrium. Those individuals that did not, tended to choose outcomes that were closer to an equal split of the available surplus. As in the example of Savage's play when faced with the Allais paradox, we see that the way an individual models a situation can have a large effect on the observed outcome. This evidence suggests that even when individual preferences are stable, the strategies that they choose may vary greatly depending on previous experiences and on their ability to understand the consequences of their actions.

The appropriate way to think about or model heuristics is consequently a very difficult unanswered question. MacLeod (1995) explores the use of experience with different states of the world as a way to update heuristic choice, resulting in a model of learning that is consistent with empirical learning curves (Newell and Rosenbloom (1981)). Though the results are still very crude, the model of heuristic search highlights the necessary trade-off between search, and judgement or intuition that Damasio (1995) links to processes that occur in the frontal lobe of the brain. The amount of reason that one would use in a particular decision situation depends on the environment. I now turn to this issue in the context of a very simple bargaining game.

4.1 Emotion, Fairness and Conflict

Once one drops the assumption that individuals are fully rational, one is faced with the puzzle of how to model the "boundedly rational" behavior we observe. A solution to this problem is found in the work of William Hamilton, George Willams, Robert Trivers and John Maynard Smith[17]. They argue that is it possible to explain observed behavior, including apparently irrational behavior

[17] For an excellent discussion of the contributions of these scientists, as well as a full set of citations see the book by Wright (1994).

such as altruism, based on the selection of the most fit set of behaviors in interacting populations facing evolutionary selection pressures. This has led more recently to the application of models with evolutionary selection to decision making.

Frank (1988) has elaborated this idea in the context of economic exchange relationships, and argues that the ability to signal anger can permit more efficient levels of exchange. In his model emotions are viewed as commitment devices that are distinct from rational thought. In a world with bounded rationality there does not exist a clear demarcation between search behavior (rational thought) and immediate action (emotion). Rather decision making in general involves trade-offs between thought and action. Given that thought is costly, then we should observe it only if it is successful.

Carmichael and MacLeod (1996) look at this question using an extension of the Nash demand game that takes a highly stylized view of the trade-off between thought and action. We suppose that individuals meet repeatedly to agree upon the division of an amount S. Before meeting they may decide on a non-negotiable reservation demand that we call an emotional commitment. If the demands of players 1 and 2 are given by σ^1 and σ^2, then we suppose that after agents meet they bargain 'rationally' over $S - \sigma^1 - \sigma^2$. By rational we simply mean that there is some bargaining protocol that results in each agent receiving a fixed share of the surplus and a final payoff of:

$$U^i = \sigma^i + \frac{\theta^i}{\theta^i + \theta^j}\left(S - \left(\sigma^i + \sigma^j\right)\right), \tag{6}$$

if $S - \sigma^1 - \sigma^2 > 0$ and zero otherwise. In this model θ^i represents the level of rationality or intelligence that a person possesses.. Thus the game has the characteristic that if neither agent makes any emotional demands the result is efficient with the more able individual receiving a larger share. Demanding one more dollar *ex ante* ensures that one gains that dollar, however it may also increase the likelihood of a breakdown with zero for both players.

We show that the efficient and evolutionary stable equilibrium has each agent demanding $S/2$[18]. In this anonymous market fair division is the one outcome that ensures an absence of conflict, with emotional commitment permitting the weaker player to obtain a larger share of the surplus. In the context of an evolutionary game the equal split result is not surprising. If all individuals are the same, and only the most fit survive, then no asymmetric outcome that gives more to one group than another can be part of an equilibrium. It does serve to highlight the role that fairness plays in reducing conflict between individuals, a result that is consistent with the interpretation of the laboratory evidence presented in Speigel, Currie, Sonnenschein and Sen (1994) and Smith (1994). They argue that fairness seems to be motivated by a desire to achieve a speedy agreement and avoid conflict.

[18] See Ellingsen (1995) for a similar result that uses a different notion of rational behavior. See also Young (1993) for a theory of social convention that arrives at a similar result. In his model bounded rationality is used to select a unique equilibrium, while in our model the level of rationality is endogenous (albeit using a highly stylized model).

The situation changes dramatically if we add some uncertainty in terms of the contribution to the surplus. Suppose that each agent has an endowment e_i, and the surplus to be divided is given by

$$S = e_1 + e_2. \tag{7}$$

Further suppose that at the time a person makes an 'emotional commitment' the size of the other person's endowment is unknown. In that case the unique efficient equilibrium has each person demanding their endowment. In essence individuals use their endowments as a starting point for negotiation (see Kahneman, Knetsch and Thaler (1986) Kahneman, Knetch and Thaler (1990) for some laboratory evidence).

The result can be extended to multi-product trade. The equilibrium demands for a single agent are illustrated in Figure 3. In this case each agent demands at least her endowment, and strictly more than her indifference curve away from the endowment. This ensures that when meeting another person there is also an agreement to consume at least ones endowment. However it has the implication that observed preferences are intransitive, even though underlying preferences are transitive.

This may help explain observed endowment effects that result in intransitivities in trade. Endowments E_1 and E_2 in Figure 4 lie on the same indifference curve. Consider the amount of money that a person with E_1 is willing to pay to increase her consumption of the good to the level E_2, as is indicated by the bundle E_3. This amount is less than the amount necessary to get an individual at E_2 to give up her consumption good to the level of E_1 Therefore individuals reveal a greater preference for goods that are part of their endowment, an effect that is consistent with the evidence documented by Kahneman et al. (1990).

5 Conclusion

In this essay I have argued that many economic institutions are shaped by the necessity of making decisions in the large, complex world of day to day life. Formally, I have suggested that we may think about these decision problems by explicitly considering the implications of relaxing Savage's "small world" assumption. I have illustrated two applications of this idea to problems of contract.

In the small world optimal contracts are very complex and should incorporate efficiently all information relevant to a relationship[19]. I have argued that in many cases the employment relationship is so complex that it is more efficiently governed using an incomplete contract that leaves unspecified many important aspects of the relationship. In these cases it is more efficient to judge the quality of a worker's performance after the state of the world has been revealed.

A second application to trading relationships illustrates a very different kind of decision making. Environmental complexity implies that rational thought is

[19] For example see Holmstrom (1979) and Holmström (1982).

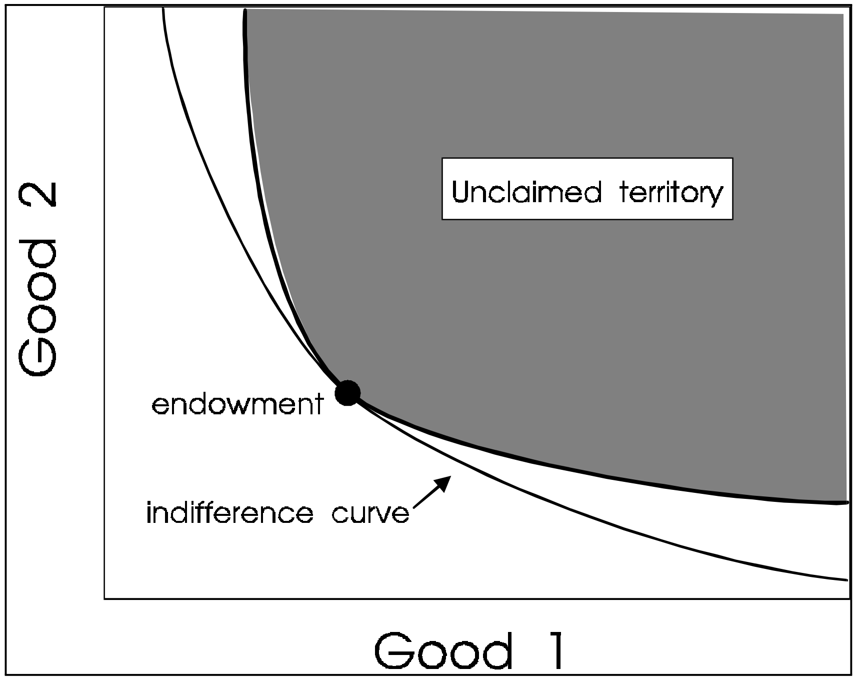

Figure 3: Demands with Territorial Bargaining

itself costly and hence subject to the pressure of natural selection. In work with Lorne Carmichael, we have shown that notions of private property arise naturally in a model where the level of rationality is endogenous. As Epstein (1995) has recently emphasized, simple rules such as the respect for private property, provide the basis for the efficient organization of production in a complex world.

This approach has some controversial implications for empirical work. Taking the model of rational choice seriously implies that the empirical specification of an economic model should be derived from a tightly specified model of optimal choice. *Homo economus* herself in such a model constructs a simple model of her environment that implies a great deal of sensitivity to new information. It is of course recognized that individuals do in fact make mistakes, however these are usually assumed to be an order of magnitude less than the variations in behavior implied by optimal choice.

If the arguments in the essay are correct, this implies that individuals are likely to be making large and possibly systematic errors, whose magnitude varies with the complexity of the environment. On average, even in a complex environment, we expect individuals to respond to economic incentives, and hence we should expect the sign of a response to treatment[20], as estimated with a linear regression model, to agree with the prediction of standard choice mod-

[20]For example a tax policy change or a change in relative prices.

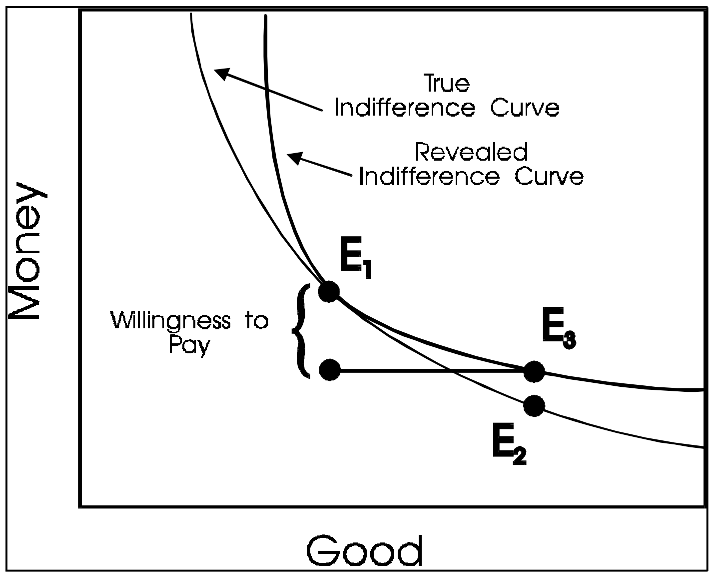

Figure 4: Territorial Bargaining and the Endowment Effect

els. The fact that errors may be large and unpredictable may help explain why it has been so difficult to progress from linear regression models to non-linear structural models that are very sensitive to the nature of decision errors. Hence the problem of bounded rationality, and the relaxation of Savage's small world assumption in particular, may have potentially important implications for both the theory and practice of economics.

References

Alchian, Armen and Harold Demsetz, "Production, Information Costs, and Economic Organization,," *American Economic Review,*, December 1972, *62* (5), 777–795.

Anderlini, Luca and Leonardo Felli, "Incomplete Written Contracts: Undescribable States of Nature," *The Quarterly Journal of Economics*, November 1994, *109* (439), 1085–1124.

Baker, George, "Incentive Contracts and Performance Measurement," *Journal of Political Economy*, June 1992, *100*, 598–614.

Baron-Cohen, Simon, *Mindblindness: An Essay on Autism and Theory of Mind*, Cambridge, MA: MIT Press, 1995.

Bernheim, B. Douglas and Michael D. Whinston, "Incomplete Contracts and Strategic Ambiguity," March 1995. mimeo Harvard University.

Binmore, Ken, "De-Bayesing Game Theory." In Binmore, Kirman and Tani, eds (1993).

——, **Alan Kirman, and Piero Tani, eds**, *Frontiers of Game Theory*, Cambridge, MA.: The MIT Press, 1993.

Brown, Charles, "Firm's Choice of Method of Pay," *Industrial and Labour Relations Review*, 1990, *43* (3), 165–182.

Camerer, Colin, "Individual Decision Making." In Kagel and Roth, eds (1995) chapter 8.

Camerer, Colin F., Eric J. Johnson, Talia Rymon, and Sankar Sen, "Cognition and Framing in Sequential Bargaining for Gains and Losses." In Binmore et al., eds (1993).

Carmichael, H. Lorne and W. Bentley MacLeod, "Territorial Bargaining," 1996. mimeo, Queen's University.

Damasio, Antonio R., *Descartes' Error*, New York, NY: Avon Books, 1995.

Deaton, Angus, *Understanding Consumption* Claredon Lectures in Economics, Oxford, UK: Clarendon Press, 1992.

Dye, Ronald A., "Costly Contract Contingencies," *International Economic Review*, February 1985, *26* (1), 233–250.

Ellingsen, Tore, "The Evolution of Bargaining Behavior," June 1995. mimeo Stockholm School of Economics.

Epstein, Larry G., "Behavior Under Risk: Recent Developements in Theory and Applications," in Jean-Jacques Laffont, ed., *Advances in Economic Theory: Sixth World Congress*, Vol. I, Cambridge, UK: Cambridge University Press, 1992, chapter 4.

Epstein, Richard A., *Simple Rules for a Complex World*, Cambridge, MA, USA: Harvard University Press, 1995.

Feigenbaum, Edward A. and Julian Feldman, eds, *Computors and Thought*, Menlo Park, CA: AAAI Press/The MIT Press, 1995 (orginally published 1963).

Frank, Robert H., *Passions within Reason*, New York, NY, U.S.A.: W. W. Norton & and Company, 1988.

Garey, Michael R. and David S. Johnson, *Computers and Intractability*, New York, USA: W. H. Freeman and Co., 1979.

Gibbons, Robert, "The Employment Relationship," 1995. mimeo, survey paper prepared for the Seventh World Congress of the Econometric Society, Tokyo, August 1995.

Gilboa, Itzhak and David Schmeidler, "Case-Based Decision Theory," *Quarterly Journal of Economics*, August 1995, *110* (3), 605–640.

Goldin, Claudia, "Monitoring Costs and Occupational Segregation by Sex: A Historical Analysis," *Journal of Labor Economics*, 1986, *4* (1), 1–27.

Grossman, Sanford J. and Oliver D. Hart, "The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration,," *Journal of Political Economy,*, August 1986, *94* (4), 691–719.

Hart, Oliver, *Firms, Contracts and Financial Structure*, Oxford, UK: Oxford University Press, 1995.

Holmstrom, Bengt, "Moral Hazard and Observability," *Bell Journal of Economics*, 1979, *10*, 74–91.

Holmström, Bengt, "Moral Hazard in Teams," *Bell Journal of Economics*, 1982, *13*, 324–40.

___ **and Paul Milgrom**, "Multi-Task Principle-Agent Analysis: Incentive Contracts, Asset Ownership, and Job Design," *Journal of Law, Economics, and Organization*, 1991, *7*, 24–52.

Johnson, Eric J., Colin Camerer, Sankar Sen, and Talia Rymon, "Limited Computation and Fairness in Sequential Bargaining Experiments," January 1996. Unpublished Manuscript, The Wharton School, University of Pennsylvania.

Johnson-Laird, Philip N., *Mental Models*, Cambridge, MA: Harvard University Press, 1983.

Kagel, John H. and Alvin E. Roth, eds, *The Handbook of Experimental Economics*, Princeton, NJ, U.S.A.: Princeton University Press, 1995.

Kahneman, Daniel, Jack Knetch, and Richard Thaler, "Experimental Tests of the Endowment Effect and the Coase Theorem," *Journal of Political Economy*, 1990, *98*, 1325–48.

—— **, Jack L. Knetsch, and Richard Thaler**, "Fairness as a Constraint on Profit Seeking: Entitlements in the Market," *American Economic Review*, September 1986, *76* (4), 728–741.

Kalai, Ehud and William Stanford, "Finite Rationality and Interpersonal Complexity in Repeated Games," *Econometrica*, March 1988, *56* (2), 397–410.

Kosslyn, Stephen M. and Olivier Koenig, *Wet Mind*, New York NY: Free Press, 1995.

Lazear, Edward P., *Personnel Economics*, Cambridge Mass.: MIT Press, 1995.

Lucas, Robert E., *Econometric Policy Evaluation: A Critique*, Vol. 1 of *Carnegie-Rochester Conference Series on Public Policy*, Amsterdam: North-Holland,

Machina, Mark J., "Choice Under Uncertainty: Problems Solved and Unsolved," *Journal of Economic Perspectives*, 1987, *1*, 121–54.

MacLeod, W. Bentley, "Complexity, Bounded Rationality and Heuristic Search," February 1995. Mimeo, Université de Montréal.

—— , "Contract, Complexity and the Employment Relationship," February 1996. mimeo Boston College.

—— **and Daniel Parent**, "Jobs and Incentives," 1996. Mimeo, C.R.D.E., Montreal.

—— **and James M. Malcomson**, "Implicit Contracts, Incentive Compatibility, and Involuntary Unemployment,," *Econometrica*, March 1989, *57* (2), 447–480.

—— **and** —— , "Motivation and Markets," January 1996. mimeo Boston College.

Maskin, Eric and Jean Tirole, "Unforeseen Contingencies, Property Rights, and Incomplete Contracts," Technical Report, Harvard University 1995.

Milgrom, Paul and John Roberts, *Economics, Organizations and Management*, Englewood Cliffs.: Prentice Hall, 1992.

Minsky, Marvin, "Steps Towards Artificial Intelligence." In Feigenbaum and Feldman, eds (1995 (orginally published 1963)) pp. 406–452.

Neuman, John Von and Oskar Morgenstern, *Theory of Games and Economic Behavior*, Princeton, NJ: Princeton University Press, 1944.

Newall, Allan and Herbert Simon, *Human Problem Solving*, Englewood Cliffs, NJ, USA: Prentice Hall, 1972.

Newall, Allen and Herbet A. Simon, "GPS, A Program that Simulates Human Thought." In Feigenbaum and Feldman, eds (1995 (orginally published 1963)) pp. 279–293.

Newell, Allan, *Unified Theories of Cognition*, Cambridge, Mass.: Harvard University Press, 1990.

Newell, Allen and Paul S. Rosenbloom, "Machanisms of Skill Acquisition and the Law of Practice," in John R. Anderson, ed., *John R. Anderson, ed.*, Lawrence Erlbaum Associates, Publishers Hillsdale, New Jersey 1981, pp. 1–55.

Nilsson, Nils J., *Principles of Artifical Intelligence*, Morgan Kaufmann Publishers, Inc., 1980.

Osborne, Martin J. and Ariel Rubinstein, *A Course in Game Theory*, Cambridge Mass.: The MIT Press, 1994.

Pearl, Judea, *Heuristics*, Reading, Mass.: Addison-Wesley Publishing Company, 1984.

Sargent, Thomas, *Bounded Rationality in Macroeconomics*, Oxford, U.K.: Oxford: Claredon Press, 1993.

Savage, Leonard J., *The Foundations of Statistics*, New York, N.Y.: Dover Publications, 1972.

Segal, Ilya, "Complexity and Renegotiation: A Foundation for Incomplete Contracts," 1995. Harvard University, Cambridge MA.

Simon, Herbert A., "A Formal Theory of the Employment Relationship," *Econometrica*, July 1951, *19*, 293–305.

—— , "A Behavioral Model of Rational Choice," *Quarterly Journal of Economics*, 1955, *69*, 99–118.

—— , *Models of Bounded Rationality, Volume 2*, Cambridge, Mass.: MIT Press, 1982.

Smith, Vernon L., "Economics in the Laboratory," *Journal of Economics Perspectives*, Winter 1994, *8* (1), 113–131.

Speigel, Matthew, Janet Currie, Hugo Sonnenschein, and Arunava Sen, "Understanding When Agents Are Fairmen or Gamesmen," *Games and Economic Behavior*, 1994, *7*, 104–115.

Tirole, Jean, "Incomplete Contracts: Where do we Stand?," 1994. Walras-Bowley Lecture, Summer meetings Econometric Society.

Townsend, Robert, "Optimal contracts and competitive markets with costly state verification," *Journal of Economic Theory*, 1979, *22*, 265–293.

Wright, Robert, *The Moral Animal*, New York, NY: Vintage Books, 1994.

Young, H. Peyton, "The Evolution of Conventions," *Econometrica*, 1993, *61*, 57–84.