

Uncertainty, Information Acquisition, and Economic Equilibria

by

Nir Vulkan

Submitted to the Department of Economics
in fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

UNIVERSITY COLLEGE LONDON

May 1996

© Nir Vulkan, 1996

The author hereby grants to UCL permission to reproduce and to distribute copies of this thesis document in whole or in part.

ProQuest Number: 10017770

All rights reserved

INFORMATION TO ALL USERS

The quality of this reproduction is dependent upon the quality of the copy submitted.

In the unlikely event that the author did not send a complete manuscript and there are missing pages, these will be noted. Also, if material had to be removed, a note will indicate the deletion.



ProQuest 10017770

Published by ProQuest LLC(2016). Copyright of the Dissertation is held by the Author.

All rights reserved.

This work is protected against unauthorized copying under Title 17, United States Code.
Microform Edition © ProQuest LLC.

ProQuest LLC
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106-1346

Abstract

Chapter one introduces the thesis, and the relationships between the different chapters. The second chapter presents a model of macroeconomic co-ordination failures based on local interaction where firms have some market power. The economy exhibits a multiplicity of Pareto-ranked equilibria. We show that the introduction of uncertainty about the competitive advantage of firms generates an endogenous equilibrium selection process. This suggests a possible solution for the multiplicity problem in macroeconomic models of co-ordination.

The third chapter introduces the notion of limited attention to economic modelling. Games are drawn from a large set, and players choose how much information about the game they wish to gather. In addition we assume that more information is expensive. Information can either be acquired secretly or publicly. We solve the model for both cases and identify the conditions under which the outcomes of the two models differ.

Chapter four is an application of the ideas studied in the third chapter. Here, we consider the problem of firms having to decide whether to enter a new market or not. Before they take that decision, they have the opportunity to buy information about several variables that might affect the profitability of this market. Our model differs from the existing literature on endogenous information acquisition in two respects: (1) there is uncertainty about more than one variable, and (2) information is acquired secretly. When the cost of acquiring information is small, entry decisions will be made as if there was perfect information. Equilibria where each firm acquires only a small amount of information are more robust than the socially undesirable equilibria where all firms gather all information. Examples illustrate the importance of assumptions (1) and (2).

The last chapter is somewhat different from the previous ones. In the 1950s and 1960s probability learning experiments showed that people adopt *probability matching strategies* in repeated choice situations, rather than strategies that maximise expected utility. These findings were compatible with the predictions of stochastic learning theories like Estes' and Bush and Mosteller's. As economists' interest in experiments and learning theories grows, these findings are of relevance today. We survey the known results and assess their relationship with different theories. We look at the many results gathered for the special case of a repeated, binary choice decision experiments. We show it is unlikely that probability matching can serve as the sole prediction of asymptotic behaviour, but that any theory of behaviour in repeated decision situations should include it in its predictions.

Acknowledgments

First, I would like to thank my supervisors, Ken Binmore and Tilman Börgers, who contributed greatly in my transformation to the "proper economist" that I am now. I consider myself lucky to have had them as my teachers, and I am grateful for their time and patience. I thank Tilman for showing me how to be self-critical and how to separate grandiose statements from real work. Ken, for always being there to give me support in times when all seemed lost, and for being an endless source of new ideas. Ken also provided (through his new ELSE centre) financial support in my final year in UCL. I thank him and the Centre for Economic Learning and Social Evolution for that support. Financial support from the department, through teaching assistantships, is gratefully acknowledged.

During my years in UCL I was influenced by the numerous conversations I had with other faculty members and graduate students in the department. I would like to thank three of them who had a major impact on me and on this thesis: Pasquale Scaramozzino who co-authored the paper chapter 2 is based on; Sjaak Hurkens who co-authored an early version of chapter 4, and my good friend and colleague Isaac Alfon, who made a special effort in the last three years to show me what *real* economists do in the real world.

A special thank-you to Fiona for the last three years, and to my parents, Eti and Uzi Vulkan, for always encouraging me to believe in myself and for their support throughout my long years of study.

Contents

1	Introduction	6
1.1	Games and Information	6
1.2	Global-games and Macroeconomic Equilibria	8
1.3	Meta-games and Limited Attention	10
1.4	Information and Equilibria in Experiments	13
2	Uncertainty and Endogenous Selection of Macroeconomic Equilibria	18
2.1	Introduction	18
2.2	The model	22
2.3	Multiple equilibria and dominant strategies	24
2.4	Uncertainty and equilibrium selection	28
2.4.1	Locally correlated signals	32
2.5	Conclusions	35
3	Modelling Limited Attention	41
3.1	Introduction	41
3.2	The Model	44
3.3	Results for the Simultaneous Move Game	46
3.4	Results for the Two-Stage Game	51
3.5	Discussion	57
4	Information Acquisition and Entry	61
4.1	Introduction	61
4.2	The Model	65

4.3	Information Structures	68
4.4	Two Firms	70
4.5	Observable Information Acquisition	76
4.6	Conclusions	79
5	An Economist's Perspective on Probability Matching	82
5.1	Introduction	82
5.2	Theoretical Background	84
5.3	Experimental Results	88
5.4	Discussion	96
5.5	Other Matching Behaviours	99
5.6	Conclusions	101

Chapter 1

Introduction

1.1 Games and Information

Ever since the publication of Von Neumann and Morgenstern's book in 1944, game-theory has been successfully applied in economics. It provided a unified language for situations in which strategic thinking and decision-making are required. It helps in our understanding of the relationships between models that for many years seemed contradictory (Bertrand (1883), Cournot (1838), and Von Stackelberg (1934)) by providing a convenient taxonomy for economic situations based on their strategic structure.

Successful as it may be, game theory is still at best a metaphor for the real world. As such it omits many of the details that would otherwise come into consideration. It is therefore worthwhile to test whether the theoretical results are robust to small changes in the model. When a small change to the assumptions results in big changes for the solutions, then those particular assumption must be closely checked. It could prove to be, in the right context, an important feature of the economic reality. An example of such a finding is the importance of timing in models of competition. The outcomes of a two-firm quantity competition are very different when we let both firms act simultaneously (as in Cournot's analysis), from the outcomes when one of the firms can move first (as presented by Von Stackelberg). This remains true even when the length of time between the moves becomes very small. These days the concept of a *first-mover* is common in economic theory.

The robustness of the game-theoretical model can be questioned with respect to the assumptions being made regarding the information and beliefs possessed by the players. In reality, economic agents are likely to face uncertainties. It is therefore worthwhile to investigate whether complete information results are robust to small amounts of incomplete information. These can be incorporated with the model by introducing a move by Nature, as first proposed by Harsanyi (1966/67). Harsanyi was also the first, in his 1973 paper to investigate robustness of equilibria. In his model of games with randomly distributed payoffs he shows that if players are slightly uncertain about the payoffs of their opponents, then all equilibria (including the mixed equilibria) are robust. More recently Carlsson and Van Damme (1993) investigated their, somewhat similar, model of *global games* where players are only partially informed about the payoffs in the form of a noisy signal. Those signals have some small support, and are correlated amongst players. When the noise tends to zero, global games converge to games of perfect information. However, they show that this is not necessarily the case for the equilibria of the game. If the underlying game exhibits a multiplicity of equilibria, then some equilibria of the game with perfect information will not be selected even if the noise goes to zero. Players will avoid playing the "risky" ones, and will instead co-ordinate on the *risk-dominant* equilibrium. (the concept of risk-dominance was introduced by Harsanyi and Selten (1988).) The reason that Carlsson and Van Damme find that only some, and not all, equilibria are robust, as in Harsanyi (1973), is that in their model noise is correlated amongst players. A player can learn from his own signal something about the payoffs of his opponents. A very general approach to the robustness of equilibria to incomplete information has recently been developed by Kaji and Morris (1995).

This thesis contributes to the understanding of the robustness of the game-theoretical model, by extensively investigating the idea that players are not able to distinguish between similar games. This is done by adopting Carlsson and Van Damme's approach to a macroeconomic example (chapter 2), and by investigating the somewhat similar idea of *meta-games* (chapters 3 and 4). We also look (in chapter 5) at the relationship between

the outcomes of experiments and the information provided to the subjects.

1.2 Global-games and Macroeconomic Equilibria

In chapter 2, we present a simple model of competition within a local interaction environment. We assume that firms strategically interact with their immediate neighbours while neglecting the effect of their actions on the economy as a whole. The demand for each of the firms' products is equal to the total output divided by the number of firms. The economy therefore shows a demand externality. The underlying game is supermodular and so exhibits strategic complementarities. (see Cooper & John (1988).) Supermodular games had proved to be a useful tool in our understanding of the micro foundations of macro models by offering an environment where agents do have some market power, while neglecting the effects of their own behaviour on the aggregates. (See Silvester 1993.) A local-interaction supermodular game is particularly attractive as it also allows all sorts of strategic effects within neighbourhoods. If the game is supermodular and the payoffs for each player are monotonic in the strategies of the other players, then if the game exhibits a multiplicity of symmetric equilibria (as in our model), these can be Pareto-ranked. Players, therefore, face a co-ordination problem. The idea that the economy may exhibit a multiplicity of equilibria, and that agents may fail to co-ordinate on a high level activity, goes back at least to Kaldor (1940).

With very few exceptions the macro literature restricts itself to concluding that models with strong micro foundations will typically exhibit a multiplicity of equilibria. In many cases, it is taken for granted that this is as far as we can get by using economic theory. Sometimes, the suggestion is made that the role of the social planner is to create a co-ordination device that enables players to get to the better equilibrium. However a theory that predicts more than one outcome is problematic both from a theoretical and from an empirical point of view. (See Harsanyi and Selten 1988, p.12 and Jovanovic 1989.) For these reasons we believe that the analysis should be carried further. Every

theoretical model is only as good as it is robust to some realistic perturbations, and macro models are no exception. It is reasonable to assume that firms will observe features of their economic environment with some idiosyncratic noise. In our model, firms are uncertain about how much extra demand they will get by being more competitive than their neighbours. Using a version of Carlsson and van Damme's approach we show that the number of equilibria in the new, noisy model, is reduced to one. This suggests that indeterminacy might not be such an important feature of macro models, and that we should perhaps pay more attention to the risk-dominance criterion. This is not all good news: on the one hand we show that co-ordination could spontaneously arise in the economy; on the other, that such co-ordination will be determined by considerations of risk rather than efficiency. If the risk-dominant equilibrium is not Pareto-efficient, then the model suggests that a social planner must do more than provide players with a co-ordination device: all uncertainties must first be completely eliminated in order for the Pareto-efficient outcome to be rationalisable.

We first analyse a deterministic model in which all firms are able accurately to calculate the payoffs given the value of a competitiveness parameter, θ . We show that, for some values of θ , all firms have a dominant strategy and that for others, there is a multiplicity of symmetric equilibria. Once noise is introduced, it is shown that there will be an endogenous selection process in the form of iterated elimination of strictly dominated strategies. Firms will, almost always, co-ordinate on the risk-dominant equilibrium, even if there exists another equilibrium at which everybody is better off. We then show that in order to generate this process in our model (or in any model of local interaction), it is sufficient that signals are locally correlated. That is, each firm faces some (small) amount of uncertainty about the competitive advantage in its neighbourhood, and expects its neighbours to be equally uncertain. However, they do not have to have specific expectations regarding the beliefs of each and every firm out of their immediate vicinity. This suggests an interesting way of extending Carlsson and van Damme's results for games with many players. Without local interaction, extending the method to an

n -player game requires that players have specific beliefs about the information obtained by *each* of the other players, which seems to be a very strong requirement.

1.3 Meta-games and Limited Attention

Consider again a situation where the actual game to be played is determined by a move of Nature. Suppose that, *a priori*, any of the states in a given set can be realized. (given some probability measure over that set.) Suppose further, that players can obtain *any* amount of information, but that there is a cost function associated with the gathering of information. Information is obtained in the form of a partition of the possible states of the world. A more refined partition will provide more accurate information, but will also be more expensive. Players will therefore face a trade-off. We use the term *meta-game* to refer to the game in which players choose both a partition, and a conditional strategy. Meta-games differ from Harssanyi's model and from Carlsson and van Damme's global-games in that uncertainty is endogenous instead of being exogenously fixed. That is, players can manipulate the amount of uncertainty they face.

If we fix an information-gathering cost function (*i.e.* a non-negative function defined over the set of all possible partitions), and if we let the values returned by this cost function go to zero then, in the limit, our model converges to a game of perfect information. It is therefore interesting to see whether the equilibria will also converge or, as in Carlsson and van Damme's case, will yield different behaviour in the limit. The answer is twofold, depending on the way information is gathered. It could be gathered privately or publicly. In the first case the meta-game is a one-stage game, where agents can only condition their strategy choice on their own information. In the second case, they can also condition this strategy choice on the information partitions chosen by their opponents. If the game is modelled as having only one stage, then, when costs tends to zero, players will play *as though* they have perfect information. Hence the results converge to those of the limiting game. If, however, we choose to analyse a two-stage game, we find

that behaviour in the limit is very different from that in the case of perfect information.

We interpret a game of endogenous information-acquisition as a model of limited attention. The model presented in chapter 3 therefore belongs to the literature on bounded rationality. Note that in economics the term bounded rationality is used in somewhat different ways. (see Auman (1986) for a survey and a classification.) Here we define bounded rationality as optimal behaviour that takes account of cognitive restrictions placed on the information-processing capacity of the individual. (as in Abreu and Rubinstein (1989), Dow (1993), Rubinstein (1986), or Rubinstein (1993) to name only a few.) Our model fits into that category because attention is a scarce cognitive resource, but one that can nevertheless be manipulated. If two situations are very similar and if it is cognitively "expensive" to distinguish between them, then we expect that agents will not always do so.

We analyse the situation where the actual game to be played is drawn from a continuous set of games. We compare the outcomes of the one and two-stage games, and find conditions under which the two will differ. We also show that those conditions are very sensitive to the players' ability to utilize their information of the other players' first-stage strategies. In other words, to sustain those outcomes that do not converge to the limiting game, players need to be very attentive to what their opponents are doing. This suggests that a two-stage model is only appropriate when such a level of attentiveness can be reasonably achieved.

We believe that the general framework presented in chapter 3 can be useful for economic modelling in places other than the theory of bounded rationality. Firms, in many cases, can manipulate the amount of information they obtain before making important decisions. More information will make the firm better informed about some features of the economic environment, but will also be costly. Examples include decisions on whether to hire the services of a marketing consultant, or whether to establish a research department. In fact, models of endogenous information acquisition seem to be a very natural way to approach such problems. It is therefore not surprising that there exists

some literature on the subject. (examples include Chang and Lee (1992), Fershtman and Kalai (1993), Hwang (1993, 1995), Li *et al.* (1987), Matthews (1984), Milgrom (1981), Ponsard (1979), and Vives (1988).) In particular, we consider the decisions of firms as to whether or not to enter a new market. Firms face uncertainties about several parameters of the economy. For example, they face uncertainty about their future demand and their overall costs of production. Firms are able to reduce these uncertainties by hiring marketing services at a cost. The model presented in chapter 4 differs from the existing literature in two important ways: First, we consider a one-stage game. (all the other models we know of, with the exception of Matthews (1984) consider a two-stage game.) The reasons for this modelling decision should be clear from the results of chapter 3. In a two-stage model, firms are able to condition their own behaviour on the choice of information partition made by their opponents. While it is reasonable to assume that firms are able to find out whether other firms are involved in market research, it is unreasonable to assume that they know exactly *what kind* of information they are receiving.

A second important difference between our work and the existing literature is that we model firms as facing uncertainty about more than one variable. (the exception here is Fershtman and Kalai (1993) who study the behaviour of a single firm in several markets, where demand for each of them is uncertain.) It is usually assumed that multi-dimensional uncertainty can be reduced to a single dimension without loss of generality. We show that this is not the case. In the one-dimensional case, if information becomes cheap then at least one firm will choose to obtain it, in equilibrium. If there is uncertainty about two or more variables, then, even if the costs of acquiring information about a certain variable is equal to zero, this still does not imply that there will be a firm who will choose to learn about it.

The models presented in chapters 3 and 4 are closely related. Chapter 4 is an application of the methods studied in chapter 3. In chapter 4 we address a less general problem which enables us to find a more detailed solution. It will be necessary to adapt some of the results from section 3.3 (Propositions 1 and 2) in section 4.2. This is because the

structure of the underlying set of games is very different - a one dimensional continuous set in chapter 3 versus a multi-dimensional discrete set in 4. The logic of the proofs is quite similar, though. Other than that, the two models belong to different fields of economic theory. Chapter 3 relates to the literature of bounded rationality, while chapter 4 presents an industrial organization model. The two chapters also address different aspects of the abstract mathematical problem of information gathering in games.

1.4 Information and Equilibria in Experiments

The terms "uncertainty" and "information" are interpreted differently in chapter 5 than in previous chapters. In both global games and meta-games we assumed that players know the rules governing the initial move by Nature. That is, the support and distribution of the set of possible games are assumed to be common knowledge. In chapter 5 we look at cases where players know their *ex post* payoff, but can only try and learn (over time) the rules with which Nature chooses them. Here, again, it is shown that players' behaviour is sensitive to the exact structure of information available to them.

Economists study the equilibrium of a given situation because they assume that it will be played if agents have enough time to learn how to behave in such a situation. (see Binmore (1987).) The process by means of which equilibrium is achieved is assumed to be some sort of a trial-and-error process, where a positive feedback will encourage agents to repeat strategies that yield large payoffs. Experiments from the 1950s and 1960s suggested this is not necessarily what happens in repeated choice situations. Psychologists often used these findings to criticize the underlying rationality assumptions of neo-classical economic theories. Mainstream economics dealt with the problem mostly by ignoring it. From a theoretical point of view, stochastic learning theories (like Estes (1957) and Bush and Mosteller (1955)) looked directly at the process of trial-and-error learning, and found that it would not always converge to the strategy that maximizes expected payoff. However, those theories assumed very little regarding players' information

and expectations.

Today, as the interest of economists is shifted towards the actual learning processes (and population learning, like the replicator dynamics), we can learn something from those early findings. We have now more complex theories of stochastic learning that refer to the expectations of players at each stage of the process (Börgers and Sarin (1995)). They found that whether agents will end up playing the strategy that maximizes their *ex ante* expected payoffs, will depend on the relationship between payoffs and expectations (the feedback). If feedbacks are non-negative, then the learning process will be similar to that of a population replicator dynamics. Agents will therefore end up playing the equilibrium. However, if it is possible for agents to receive negative feedbacks, then behaviour in the limit could be very different. Their theory therefore provides us with clearer ideas of when and why equilibrium is achieved. In chapter 5 we look at survey and review those early experiments from the 1950s and 1960s, paying special attention to the information provided to the subjects.

References

- ABREU, D. AND A. RUBINSTEIN (1988), "The Structure of Nash Equilibrium in Repeated Games with Finite Automata", *Econometrica*, pp. 1259-1281.
- AUMANN, R. J. (1986), "Rationality and Bounded Rationality", *The 1996 Nancy L. Schwartz Lecture*, Kellogg School of Management, Northwestern University.
- BERTRAND, J. (1883), "Theorie Mathematique de la Richesse Social", *Journal des Savants*, pp. 499-508.
- BINMORE, K. (1987), "Modelling Rational Players, I," *Economics and Philosophy*, **3**, pp. 179-214.
- BÖRGERS, T. AND R. SARIN (1995), "Learning through Reinforcement and Replicator Dynamics", University College London Discussion Paper, **19**.

- BUSH, R. AND F. MOSTELLER (1955), *Stochastic models for learning*. New York: Wiley.
- CARLSSON, H. AND E. VAN DAMME (1993), "Global Games and Equilibrium Selection", *Econometrica* **5**, pp. 989-1018.
- CHANG, C.H. AND C.W. LEE (1992) "Information Acquisition as Business Strategy", *Southern Economic Journal*, **58**, pp. 750-761.
- COOPER, R. AND A. JOHN (1988), "Coordinating coordination failure in Keynesian models", *The Quarterly Journal of Economics* **3**, pp. 441-463.
- COURNOT, A. (1838), *Recherches sur les principes mathematiques de la theorie des richesses*. Translated into English by N. Bacon as *Researches in the mathematical principles of the theory of wealth*. London: Haffner, 1960.
- DOW, J. (1993), "Search Decisions with Limited Memory", *Review of Economic Studies*, **58**, pp. 1-14.
- ESTES, W. (1957), "Theory of Learning with Constant, Variable, or Contingent Probabilities of Reinforcement", *Psychometrika*, **22**.
- FERSHTMAN, C. AND E. KALAI (1993), "Complexity Considerations and Market Behaviour", *RAND Journal of Economics*, **24**, pp. 224-235.
- HARSANYI, J. (1967-68), "Games with incomplete information played by Bayesian players, I, II, and III", *Management Science*, **14**, pp. 159-182, 320-324, 486-502.
- HARSANYI, J. (1973), "Games with Randomly Disturbed Payoffs: A New Rational for Mixed Strategy Equilibrium Pointss", *International Journal of Game Theory* **2**, pp. 1-23.
- HARSANYI, J. AND R. SELTEN (1988), *A general theory of equilibrium selection in games*. Cambridge, Mass.: MIT Press.

- HWANG, H. (1993). "Optimal Information Acquisition for Heterogeneous Duopoly Firms", *Journal of Economic Theory*, **59**, pp. 385-402.
- HWANG, H. (1995) "Information Acquisition and Relative Efficiency of Competitive, Oligopoly, and Monopoly Markets," *International Economic Review*, **36**, pp. 325-340.
- JOVONOVIC, B. (1989), "Observable Implications of Models with Multiple Equilibria", *Econometrica*, **57**, pp. 1431-1437.
- KAJI, A. AND S. MORRIS (1995), "The Robustness of Equilibria to Incomplete Information", University of Pennsylvania, mimeographed.
- KALDOR, N. (1940), "A Model of the Trade Cycle", *The Economic Journal*, **197**, pp. 78-92.
- LI, L., MCKELVEY, R.D., AND T. PAGE (1987), "Optimal Research for Cournot Oligopolists", *Journal of Economic Theory*, **42**, pp. 140-166.
- MATTHEWS, S. (1984), "Information Acquisition in Discriminatory Auctions", in: *Bayesian Models in Economic Theory*, eds. M. Boyer and R.E. Kihlstrom. Elsevier.
- MILGROM, R. (1981), "Rational Expectations, Information Acquisition and Competitive Bidding", *Econometrica*, **49**, pp. 921-944.
- VON NEUMANN, J. AND O. MORGENSTERN (1944), *Theory of games and economic behaviour*, Princeton University Press, Princeton NJ.
- PONSSARD, J. P. (1979), "The Strategic Role of Information on the Demand Function in an Oligopolistic Market", *Management Science*, **25**, pp. 243-250.
- RUBINSTEIN, A. (1986), "Finite Automata Play the Repeated Prisoners' Dilemma", *Journal of Economic Theory*, **39**, pp. 83-96.

- RUBINSTEIN, A. (1993), "On Price Recognition and Computation Complexity in a Monopolistic Model", *Journal of Political Economy*, **31**, pp. 473-484.
- SILVESTRE, J. (1993), "The market-power foundations of macroeconomic policy", *Journal of Economic Literature*, **31**, pp. 105-141.
- VON STACKELBERG, H. (1934), *Marktform und Gleichgewicht*, Vienna: Julius Springer.
- VIVES, X. (1988), "Aggregation of Information in Large Cournot Markets", *Econometrica*, **56**, pp. 851-876.

Chapter 2

Uncertainty and Endogenous Selection of Macroeconomic Equilibria

2.1 Introduction

One of the most actively pursued areas of research in modern macroeconomics is the attempt to provide rigorous microeconomic foundations for macroeconomic relationships. It is increasingly being recognised that a satisfactory explanation of aggregate relationships must rely upon an analysis of the interactions amongst individual agents (see for instance the contributions in Mankiw and Romer, 1991, and in Dixon and Rankin, 1995). Most research has focused on the investigation of the properties of non-Walrasian models, in which agents wield market power and are therefore able to influence the market in which they operate. The resulting market failures can be responsible for multiple Nash equilibria, whereby the economy could settle at more than one level of aggregate activity.


The idea that the economy may exhibit a multiplicity of equilibria, and that agents may fail to co-ordinate their actions on a high level of activity, goes back at least to Kaldor (1940). In his original analysis, this possibility arises due to non-linearities in the *ex ante* savings and investment functions. Recent accounts of co-ordination problems have emphasised the role of incomplete markets, increasing returns, and search costs as

possible sources for the lack of co-ordination (see Silvestre, 1993, for a survey). Often, the underlying model of an economy with multiple equilibria can be described as a supermodular game (there is a strategic complementarity among the agents' payoffs), and its Nash equilibria can be Pareto-ranked (Cooper and John, 1988; Milgrom and Roberts, 1990).

A fundamental problem with macroeconomic models with multiple equilibria is that the market outcome is left undetermined. This has serious implications for the possible interpretation of the model. Harsanyi and Selten argue that "a theory that can only predict that the outcome of a noncooperative game is an equilibrium point - without specifying which equilibrium point it is - is an extremely weak and uninformative theory" (Harsanyi and Selten, 1988, p. 12). Models with multiple equilibria can also be very difficult to implement empirically: Jovanovic (1989) shows that the structure of these models can only be identified under restrictive conditions.

The usual approach to reduce the number of feasible equilibria is to impose additional restrictions on the behaviour of the agents, in order to rule out some of the Nash equilibria of the economy (as in most of the refinements concepts). In the present paper, we follow a different route. We suggest that an equilibrium selection process in macro models can be made possible based on the robustness of equilibria to the introduction of small amount of incomplete information into the model. We exploit the notion that, in a decentralised market economy, agents mainly interact with their neighbours. If even a very small amount of noise is introduced in the economy, and if agents receive stochastic signals which are correlated among neighbours, this could lead agents to endogenously coordinate their decisions. This produces the effect of reducing the set of feasible equilibria and, under some conditions, can lead to a unique equilibrium for the economy.

The main reason for the selection of equilibria is closely related to the notion of risk dominance, as analysed by Harsanyi and Selten (1988). The intuition for our results can be understood by an example. Suppose that we start from a situation in which the fundamentals are consistent with only one strategy (*i.e.* a dominant strategy). Let



the fundamentals of the economy be slightly perturbed, in such a way that additional strategies are now also possible in equilibrium. Under such conditions, it is clear that the first strategy is "*less risky*" than the latter ones. Agents will look at their own signal and will make inferences about the possible actions of the other players. The less risky strategy would also become the *dominant* strategy following the perturbation.

A key point of the analysis is that when strategic interactions only take place amongst neighbours, locally correlated idiosyncratic signals are sufficient to generate equilibrium selection. This is a major improvement over the regular requirements for co-ordination in models without local interactions, in which the signals must be correlated amongst *all* players.

It should be noted that no such endogenous co-ordination mechanism exists in a completely deterministic framework. The failure to co-ordinate could therefore be a rather fragile feature of some macroeconomic models. It is also important to note that the rationality requirements behind iterated elimination of strictly dominated strategies, which lead to equilibrium selection, are less stringent than those for Nash equilibrium. Hence, although the predicted outcomes are a sub-set of the Nash equilibria, the requirements are no less plausible than for the conventional Cournot-Nash equilibrium.

The results in this paper are related to the literature on equilibrium selection in game theory. In the latter, it is suggested that an endogenous selection of equilibrium may follow from rational behaviour, in the presence of uncertainty about some features of the game. A fundamental contribution is the paper by Carlsson and van Damme (1993a), which considers the equilibrium selection process in a very general (2×2) game in the presence of uncertainty about payoffs. Other contributions are Shin (1995), who considers a search model with idiosyncratic noise, and Morris (1995), who analyses a work-shirk model with uncertainty about timing. Carlsson and van Damme (1993b) analyse an n -player stag-hunt game with payoff uncertainty.

Relative to the existing game-theoretic literature, the present paper has the following innovative features. First, the equilibrium selection process takes place in a macro model.

Second, the results by Carlsson and van Damme (1993a) are generalised to an economy formed of n agents with a continuum of strategies (this is however achieved at the cost of a sacrifice in the generality of the results, which now hold for the particular game we consider). Third, we show that the Carlsson-van Damme's findings can be applied to models with strategic complementarities. Forth, we show that signals need only be correlated among neighbours (and not among all agents in the economy). This makes it more appropriate to apply an endogenous selection process to a large economy.

We develop our arguments by using a model of local oligopoly based on Salop's (1979) circular economy. The model exhibits strategic complementarities, and has a multiplicity of Pareto-ranked Nash equilibria. Agents only interact locally, and observe stochastic signals which are correlated across neighbours.

Formally, each firm is assumed to be exogenously located on a circle (which represents varieties of horizontally differentiated commodities), and to be exposed to competition only *vis a vis* its immediate neighbours (its potential market rivals). For simplicity, we abstract from price decisions. We instead assume that the level of activity of a firm directly affects the intensity of competition with its rivals. More precisely, a higher level of production is also associated with a higher level of expenditures on advertising. This captures a "business stealing" effect (see *e.g.* Mankiw and Whinston, 1986), whereby a firm producing at a higher level of activity can attract customers away from its rivals.

The firms' investment and production decisions affect the competitive conditions in their local neighbourhood, but do not influence demand farther away in the preference space. However, the demand for each firms' output depends on the investment and production decisions of *all* firms in the economy. The payoff of each firm is thus a function both of the strategies chosen by its neighbours (*via* oligopolistic competition), and of the behaviour of all other firms in the economy (which affects the total demand for the firm's product). There is thus an aggregate demand externality.

Local markets exhibit idiosyncratic features, which might result in a different strategic advantage from a more aggressive behaviour over the various locations on the preference

space. Firms could also differ with respect to their information. In general, a firm's information set includes both common knowledge and private information. The common knowledge comprises past history, and can be thought of as reflecting the public perception about the effectiveness of choosing a more aggressive investment strategy. The private information is related to the way firms observe their local environment, and could be affected by different marketing technologies and preliminary consumer surveys.

The balance between common knowledge and private information can affect individual and aggregate behaviour. This paper shows that even a small amount of idiosyncratic uncertainty could lead to endogenous co-ordination amongst agents.

The structure of the paper is as follows. The next section presents the structure of the model and motivates our analysis. Section 3 describes the properties of the economy in the absence of idiosyncratic noise. Section 4 analyses the more general model and discusses the equilibrium selection process. Section 5 concludes.

2.2 The model

A continuum of consumers is uniformly located over a circle, whose measure is normalised to unity. Their position on the circle represents their preferences. There is a large number n of firms in the economy, located at uniformly spaced points on the circle. Firms do not choose their location, *i.e.* their variety, but decide on their level of output. A higher level of production is associated with higher expenditure on advertising. Each consumer supplies an equal share of labour to every firm in the economy, and receives an equal share of the total wage payments of every firm. Workers inelastically supply one unit of labour. Firms are uniformly owned by consumers, who act as their shareholders. Total demand in the economy is given by the sum of labour income and of firms' expenditures on advertising.

Total costs are $c = w + a$, where $\underline{c} = w$ and $\bar{c} = w + \bar{a}$. Output increases with costs, *i.e.* $y = f(c)$, $f'(\cdot) > 0$. We denote $\underline{y} = f(\underline{c})$ and $\bar{y} = f(\bar{c})$. We assume that if firm i

invests more than its neighbour $i+1$ then it attracts a larger share of the demand over the arc $(i, i+1)$. Formally, we let

$$(1) \quad \theta_i(|a_i - a_{i+1}|) = \frac{1}{2} + \frac{\bar{\theta} - 1/2}{\bar{c} - \underline{c}} |a_i - a_{i+1}|$$

where $\bar{\theta} \in [1/2, 1]$. If $a_i \geq a_{i+1}$ then firm i receives a fraction θ_i of the demand in the arc, otherwise it receives $1 - \theta_i$. If $a_i = \bar{a}$ and $a_{i+1} = 0$ then $\theta_i = \bar{\theta}$ denotes the maximum *ex ante* demand on the arc. Before investment decisions are being made, firms observe $\bar{\theta}$. However, this observation can be noisy. Formally, we assume that each firm observes $\bar{\theta}_i = \bar{\theta} + v_i$ where v_i is uniformly distributed over $[-v, v]$. One of the main purposes of this paper is to show that even when v is arbitrarily small the structure of equilibria in the economy is drastically changed. If $v = 0$ then all firms have the same expectations about $\bar{\theta}$. In general, these expectations could be different. *Ex post*, $\bar{\theta}$ could be different on different arcs.

In game-theoretic terms, we have a simultaneous-move game with a continuum of strategies. The payoff of firm i is a function of the strategy profile of the other $n-1$ firms. However, since firms do not interact directly with firms other than their immediate neighbours, a sufficient statistic for the behaviour of the remaining $n-3$ firms is their average behaviour, which can be approximated by total output in the economy divided by the measure of the set of firms. Let $Y \equiv \sum_{j=1}^n y_j$ denote total output. Then the expected payoff for firm i is:

$$(2) \quad \pi_i = \frac{Y}{n} [\theta_i \cdot I(a_i \geq a_{i+1}) + (1 - \theta_i) \cdot I(a_i < a_{i+1}) \\ + (1 - \theta_{i-1}) \cdot I(a_{i-1} \geq a_i) + \theta_{i-1} \cdot I(a_{i-1} < a_i)] - c_i$$

where $I(\cdot)$ denotes the indicator function. Firms strategically interact at the local level, but neglect the economy-wide effects of their actions. The simultaneous move nature of the game, together with the aggregate demand externalities, will generate a

co-ordination game.

2.3 Multiple equilibria and dominant strategies

In this section we first analyse the model when the idiosyncratic component of the noise is absent, *i.e.* $v = 0$. We identify values of the competitiveness parameter $\bar{\theta}$ for which firms have a dominant strategy. These values play a crucial role in the analysis of section 4. We calculate the symmetric equilibria of the game and show that it can have multiple Pareto-ranked equilibria. We also show that, when firms' choice is restricted to two activity levels only, the equilibrium conditions in the market do not depend on the exact configuration of firms along the circle, *i.e.* location does not matter.

Let us first consider whether there are regions of values where \underline{c} is a dominant strategy. The best candidate behavioural profile for firm i to have an incentive to switch from \underline{c} to a higher level of investment occurs when: (1) both neighbours are investing \underline{c} ; and (2) all other firms invest \bar{c} , hence gaining an additional fraction of demand is most valuable. Firm i 's profit is $\pi_i = \frac{1}{n}[(n-3) \cdot \bar{y} + 3 \cdot \underline{y}] - \underline{c}$. Switching to $c_i = \underline{c} + \epsilon$, $\epsilon \in (0, \bar{c} - \underline{c}]$ will result in profit $\pi_i^* = 2\theta(\epsilon) \frac{1}{n}[(n-3) \cdot \bar{y} + 2 \cdot \underline{y} + f(c_i)] - (\underline{c} + \epsilon)$. The switch is not profitable iff $\pi_i^* < \pi_i$, which is equivalent to

$$(3) \quad \bar{\theta} < \frac{1}{2} \left\{ 1 + \frac{\bar{c} - \underline{c}}{\frac{1}{n}[(n-3)\bar{y} + 2\underline{y} + f(c_i)]} \right\} - \frac{\bar{c} - \underline{c}}{2\epsilon \cdot [(n-3)\bar{y} + 2\underline{y} + f(c_i)]} \cdot [f(c_i) - \underline{y}]$$

If we let $n \rightarrow \infty$ that is, if the three firms investing \underline{c} have a negligible impact on total demand, then the inequality converges to the following:

$$(3.1) \quad \bar{\theta} < \frac{1}{2} \left(1 + \frac{\bar{c} - \underline{c}}{\bar{y}} \right) \equiv \theta^1$$

Note that the right hand side of (3.1) is smaller than that of (3), hence requiring $\bar{\theta} < \theta^1$ guarantees that \underline{c} is a dominant strategy for every possible value of n . Symmetri-

cally, the condition for \bar{c} to be a dominant strategy (when both neighbours invest \bar{c} and all other firms invest \underline{c}) is

$$(4) \quad \bar{\theta} > \frac{1}{2} \left(1 + \frac{\bar{c} - \underline{c}}{\underline{y}} \right) \equiv \theta^2$$

From (3) and (4), $\theta^1 < \theta^2$ since $\bar{y} > \underline{y}$. Thus, for small values of $\bar{\theta}$ it is dominant to invest \underline{c} and for large values it is dominant to invest \bar{c} ¹. For intermediate values we will now show that there is a multiplicity of equilibria.

If all firms invest \bar{c} then the payoff for firm i is $\bar{y} - \bar{c}$. For this to be an equilibrium no firm should have an incentive to deviate by investing a lower amount. This requires that $\bar{y} - \bar{c} > 2(1 - \theta(\epsilon))\bar{y} - (\bar{c} - \epsilon)$, which holds when $\bar{\theta} > \theta^1$. Similarly, when all firms invest \underline{c} the payoff for firm i is $\underline{y} - \underline{c}$. This is an equilibrium when $\underline{y} - \underline{c} > 2\theta(\epsilon)\underline{y} - (\underline{c} + \epsilon)$, which holds when $\bar{\theta} < \theta^2$.

Assume now that all firms invest $\tilde{c} = \alpha\underline{c} + (1 - \alpha)\bar{c}$. Firm i has no incentive to change its level of investment when both the following inequalities hold: (1) $Y/n - \tilde{c} \geq 2\theta(\epsilon)Y/n - (\tilde{c} + \epsilon)$, and (2) $Y/n - \tilde{c} \geq 2(1 - \theta(\epsilon))Y/n - (\tilde{c} - \epsilon)$. These conditions are jointly met only when

$$(5) \quad \bar{\theta} = \frac{1}{2} + \frac{\bar{c} - \underline{c}}{2Y/n}$$

To summarise, there is a unique symmetric Nash equilibrium when $\bar{\theta} < \theta^1$ and when $\bar{\theta} > \theta^2$. For $\theta \in (\theta^1, \theta^2)$ there are three symmetric Nash equilibria where all firms invest \underline{c} , \bar{c} , or \tilde{c} (which depends on $\bar{\theta}$, as one can see by solving equation (5) for α).

Next, consider what happens when the economy starts from a situation in which each firm invests \bar{c} , and αn contiguous firms invest \underline{c} instead. The payoff for \bar{c} -firms in the interior region is $[\alpha\underline{y} + (1 - \alpha)\bar{y}] - \bar{c}$, and $[\alpha\underline{y} + (1 - \alpha)\bar{y}] - \underline{c}$ for interior \underline{c} -firms. Firms are

¹When the competitive advantage is a general function of the expenditure on advertising, similar conditions hold, although the actual values are different.

better off if they operate in an environment in which their neighbours are not aggressive. For firms at the edges, the payoff is $(1/2 + \bar{\theta}) \cdot [\alpha \underline{y} + (1 - \alpha) \bar{y}] - \bar{c}$ if they invest \bar{c} and $(1/2 + (1 - \bar{\theta})) \cdot [\alpha \underline{y} + (1 - \alpha) \bar{y}] - \underline{c}$ if they invest \underline{c} . In order for the configuration to be an equilibrium, firms at the edges must be indifferent between the two extreme strategies. This implies

$$(6) \quad \bar{\theta} = \frac{1}{2} \left[1 + \frac{\bar{c} - \underline{c}}{\alpha \underline{y} + (1 - \alpha) \bar{y}} \right] \in (\theta^1, \theta^2)$$

or

$$(7) \quad \alpha = \frac{1}{\bar{y} - \underline{y}} \left(\bar{y} - \frac{\bar{c} - \underline{c}}{2\bar{\theta} - 1} \right) \in (0, 1)$$

If firms are restricted to only two levels of investment; \bar{c} and \underline{c} , we are able to prove the following claim.

Claim 1. If $\bar{\theta} \in [\theta^1, \theta^2]$, there is a pure-strategy Nash equilibrium in which a proportion α of firms invest \underline{c} and the remaining $(1 - \alpha)$ invest \bar{c} , where α is given by equation (7).

Proof. See Appendix.

Claim 1 shows that the previous result does not depend on the fraction of low-investing firms being contiguous. The above conditions are independent of the exact configuration of firms along the circle, *i.e.*, location does not matter. However *ex-post*, each firm in the economy prefers to have less competition in its environment. Expected profit to firm i if both neighbours invest \underline{c} is $\alpha \underline{y} + (1 - \alpha) \bar{y} - \underline{c}$; if both its neighbours invest \bar{c} it is $\alpha \underline{y} + (1 - \alpha) \bar{y} - \bar{c}$; and it is $\alpha \underline{y} + (1 - \alpha) \bar{y} - (\underline{c} + \bar{c})/2$ if one neighbour invests \bar{c} and the other \underline{c} .

When all firms invest $\alpha \underline{c} + (1 - \alpha) \bar{c}$, the firms' total surplus is

$$(8) \quad W = n[\alpha \cdot (\underline{y} - \underline{c}) + (1 - \alpha) \cdot (\bar{y} - \bar{c})]$$

It follows that, when $\bar{y} - \bar{c} > \underline{y} - \underline{c}$, firms' surplus is maximised iff all firms invest \bar{c} . The opposite holds when $\bar{y} - \bar{c} < \underline{y} - \underline{c}$.

Note that, under constant or increasing returns to scale, the condition $\bar{y} - \bar{c} > \underline{y} - \underline{c}$ is always satisfied. Therefore the optimum is achieved when all firms invest \bar{c} (and generally increases with the number of firms investing \bar{c}).

If $\bar{\theta} \in (\theta^1, \theta^2)$ then there could be a co-ordination failure, with the firms implementing a strategy which is individually rational but socially inefficient. All firms could be made better off if it were possible to co-ordinate their activity to the high productivity equilibrium.

The social optimum is achieved when firms can co-ordinate their investment to the level of activity characterised by the highest productivity. When this happens, in equilibrium there will be no net competitive advantages among firms, because they all undertake the same level of investment.

Note that the game is supermodular in pure strategies, according to the definition of Milgrom and Roberts (1990). Supermodularity is an extension of the notion of strategic complementarities. These arise if "an increase in one player's strategy increases the *optimal strategy* of the other players" (Cooper and John, 1988, p. 442), which requires that the set of strategies be ordered. If payoffs are monotonic in the strategies of the other players and the supermodular game exhibits a multiplicity of symmetric equilibria, then these can be Pareto-ranked. Hence one has a *co-ordination game*, whereby a decentralised economy can find itself in a "bad" equilibrium. Individual, non-co-operative rationality prevents the economy from moving to a better equilibrium, even when such an equilibrium exists.

2.4 Uncertainty and equilibrium selection

In this section we assume that $\nu > 0$, that is, there is an idiosyncratic component in the uncertainty with regard to the competitive advantage of each firm. There are several ways to introduce uncertainty which could be relevant for macroeconomic modelling. According to a first interpretation, the firm perceives a noisy signal about the average competitive advantage. We assume that the noise is small and uniformly distributed about the true value of the parameter. A second interpretation is that firms observe a signal about their local competitive advantage, and these signals are correlated between neighbours (but not necessarily farther away on the variety space). The firms' exact behaviour will depend on the specific assumptions about the noise. Under both assumption an endogenous selection process will take place: in the first case firms will almost always co-ordinate on a certain equilibrium, in the second case there could still remain a region of indeterminacy.

In the presence of noise, the firm's behaviour should be modelled as a function both of its own signal and of its posterior belief regarding the possible signals of the other firms. Each firm makes inferences about the possible behaviour of the other firms given the possible signal they might receive, and chooses its reaction function to maximise its expected payoff. For signals far enough from the boundaries (that is $\bar{\theta}_i - \frac{1}{2} > \nu$ and $1 - \bar{\theta}_i > \nu$), the symmetry of the noise implies that firm i 's best predictor for the true value of the competitive parameter is its own signal. Moreover, the signal is also the best predictor for the neighbour's signal. These properties are crucial for the results.²

Firms have a dominant strategy for extreme values of the competitive advantage. Taking this into account when calculating their optimal reaction function, some strategies become dominant over a larger region of $\bar{\theta}$. In fact, if the noise is sufficiently small, then the global game is dominant solvable, *i.e.* iterated elimination of strictly dominated strategies will lead firms to co-ordinate their behaviour for any possible value of the

²In general to use Carlsson and van Damme's method we need to consider the posterior beliefs of players. However, in our model firms' payoffs are linear in $\bar{\theta}$, hence it is sufficient to look at expectations.

competitive advantage.

Suppose firm i 's signal is $\bar{\theta}_i$, where $\bar{\theta}_i - \theta^1 \gg \epsilon$. The firm knows with certainty that $\bar{\theta}$ is larger than θ^1 . However, if its behaviour is part of a consistent plan, it must take into account what its neighbours will do when their signal is $\bar{\theta}_i - \epsilon$. The neighbours' behaviour in turn depends on what firm i will do for $\bar{\theta}_i - 2\epsilon$, *etc.* Hence, each firm must consider the optimal behaviour for signals smaller than θ^1 .

For the first interpretation, following a logic similar to Carlsson and van Damme (1993a) we are able to prove the following proposition.

Proposition 1. If $v > 0$ then iterated elimination of strictly dominated strategies results in each firm investing \underline{c} if $\bar{\theta}_i < (\theta^1 + \theta^2)/2$, and \bar{c} if $\bar{\theta}_i > (\theta^1 + \theta^2)/2$.

Proof of Proposition 1.

The investment behaviour of each firm can be described by a function $f_i[-\frac{1}{2}, 1] \rightarrow [0, \bar{a}]$. As all firms are restricted to investing \underline{c} when $\bar{\theta}_i < \theta^1$ (and in particular firms $i-1$ and $i+1$), then the expected payoffs to firm i when it observes $\bar{\theta}_i = \theta^1$ are:

$$\frac{1}{2} \left(\frac{Y}{n} - \underline{c} \right) + \frac{1}{2} A$$

when it invests \underline{c} , and

$$\frac{1}{2} \left(2\theta(\epsilon) \frac{Y}{n} - (\underline{c} + \epsilon) \right) + \frac{1}{2} B$$

when it invests $\underline{c} + \epsilon$, where A and B are defined as:

$$A = \int_{[0, v]^2} \{ [1 - \theta_i(|f_{i+1}(\bar{\theta}_{i+1})|)] + [1 - \theta_{i-1}(|f_{i-1}(\bar{\theta}_{i-1})|)] - \underline{c} \} \frac{1}{4 \cdot v^2} d\bar{\theta}_{i-1} d\bar{\theta}_{i+1}$$

$$B = \int_{[0, \nu]^2} \{ [1 - \theta_i(|\epsilon - f_{i+1}(\bar{\theta}_{i+1})|)] + [1 - \theta_{i-1}(|\epsilon - f_{i-1}(\bar{\theta}_{i-1})|)] \} - (\underline{c} + \epsilon) \frac{1}{4 \cdot \nu^2} d\bar{\theta}_{i-1} d\bar{\theta}_{i+1}$$

The inequality $A \geq B$ follows from the definition of $\theta(\cdot)$ and from substituting $\bar{\theta}_i = \theta^1$ from equation (3). In order to show that \underline{c} dominates all other investment strategies it remains to prove the following inequality:

$$\frac{Y}{n} - \underline{c} > 2\theta(\epsilon) \frac{Y}{n} - (\underline{c} + \epsilon)$$

Substituting $\bar{\theta}_i = \theta^1$ from (3) we obtain that the above inequality is equivalent to

$$\bar{y} > \frac{Y}{n}$$

The last inequality holds since Y/n (the expected average demand) is at most equal to $(\underline{y} + \bar{y})/2$ (because half of the firms are expected to receive signals smaller than θ^1 , and therefore to invest \underline{c}), and in particular is smaller than \bar{y} .

Therefore iterated elimination of strictly dominated strategies forces all firms to invest \underline{c} for $\bar{\theta}_i = \theta^1$. Denote by $\tilde{\theta}$ the smallest value of $\bar{\theta}_i$ for which iterated elimination of strictly dominated strategies does not force firms to invest \underline{c} . From the above inequalities it is clear that

$$(9) \quad \tilde{\theta}(y) \geq \frac{\theta^1 + \theta^2}{2}$$

Symmetrically, all firms are restricted to investing \bar{c} when $\bar{\theta}_i > \theta^2$ (and in particular firms $i-1$ and $i+1$). The expected payoffs to firm i when it observes $\bar{\theta}_i = \theta^2$ are:

$$\frac{1}{2} \left(\frac{Y}{n} - \bar{c} \right) + \frac{1}{2} C$$

when it invests \bar{c} , and

$$\frac{1}{2}(2(1 - \theta(\epsilon))\frac{Y}{n} - (\underline{c} - \epsilon)) + \frac{1}{2}D$$

when it invests $\bar{c} - \epsilon$, where

$$C = \int_{[-v,0]^2} \{[\theta_i(|\bar{a} - f_{i+1}(\bar{\theta}_{i+1})|) + \theta_{i-1}(|\bar{a} - f_{i-1}(\bar{\theta}_{i-1})|)] - \bar{c}\} \frac{1}{4 \cdot v^2} d\bar{\theta}_{i-1} d\bar{\theta}_{i+1}$$

$$D = \int_{[-v,0]^2} \{[\theta_i(|\bar{a} - \epsilon - f_{i+1}(\bar{\theta}_{i+1})|) + \theta_{i-1}(|\bar{a} - \epsilon - f_{i-1}(\bar{\theta}_{i-1})|)] - (\bar{c} - \epsilon)\} \frac{1}{4 \cdot v^2} d\bar{\theta}_{i-1} d\bar{\theta}_{i+1}$$

For similar considerations as above, $C \geq D$. Investing \bar{c} dominates all other strategies when:

$$\frac{Y}{n} - \bar{c} > 2(1 - \theta(\epsilon))\frac{Y}{n} - (\bar{c} - \epsilon)$$

Substituting $\bar{\theta}_i = \theta^2$ from equation (4) we obtain

$$\underline{y} < \frac{Y}{n}$$

The last inequality holds since Y/n is at least equal to $(\underline{y} + \bar{y})/2$ and in particular is greater than \underline{y} .

Therefore iterated elimination of strictly dominated strategies forces all firms to invest \bar{c} for $\bar{\theta}_i = \theta^2$. Denote by $\hat{\theta}$ the largest value of $\bar{\theta}_i$ for which iterated elimination of strictly dominated strategies does not force firms to invest \bar{c} . From the above inequalities it is clear that

$$(10) \quad \hat{\theta}(y) \leq \frac{\theta^1 + \theta^2}{2}$$

Additionally, $\tilde{\theta} \leq \hat{\theta}$. Combining this with equations (9) and (10) we obtain

$$\frac{\theta^1 + \theta^2}{2} \leq \tilde{\theta}(y) \leq \hat{\theta}(y) \leq \frac{\theta^1 + \theta^2}{2},$$

That is, firms' behaviour is as described in Proposition 1.

QED

Aggregate behaviour. If θ is such that the support of individual firms' posteriors $[\theta \pm 2v]$ does not include the value $(\theta^1 + \theta^2)/2$, then all firms in the economy co-ordinate on the risk-dominant equilibrium. Otherwise, the proportion of firms investing \underline{c} is $1/2 + [(\theta^1 + \theta^2)/2 - \theta]/2v$ if $\theta \leq (\theta^1 + \theta^2)/2$, and $[\theta - (\theta^1 + \theta^2)/2]/2v$ otherwise. In any case, if v is small, firms will almost always co-ordinate.

Comment. This case shows that, if shocks are correlated among firms, then even a very small amount of uncertainty will lead firms to endogenously co-ordinate (almost always). Note however that the selection process is not guided by Pareto optimality.

2.4.1 Locally correlated signals

In this case, the information structure is related to the local interaction structure in the economy. The signal perceived by firms is correlated amongst neighbouring firms and uncorrelated otherwise. This can be rationalised on the grounds that neighbouring firms compete over overlapping segments of the market.

In order to analyse firms' behaviour in this setting, we first need to specify their expectations about the "average behaviour" in the economy. We can now make use of the fact that firm i 's payoff depends on the behaviour of its immediate neighbours and on the average output of the other $n-3$ firms in the economy. Let the competitiveness

parameter on arc $(i-1, i)$ be $\theta_i = \bar{\theta} + \eta_i$, where $\bar{\theta}$ is a r.v. with expected value $\hat{\theta}$ and variance $\sigma_{\bar{\theta}}^2$, and η_i is an orthogonal i.i.d. with expected value $E(\eta_i) = 0$ and variance σ_{η}^2 . Firm i receives a signal equal to the average of the competitiveness parameters on both sides: $\bar{\theta}_i = (1/2)(\theta_i + \theta_{i+1}) = \bar{\theta} + (1/2)(\eta_i + \eta_{i+1})$. We are then able to prove the following Proposition.

Proposition 2. If $\hat{\theta} < \theta^1$, then all firms invest \underline{c} if $\bar{\theta}_i < \theta^2$, and \bar{c} otherwise. If $\hat{\theta} > \theta^2$, all firms invest \bar{c} for $\bar{\theta}_i > \theta^1$ and \underline{c} otherwise. If $\hat{\theta} \in (\theta^1, \theta^2)$, each firm will invest \underline{c} for $\bar{\theta}_i < \hat{\theta}$ and \bar{c} otherwise.

Proof of Proposition 2.

The solution for the signal extraction problem is: $E(\bar{\theta}_{i+1}|\bar{\theta}_i) = E(\bar{\theta}_{i-1}|\bar{\theta}_i) = \gamma \cdot \bar{\theta}_i + (1 - \gamma) \cdot \hat{\theta}$ where $\gamma \equiv \frac{1/4 \cdot \sigma_{\eta}^2}{\sigma_{\bar{\theta}}^2 + 1/2 \cdot \sigma_{\eta}^2} \in (0, 1)$.

As in the proof of Proposition 1, all firms are restricted to investing \underline{c} when $\bar{\theta}_i < \theta^1$ (and in particular firms $i-1$ and $i+1$). The expected payoffs to firm i when it observes $\bar{\theta}_i = \theta^1$ are:

$$\left[\frac{1}{2} - g(\gamma)\right] \cdot \left(\frac{Y}{n} - \underline{c}\right) + \left[\frac{1}{2} + g(\gamma)\right] \cdot A$$

when it invests \underline{c} , and

$$\left[\frac{1}{2} - g(\gamma)\right] \cdot \left(2\theta(\epsilon)\frac{Y}{n} - (\underline{c} + \epsilon)\right) + \left[\frac{1}{2} + g(\gamma)\right] \cdot B$$

when it invests $\underline{c} + \epsilon$, where $g(\gamma) \in \left(-\frac{1}{2}, \frac{1}{2}\right)$ is the correction to the probability that firm i attaches to the event that its neighbours observed a signal smaller than $\bar{\theta}_i$, and A and B are as defined in the proof of Proposition 1, with the difference that here we integrate over $[0, \eta]^2$ instead of over $[0, v]^2$.

The expected payoff from investing \underline{c} is greater than from investing $\underline{c} + \epsilon$ if

$$\bar{y} > \frac{Y}{n}$$

Here, unlike in Proposition 1, the expected average income, Y/n , is obtained by replacing $\hat{\theta}$ into equation (5). Let \hat{y} denote expected average income. Iterated elimination of strictly dominated strategies will force firms to invest \underline{c} until the appropriate inequalities (see the previous proof) no longer hold. This implies that firms will invest \underline{c} for $\bar{\theta}_i \leq \hat{\theta}$. Similarly, all firms will invest \bar{c} for $\bar{\theta}_i \geq \hat{\theta}$. Hence the result.

QED

Proposition 2 says that, if firms' expectations about total output are very low, they will co-ordinate on the low investment equilibrium in the region of multiplicity. If expectations are very high, they will co-ordinate on the high investment equilibrium. For intermediate values of the expectations, the endogenous selection will ensure co-ordination: the critical threshold will be consistent with the economy-wide expected value of the signal, $\hat{\theta}$.

Aggregate behaviour. The aggregate behaviour depends on the statistical distribution of $\bar{\theta}$. The proportion of firms investing \underline{c} is given by $Prob(\bar{\theta} \leq \hat{\theta})$. As in the case of Proposition 1, if η is small, firms will almost always co-ordinate with their immediate neighbours.

The previous result depends on all firms sharing the same expectation regarding average output. If this assumption is removed, the result does not hold. Think, for example, on firms as receiving an additional signal about the state of the economy in the form of \hat{y}_i . Behaviour in this multi-dimensional signal space is, in general, much

more complicated than described before. It is still possible to see, using the previous calculations, that the following holds: if $\hat{\theta}_i(\hat{y}_i)$, and $\bar{\theta}_i$ are not too "far apart", iterated elimination of dominated strategies will force all firms to switch from \underline{c} to \bar{c} at $\tilde{\theta} = (\theta^1 + \theta^2)/2$. Hence, the relative proportion of firms investing \underline{c} and \bar{c} is $F(\tilde{\theta})$ and $1 - F(\tilde{\theta})$, where $F(\cdot)$ is the cumulative distribution function. This, however, is not true when the two observations are "far apart". Assume, for example, that $\bar{\theta}_i = \theta^1 + \gamma$ and $\hat{\theta}_i = \theta^2 - \gamma$ (where γ is small). Two conflicting forces operate on the decision maker: on the one hand, \bar{c} is still riskier than \underline{c} with regard to their local competition; on the other, the expected aggregate income in the economy is high, thus making a switch to \bar{c} more profitable. Firms will invest \underline{c} for values of individual signals less than θ^1 or slightly above it, and \bar{c} for values greater than θ^2 or slightly smaller (the exact boundaries depend on the signal about the state of the economy). Iterated elimination of strictly dominated strategies leaves a region of indeterminacy. However, we still obtain some endogenous co-ordination over regions with a multiplicity of Nash equilibria.

2.5 Conclusions

This paper presents a macroeconomic model of co-ordination failures based on local oligopoly. The key parameter for firms is the competitive advantage they can gain over their neighbours by undertaking higher levels of investment. For a non-singular range of values of the competitive advantage parameter the economy exhibits multiple equilibria. The decentralised market outcome could be socially inefficient because of the firms' failure to co-ordinate on a high-productivity equilibrium. It is shown that the neighbourhood structure described in this paper can be responsible for multiplicity of equilibria and market failures.

In the absence of idiosyncratic noise, the set of possible equilibria depends on the competitiveness parameter, θ . Either firms have a dominant strategy, or there is a multiplicity of equilibria. If firms have only two investment strategies, the proportions engaging in a

high or low level of investment depend on the exact value of θ , but are independent of the exact configuration of firms in the economy. In the absence of explicit co-ordination devices, the economy could settle on any of the possible equilibria.

However, if one introduces a stochastic element in the economy, and allows firms to observe imperfect signals about the competitiveness conditions in the local output market, an endogenous equilibrium selection process could take place. In particular, when firms' noisy signals are correlated between neighbours, iterated elimination of strictly dominated strategies significantly reduces the set of possible market outcomes. Firms choose a low level of investment for "small" signals, and a high level for "large" signals. This would correspond to the adoption of the risk-dominant strategy, in the sense of Harsanyi and Selten (1988). If the firms' expectation about the state of the economy is the same for all firms, then they will all switch from low to high levels of investment at a critical value of θ . This value is the unique θ for which the expected average output for the economy is equal to its value in the intermediate investment equilibrium identified in section 3. If firms' expectations about the state of the economy depend on an additional signal, then the switch from low to high investment will occur at $(\theta^1 + \theta^2)/2$, if this signal is not too different from their own observation of θ .

The model suggests that co-ordination could spontaneously emerge in the presence of noise. Expectations play a crucial role in determining the macroeconomic outcome. Indeterminacy may not be an endemic feature of macroeconomic models with multiple equilibria.

Appendix

Proof of Claim 1.

In Section 3 we have already shown that the strategy profile where all firms invest \underline{c} is a Nash equilibrium if and only if $\bar{\theta} \leq \theta^2$, and that the case where all firms invest \bar{c} is a Nash equilibrium if and only if $\bar{\theta} \geq \theta^1$. To complete the proof of Proposition 1, it remains to be shown that, for $\bar{\theta} \in (\theta^1, \theta^2)$, every configuration in which a proportion $\alpha(\bar{\theta}) = [1/(\bar{y} - \underline{y})][\bar{y} - (\bar{c} - \underline{c})/(2\bar{\theta} - 1)]$ of firms invest \underline{c} , and $(1 - \alpha(\bar{\theta}))$ invest \bar{c} is a Nash equilibrium.

Each firm in the economy can be in exactly one of the following six configurations of investment behaviour and neighbourhood structure:

- (1) the firm invests \underline{c} and both its neighbours invest \underline{c} ;
- (2) the firm invests \underline{c} and both its neighbours invest \bar{c} ;
- (3) the firm invests \underline{c} , one of its neighbours invests \underline{c} and the other \bar{c} ;
- (4) the firm invests \bar{c} , one of its neighbours invests \underline{c} and the other \bar{c} ;
- (5) the firm invests \bar{c} and both its neighbours invest \underline{c} ;
- (6) the firm invests \bar{c} and both its neighbours invest \bar{c} .

We next show that, if the relationship between $\bar{\theta}$ and α is as in equation (7), the firm will have no incentive to change its behaviour in any of the possible configurations.

Cases (1) and (2): Firm i 's payoff from investing \underline{c} is $\pi_i = (1/n)(Y) - \underline{c}$, whereas if it invests \bar{c} it will receive $\pi_i = (2\bar{\theta}/n)(Y) - \bar{c}$. The difference (the incentive to deviate) is:

$$\frac{Y}{n}(1 - 2\bar{\theta}) - (\underline{c} - \bar{c}) = \frac{\alpha n \underline{y} + (1 - \alpha)n \bar{y}}{n} \left(1 - 1 - \frac{\bar{c} - \underline{c}}{\alpha \underline{y} + (1 - \alpha)\bar{y}}\right) - (\underline{c} - \bar{c}) = 0$$

Therefore the firm has no incentive to change its investment strategy in any of these

two cases.

Cases (3) and (4): Firm i 's payoff from investing \underline{c} is $\pi_i = (1/n)[(3/2) - \bar{\theta}](Y) - \underline{c}$, whereas if it invests \bar{c} it will receive $\pi_i = (1/n)[(1/2) + \bar{\theta}](Y) - \bar{c}$. The difference (the incentive to deviate) is:

$$\frac{Y}{n} \left(\frac{3}{2} - \bar{\theta} - \frac{1}{2} - \bar{\theta} \right) - (\underline{c} - \bar{c}) = \frac{\alpha \underline{y} + (1 - \alpha) n \bar{y}}{n} \left(1 - 1 - \frac{\bar{c} - \underline{c}}{\alpha \underline{y} + (1 - \alpha) \bar{y}} \right) - (\underline{c} - \bar{c}) = 0$$

Therefore the firm has no incentive to change its investment strategy in any of these two cases.

Cases (5) and (6): Firm i 's payoff from investing \underline{c} is $\pi_i = [2(1 - \bar{\theta})/n](Y) - \underline{c}$, whereas if it invests \bar{c} it will receive $\pi_i = (1\theta/n)(Y) - \bar{c}$. The difference (the incentive to deviate) is:

$$\frac{Y}{n} (2 - 2\bar{\theta} - 1) - (\underline{c} - \bar{c}) = \frac{\alpha \underline{y} + (1 - \alpha) n \bar{y}}{n} \left(1 - 1 - \frac{\bar{c} - \underline{c}}{\alpha \underline{y} + (1 - \alpha) \bar{y}} \right) - (\underline{c} - \bar{c}) = 0$$

Therefore the firm has no incentive to change its investment strategy in any of these two cases.

References

- CARLSSON, H., AND E. VAN DAMME (1993a), "Global Games and Equilibrium Selection", *Econometrica* **5**, pp. 989-1018.
- CARLSSON, H., AND E. VAN DAMME (1993b), "Equilibrium Selection in Stag-Hunt Games", in Ken Binmore, A. Kirman and P. Tani, *Frontiers of Game Theory*, MIT Press, Cambridge (MA.), pp. 237-253.
- COOPER, R., AND A. JOHN (1988), "Coordinating Coordination Failures in Keynesian Models", *The Quarterly Journal of Economics* **3**, pp. 441-463.
- DIXON, H., AND N. RANKIN (eds.) (1995), *The New Macroeconomics*, Cambridge

University Press, Cambridge.

HARSANYI, J., AND R. SELTEN (1988), *A General Theory of Equilibrium Selection in Games*, MIT Press, Cambridge (MA).

JOVANOVIC, B. (1989), "Observable Implications of Models with Multiple Equilibria", *Econometrica* **6**, pp. 1431-1437.

KALDOR, N. (1940), "A Model of the Trade Cycle", *The Economic Journal* **197**, pp. 78-92.

MANKIW, G., AND D. ROMER (eds.) (1991), *New Keynesian Economics*, MIT Press, Cambridge (MA).

MANKIW, G., AND M. WHINSTON (1986), "Free Entry and Social Inefficiency", *Rand Journal of Economics* **1**, pp. 48-58.

MILGROM, P., AND J. ROBERTS (1990), "Rationalizability, Learning and Equilibrium in Games with Strategic Complementarities", *Econometrica* **6**, pp. 1255-1277.

MORRIS, S. (1995), "Co-operation and Timing", CARESS Working Paper No. 95-05, University of Pennsylvania.

SALOP, S. (1979), "Monopolistic Competition with Outside Goods", *The Bell Journal of Economics* **10**, pp. 141-156.

SHAKED, A., AND J. SUTTON (1987), "Product Differentiation and Industrial Structure", *The Journal of Industrial Economics* **2**, pp. 131-146.

SHIN, H. (1995), "Defining a Notion of Confidence in an Economy with Differential Information", University of Southampton, mimeographed.

SILVESTRE, J. (1993), "The Market-Power Foundations of Macroeconomic Policy", *Journal of Economic Literature* **31**, pp. 105-141.

- SOLOW, R. (1986), "Monopolistic Competition and the Multiplier", in W.P. Heller, R.M. Starr, and D.A. Starrett (eds.), *Equilibrium Analysis. Essays in Honor of Kenneth J. Arrow*, Vol. II, Cambridge University Press, Cambridge, pp. 301-315.
- WEITZMAN, M. (1982), "Increasing Returns and the Foundations of Unemployment Theory", *Economic Journal*, **92**, pp. 787-804.

Chapter 3

Modelling Limited Attention

3.1 Introduction

Traditional game theoretical analysis assumes that players have perfect information regarding the rules of the game. Various extensions of the model, where players are faced with some (fixed amount of) uncertainty, have been successfully analysed. (as in Harsanyi (1967/68), and Carlsson and Van Damme (1993).) In those models bayesian agents make their choice of best response based on their partial information. Uncertainty is fixed, and is modelled as a move of nature. Our intention is to further extend the scope of the theory by investigating games of endogenous information acquisition. That is, we look at how bayesian agents act in situations where they can *choose* how much information to obtain. In particular we are interested in the case where acquiring information is not free. Given the time and mental effort it takes to pay attention to details, it is not cost-effective to learn more about the payoff when, for example, it is already clear that one strategy is dominant. In situations like this, therefore, players have an incentive to use some sort of classification (of sets of games) based on their degree of similarity. In this paper we present a general theoretical framework that can be used to analyse situations in which people have to decide how much attention to pay to a particular game. Paying more attention to the game yields better and more precise information, which helps the players to choose the optimal action. However, paying attention is costly and the players therefore face a trade-off.

An attractive feature of a model of costly endogenous information is that its assumptions are in line with findings from cognitive psychology. Scientists there have been studying attention, and the properties of the attention allocating mechanism for the last three decades. Their findings are, in general, in support of the assumptions that attention is cognitively costly, and that individuals use some sort of classifications based on the degree of similarity to reduce the amount of attention they need to pay to each task. Clinical studies showing that anxious patients exhibit attention style that is very "actively engaged" with the external world (as in Beck & Emery (1985)) support the assumption that attention is "expensive". In economic terms, the marginal cost of paying attention to a particular task, for these people, exceeds their marginal benefit from it. In another set of experiments subjects were asked to perform simple tasks simultaneously, where the similarity of any pair of tasks could be easily measured. (for example, the task is to identify certain objects in a picture, where some of the pictures are very similar.) The results show an inverse relationship between the degree of similarity of tasks and the ability to distinguish between them. (See Eysenck and Keane (1990) for an introduction and a survey of attention experiments.) These kinds of findings support the assumption that people do classify. In general we would expect that an optimal allocation mechanism will divide attention in such a way that the marginal benefit from increasing attention to one particular task is exactly off-set by the marginal cost of it.

Experimental psychologists have been mainly concerned with one person *decision problems*. For such problems there is usually a unique solution. We introduce a general framework for modelling limited attention in *games*. We assume the existence of a one-dimensional, continuous state space, where each space corresponds to a game, and a probability measure over it. Here we can study the decisions of players whether to split the set, and where exactly to do that. Alternatively, one can study a model where the set is discrete. For some applications, like market entry games, this could be more appropriate. (see chapter 4 for such a model.) When we model games, what is optimal depends on the behaviour of the other people involved. The modeller is faced with an

additional problem of choosing whether or not a player is allowed to observe the choice of information obtained by the other players. That is, should the model be of a simultaneous move, or one where the information acquisition takes place before the (conditional) choice of strategies? Though the latter suggestion seems more natural at first, it allows for the possibility that players will condition their behaviour on the choice of information of their opponents. In models of bounded rationality this could, sometimes, be counter intuitive.

We introduce the general framework of *meta-games*: games where players choose information and (conditional) strategies, and nature chooses the state of the world. We analyse the equilibria of such meta-games, for both a one-stage and a two-stage meta-game. In particular we focus on the differences between the two cases, especially those differences that hold even when the costs of acquiring information tends to zero. When players simultaneously choose their partition and behaviour, we prove two limiting results for the case where the information costs becomes very small. We suggest a selection criteria for those equilibria. We also discuss the Nash equilibria of the meta-game in the cases where information is not cheap. However, those results depend on the exact structure of the information costs and on the payoffs of the underlying games. We then analyse the (sequential) equilibria of a two-stage game, where players first choose where to split the state space, and only then choose conditional strategies. We find conditions under which the outcomes of the sequential equilibria differ from those of the equivalent one stage game even when the cost of information becomes very small. To illustrate these differences we look at two examples: In the first, a person classifies the set of all possible opponents he might have to play a co-operation game with, as "friends" and "strangers". We show that if the game is modelled in two stages, a rational player might end-up co-operating where he would not have, had he been perfectly informed about the state of the world. Our second example demonstrates that in a two-stage game, a person might credibly threat not to use his information, therefore forcing the other player to treat him as "ignorant". (making himself, overall better off.)

The idea that individuals' choice of actions is related to their information about similar

situations, is not new to game theory and economics. Schelling (1960) argues that such an approach will greatly contribute to the predictiveness of game-theory, while Kreps (1990) suggested that any realistic learning model should take into account inferences based on similarity. Such a fictitious learning model was eventually presented by Li Calzi (1995). A more general approach to the role of similarity in decision making was presented by Rubinstein (1988). Macroeconomists have noted that sometimes firms will choose not to change their pricing policy as a result of small changes to their costs. This is because the benefit from doing so does not exceed the costs of changing labels, advertisements etc. (the so called "menu-costs", or "near rationality" arguments, see Mankiw (1985) and Akerlof and Yellen (1985)). This kind of models share with ours the property that there is some classification of the set of possible inputs and a change in action is only considered when some variables move to a different class.

The rest of the paper is organized in the following way: Section 2 introduces the general framework of meta-games, the different notions of strategies, payoffs, and equilibria. In section 3 we analyse the Nash equilibria of the simultaneous-move game. Section 4 analyses the (sequential) equilibria of a two-stage game. Section 5 concludes by discussing the appropriateness of the one and two stage games from the point of view of economic modelling.

3.2 The Model

The general framework is as follows; There is a state space Ω which is an interval in \mathcal{R} . Every state $\omega \in \Omega$ corresponds to a two player game $g(\omega) = (A_1, A_2, u_1^\omega, u_2^\omega)$. That is, the players already know the strategies available but they do not know yet the payoff functions, because those depend on the true state ω . It is assumed that the payoffs u_i^ω are continuous in ω . Nature selects the game to play according to a probability measure ρ on Ω , where ρ is common knowledge. Players can gather (partial) information about the state by choosing a partition $P \in \mathcal{P}(\Omega)$, where $\mathcal{P}(\Omega)$ denotes the set of all convex

partitions of Ω . Formally $P = \{P_1, P_2, \dots\}$ is a convex partition of Ω if it satisfies the following four conditions:

(i) $\rho(P_i) > 0$ (all i),

(ii) $P_i \cap P_j = \emptyset$ (all $i \neq j$), and

(iii) for all $\omega \in \Omega$ there exists P_i with $\omega \in P_i$.

(iv) if $\omega_1 \in P_i$ and $\omega_2 \in P_i$, then $\forall \alpha \in [0, 1], \alpha \cdot \omega_1 + (1 - \alpha) \cdot \omega_2 \in P_i$.

We say that P^* is a refinement of P if for each $P_i \in P^*$ there exists some $P_{i'} \in P$ such that $P_i \subseteq P_{i'}$. We say that P^* is a proper refinement of P if, in addition, $P \neq P^*$.

If any of the players chooses not to obtain any information, his beliefs will be governed solely by the prior distribution $\rho(\cdot)$. Otherwise his beliefs are determined by ρ conditional on P .

We incorporate the choice of a partition with the choice of action in one meta-game Γ . We will distinguish between two kinds of meta-games Γ_1 and Γ_2 ; In Γ_1 - the game is of simultaneous moves. A strategy for player i is a pair (P^i, s_i) , where P^i denotes i 's partition and $s_i : P^i \rightarrow A_i$ assigns to each information set of P^i one action. Γ_2 is a two stage game, where partitions are chosen in the first stage and (conditional) behaviour in the second stage. A strategy in Γ_2 is also a pair (P^i, s_i) , with the difference that $s_i : P_i \times \mathcal{P}(\Omega) \rightarrow A_i$ *i.e.* players are allowed to condition their second-stage behaviour not only on their own information, but also on the choice of information of the other player. The payoffs in (both) meta-game are the expected payoffs in $g(\omega)$ over all $\omega \in \Omega$ minus the cost of filtering, $c(P_i)$. We assume that $c : \mathcal{P}(\Omega) \rightarrow \mathcal{R}^+$ is bounded from above by \bar{c} and that the cost of filtering is increasing in the informativeness of the filter. That is if P' is a proper refinement of P then $c(P') > c(P)$.

More formally, fix Ω and fix the partitions of both players: P^1 and P^2 . Remember that $u_i^\omega(s_1, s_2)$ is player i 's payoff when nature chose the state ω and player j chose strategy s_j . We can now define the (expected) payoff in Γ for player i as:

$$\int_{\Omega} u_i^\omega(s_1(P^1(\omega)), s_2(P^2(\omega))) d\rho(\omega) - c(P^i)$$

As a reference point for the equilibria of Γ_1 and Γ_2 we will consider what players would do if they were perfectly informed about the state of the world. If both players are (expected) utility maximizers this implies that for every game chosen by nature they will play the equilibrium (if it exists). The following definition will be useful as a benchmark for the rest of our analysis:

Definition 1: Let $s(\omega)$ be a strategy profile in Γ . (the definition holds for both Γ_1 and Γ_2 .) We say that $s(\omega)$ is a $\Gamma - Nash$ if for almost all $\omega' \in \Omega$, $s(\omega')$ is a Nash equilibrium of $g(\omega')$.

If, for all ω , the game $g(\omega)$ has a unique Nash equilibrium then the set of all $\Gamma - Nash$ strategy profiles contains a single element. On the other hand if Ω contains an interval $[a, b]$ with $\rho([a, b]) > 0$ where for each $\omega \in [a, b]$ the game $g(\omega)$ exhibits multiple Nash equilibria, then the set of all $\Gamma - Nash$ is infinite. (and uncountable.)

3.3 Results for the Simultaneous Move Game

Acquiring information is costly, and therefore it is a strictly dominated strategy in the one stage game, Γ_1 , to choose to learn and not use the information. In an equilibrium strictly dominated strategies are never used. We will write s^* for the strategy (I, s) , where $s^*(\omega) = s(\omega)$ for all $\omega \in \Omega$, and where I is the smallest set that provides sufficient information to play s .

The set of Nash equilibria of Γ_1 depends on the exact structure of the information-cost function. If information is very cheap, or free, players will choose to obtain as much as possible of it (this is no longer true for Γ_2 , or any other game where information is obtained publicly, as we demonstrate in the next subsection.) If information is very expensive, players will choose in many cases not be informed. If information costs are high enough, any Nash equilibria of Γ_1 will imply that players will not gather any information and instead will play the Nash equilibrium of the "expected" game, $\int_{\Omega} g(\omega) d\rho(\omega)$. In the other extreme, if information becomes very cheap players will play as though they have perfect information. We are now able to prove the following limiting result.

Proposition 1 Fix Ω and $g(\omega)$. Let Γ^c denote the meta-game in which $c(\cdot)$ is fixed¹. Let s_c^* be a pure equilibrium of Γ^c . Then, if the limit exists, $\lim_{\bar{c} \rightarrow 0} s_c^*(\omega)$ is an equilibrium of $g(\omega)$ for (almost) all $\omega \in \Omega$.

Proof. Suppose on the contrary that there exists an interval $[a, b]$ where for any state $\omega' \in [a, b]$ (except maybe for $\omega = a$ and $\omega = b$) there is a player who is not choosing the optimal action in these states in the limit. For small information costs this remains true. This player can profitably deviate by refining his partition in such a way that allows him to play the best response to his opponent's strategy in the interval $[a, b]$. The extra costs of such a deviation is a function of the (minimal) number of new classes needed to implement such a strategy. Let $\kappa > 0$ the difference between $c(s^*)$ and the cost of his new, refined strategy. However, as $\bar{c} \rightarrow 0$ the increase in expected payoff from playing best response in $[a, b]$ exceeds κ . □

We are able to prove, under some additional assumptions, that the complementary statement also holds. Any selection of (pure strategy) Nash equilibria in the underlying

¹That is, the cost function is determined for every possible partition of Ω . A simple case is $c(F) = n \cdot c'$ where c' is a constant and n is the number of classes in P .

games can be combined to make a Nash equilibrium of the meta-game Γ_1 if the costs of implementing such a strategy are small enough.

Proposition 2 Fix Ω and $g(\omega)$. Let $s(\omega)$ be a Γ – *Nash* strategy profile in Γ_1 . Denote by Γ^c the meta game where $c(\cdot)$ is fixed. Assume that for almost all ω , all Nash equilibria of $g(\omega)$ are strict, and that the costs of partitions of the same size are equal. Then if $s(\omega)$ is implementable by a finite, convex partition of Ω , there exists a Nash equilibrium s_c^* of Γ^c such that $\lim_{\bar{c} \rightarrow 0} s_c^*(\omega) = s(\omega)$ for (almost) all $\omega \in \Omega$.

Proof. Deviations in which players do not change their partitions could not be profitable. This follows immediately from the definition of $s(\omega)$, as a Γ – *Nash* strategy in which players play their best response in $g(\omega)$ for almost all $\omega \in \Omega$. In addition, the assumption that all Nash equilibria are strict implies that, by changing her partition, there exists a non-empty set of intervals \mathcal{I} in which player i does not play her best response. Such a deviation might be profitable if the cost of the new partition, P' is smaller than of that implementing s^* , P . From the assumptions that P is finite, and that the costs of partitions of the same size are equal, we must conclude that in P' the player is partitioning Ω to a smaller number of sub-sets.

We next show that there exists $k > 0$ such that for all possible deviations (with properties described above) the loss in expected payoff to player i from not playing her best response in \mathcal{I} is at least k . Suppose on the contrary that such a k does not exist. Thus there must be a sequence of strategies, s_j such that the number of sub-sets in any $P(s_j)$ is smaller than that in P , and that the loss in expected payoff for player i converges to zero, as $j \rightarrow \infty$. This, in turn, implies that $\rho(\mathcal{I}) \rightarrow 0$ as $i \rightarrow \infty$. However, this means that $P(s_j)$ must converge to \bar{P} where P is a proper refinement of \bar{P} . This outcome stands in contradiction with our assumption that players have always a unique best response. (*i.e.* all Nash equilibria are strict.) We therefore conclude that such $k > 0$ exists.

As \bar{c} goes to zero, the gains from choosing a partition with less sub-sets becomes

smaller than the losses from not playing best response in \mathcal{I} . Players will therefore have no incentive to deviate from $s(\omega)$. Hence, our result. \square

Typically, Γ_1 will have many equilibria, even when the information costs function is fixed. Of course, any of the known equilibrium selection theories can be applied here, but there is an additional selection criteria which seems appealing in the framework of meta-games; Selection based on robustness to increasing information costs. For any $\Gamma - Nash$ strategy profile, s , we know from Proposition 2 that if information costs tends to zero, then s^* will become an equilibrium of Γ_1 . Suppose now that the cost of acquiring information increases. Except for the case where both players do not partition Ω at all to implement their $\Gamma - Nash$ strategy, there will be a minimal threshold value, $t(s)$, such that for costs greater than that $t(s)$, there will be a player with an incentive to deviate from s . For any two given equilibria of Γ_1 the one with the higher threshold value is more robust to an increase in information costs.

The following example illustrates the last argument:

	α	β
α	1, 1	$\omega \cdot (1 - \omega), 4$
β	4, 0.5	0, 0

Example 1. $\Omega = [-1, 2]$

Notice that when $\omega \in (0, 1)$, then $\omega \cdot (1 - \omega) > 0$ and the game $g(\omega)$ has two pure-strategy Nash equilibria: (α, β) and (β, α) . For all other values of ω only the latter is an equilibrium of $g(\omega)$. Consider now the meta-game where ρ is uniform over the interval $\Omega = [-1, 2]$. If costs of information becomes small enough the following $\Gamma - Nash$ strategies will constitute a Nash equilibrium of Γ_1 : s_1 , where players 1 follows

the strategy $\{[-1, 0], (0, 1), [1, 2]\}, (\beta, \alpha, \beta)$ while player 2 plays $\{[-1, 0], (0, 1), [1, 2]\}, (\alpha, \beta, \alpha)$. Alternatively players might end up playing a different $\Gamma - Nash$ equilibrium, s_2 where player 1 plays $\{\{\Omega\}, \beta\}$, and 2 plays $\{\{\Omega\}, \alpha\}$. (*i.e.* neither of the players gather any information.) It should be clear from the above discussion that the second equilibrium is robust to increasing information costs, while s_1 is not. For example, if $c(\{[-1, 0], (0, 1), [1, 2]\}) > \frac{1}{18}$ then player 1's best respond to 2's $\{[-1, 0], (0, 1), [1, 2]\}, (\alpha, \beta, \alpha)$ is no longer $\{[-1, 0], (0, 1), [1, 2]\}, (\beta, \alpha, \beta)$ but instead it is not to partition Ω at all and to always play β . It follows immediately that $t(s_1) \leq \frac{1}{18}$.²

When the underlying games, $g(\omega)$, exhibits multiple equilibria, the above discussion could suggest a selection criteria. That is, if players' behaviour is part of a *consistent plan* of how to play *all* games in a given set Ω , then the choice of equilibrium in the underlying game will be influenced by behaviour in the "neighbouring" games.

In the last example Γ_1 had a $\Gamma - Nash$ equilibrium that was implementable for any information-costs function. This happened because none of the players had to acquire any information in order to play his equilibrium strategy. This is a special feature of the meta-game considered in that example. Generally, this would not happen. That is, for each $\Gamma - Nash$ equilibrium, at least one of the players will have to acquire *some* information. Propositions 1 and 2 are only informative when the costs of acquiring information are very small. In most cases, for large enough costs of acquiring information, any equilibrium of Γ_1 will imply that players play differently to what they would have, if they had perfect information. A good example of how to calculate *all* equilibria of Γ_1 appears in chapter 4. This, however, is achieved at a cost; the analysis highly depends on the assumptions made regarding the exact structure of the information-costs function. If we do not commit ourselves to a particular structure, we can only partially characterize the problem of calculating the equilibria of Γ_1 for high costs. Given the strategy of his opponent, each player has to calculate his "nearest" best-response which is implemented

²It is not possible, in this example, to calculate the actual value $t(s_1)$ without making further assumptions about the costs of partitions of size 1.

at minimal costs. This, of course, depends on the actual payoffs in different states of the world. It is well worth noting that even though we are not able to prove any general results here, we are able, in chapter 4, to show some non-intuitive outcomes in the more complicated case where the state space is multi-dimensional.

3.4 Results for the Two-Stage Game

We take sequential equilibrium as the solution concept for Γ_2 . That is, we insist that players will play a continuation equilibrium in the second stage. Generally the set of outcomes supported by sequential equilibria is different from the set of outcomes supported by Nash equilibria of the same game with simultaneous move. In particular, outcomes differ when information costs become very small. In the one stage game all Nash equilibrium of the meta-game imply that players will play the Nash equilibrium of (almost) all of the underlying games. This is no longer true when we consider a two stage game. For some choices of Ω we find that all sequential equilibria are not $\Gamma - Nash$. In other cases the set of equilibrium outcomes is significantly larger than the set of all $\Gamma - Nash$. Generally, this will happen because behaviour in the first stage is observable and therefore strategical. The choice of information is also a signal, and as such it can be used as a co-ordination device or to signal the willingness to play a particular continuation equilibrium in the second stage.

We discuss two cases in which Ω will generate equilibrium outcomes that are not $\Gamma - Nash$, even when information costs are small. The first is simply the case where the set of sequential equilibrium outcomes (when information is cheap) is larger than the set of outcomes supported by $\Gamma - Nash$ strategies. This is possible because players are able to co-ordinate on playing non $\Gamma - Nash$ strategies while (credibly) threatening to move to a "bad" equilibrium in the second stage. The logic behind this argument is somewhat similar to that used in the "folk theorems" of repeated games. Player i can signal, by choosing a particular partition, that she intends to play a continuation equilibrium that

is not a $\Gamma - Nash$, but that yields high expected payoffs for both players. If Ω is such that she is able to punish her opponent, by choosing to play "bad" continuation equilibria when he chooses any other partition, then such an equilibrium of the meta-game becomes possible.

The following example (Example 2) illustrates the above claims:

Consider the following game with two strategies: c for co-operation, and d for non co-operation. Assume that the payoffs depend on the players' beliefs regarding the likelihood of playing similar games with the same opponent in the future. Formally, there is some symmetric metric relationship defined for all pairs of individuals, and payoffs depend on how "close" the two players are. We will assume that the best response for d is always non co-operation. The best response for co-operation, however, depends on the distance between the players. If they are close enough then best-response to c is c . Otherwise it is better to defect. Calculating the expected payoff each time the game is being played is costly, so a person might decide to once-and-for-all classify the set of all possible opponents as "friends" (*i.e.* those where c is best respond to c), and "strangers". The following table demonstrates payoffs for such a meta-game:

	c	d
c	$1 - \omega, 1 - \omega$	$0, 0.5$
d	$0.5, 0$	$0.25, 0.25$

Example 2. $\Omega = [0, 1]$

If $\omega > 0.5$, the strategy d strictly dominates c , and the game has a unique Nash equilibrium where both players play d . If $0.5 < \omega < 0.75$, then (c, c) Pareto-dominates the equilibrium outcome (d, d) . (that is, each of the games in this interval is a version of the "prisoners' dilemma".) If $\omega \leq 0.5$, the game has two symmetric Nash equilibria; (d, d)

and (c, c) where the latter Pareto-dominates the former. The worst (credible) punishment that players could inflict on each other is by always playing d . The expected payoff from always playing d is 0.25 or $0.25 - c(P)$ if the player chose to partition the set and not to use that information. Next, we look at all possible continuation equilibria for each pairs of partitions. Once partitions have been fixed, players can choose one of the following four strategies: (c, c) , (c, d) , (d, c) , and (d, d) . Expected payoffs are calculated by integrating the payoffs over Ω using the density function, ρ . For any pair of partitions there are exactly two continuation equilibria: one in which both players play (c, d) conditional on their signal, and another where both players play (d, d) . (the last result holds for any choice of ρ .) If only one of the players chooses not to partition Ω at all, then there is only one continuation equilibrium, (d, d) . If both players chose not to partition Ω , there are two continuation equilibria (c, c) and (d, d) .

The set of all sequential equilibria is large and complex but we are still able to identify all the outcomes of those equilibria. By choosing the appropriate pair of partitions we could obtain any pair (t_1, t_2) of expected payoffs where $t_1, t_2 \in (0.25, 0.53125)$. To see why, consider the expected payoffs for players when player 1 plays $\{[0, x], (x, 1)\}$, (c, d) , and 2 plays $\{[0, y], (y, 1)\}$, (c, d) (assume, without the loss of generality, that $0 < x < y < 1$):

$$\text{Player 1's expected payoff is: } \int_0^x (1 - \omega) d\omega + \int_x^y 0.5 d\omega + \int_y^1 0.25 d\omega = \frac{x^2}{2} - \frac{1}{4} - \frac{x}{2} + \frac{y}{4}$$

$$\text{and 2's expected payoff is: } \int_0^x (1 - \omega) d\omega + \int_x^y 0 d\omega + \int_y^1 0.25 d\omega = \frac{x^2}{2} - \frac{1}{4} - \frac{y}{4}$$

(t_1, t_2) is obtained by manipulating the values of x and y . The value 0.53125 is achieved if both players co-operate for $\omega \in (0, 0.75)$, and defect for all other values of ω .

We can now construct a sequential equilibrium for almost any such pair (t_1, t_2) where each player uses the (d, d) continuation equilibria as a threat if the other player does not co-operate. For example, players could adopt the following strategy: In the first stage

partition Ω to $\{[0, 0.75], (0.75, 1]\}$, and in the second stage play (c, d) if the other player chose *exactly* the same partition, otherwise play (d, d) . Notice that by adopting such a strategy each player guarantees himself an expected payoff of 0.53125 which is higher than in any of the equilibria where players' strategies are $\Gamma - Nash$. In particular we find that players will co-operate throughout the interval of "prisoners' dilemma"s. The exact range of values of (t_1, t_2) that are possible as outcomes of sequential equilibria of Γ_2 depends on the costs of partitioning the set to two. Let $c(P) = c_0$ when P splits Ω to two. If the splitting point is close to zero or to one, then the expected payoffs from playing the (c, d) continuation equilibrium become smaller than $0.25 + c_0$, hence players are better off not partitioning Ω at all. However, c_0 needs to be relatively large in order to rule out the equilibria where players split the set anywhere between 0.5 and 0.75, that is, the equilibria of the meta-game where players co-operate in the prisoners' dilemma.

To implement the strategy described above, players need to be very attentive to the choice of information made by their opponents. We now show that the level of "over co-operation" highly depends of that particular feature of the equilibrium strategy. To see this, we consider a permutation of the model where players are only able to tell whether the "splitting point" (that is, the exact value where Ω is partitioned) of their opponent falls within a (fixed) $\delta > 0$ from their own. In this perturbed model the outcome $(0.5312, 0.5312)$ is no longer possible as an equilibrium outcome. The best response to a strategy that switches from c to d at 0.75 is to switch between the same strategies at $0.75 - \delta$. The incentive to do that should be clear from the choice of Ω . Following similar logic, rational players will always have an incentive to move their "splitting point" closer to 0.5. The $\Gamma - Nash$ strategies will not be affected at all by this perturbation, while our other results from above will completely change, for *any* fixed value of $\delta > 0$.

The second case we discuss is a meta-game version of the well-known economic examples where information can "hurt" the player who obtained it. By not obtaining information the player can force his opponent to play a continuation equilibrium that

makes him better off than in any of the equilibria where he is informed. Whether or not that strategy (*i.e.* not acquiring any information and playing the continuation equilibrium) is a $\Gamma - Nash$ of Γ will not matter. In Γ_2 this choice of behaviour is perfectly rational.

Condition 1: There exists a strategy s_i for player i such that:

- (i) there exists a non empty set \mathcal{S} of sequential equilibria of Γ_2 where player i plays s_i ,
- (ii) none of the members in \mathcal{S} is a $\Gamma - Nash$, and
- (iii) player i is better off in any of the equilibria \mathcal{S} than in the outcomes of any of the $\Gamma - Nash$.

Proposition 3: If Ω satisfies Condition 2 then Γ_2 does not have a sequential equilibrium that is $\Gamma - Nash$.

Proof: Suppose on the contrary that there exists a sequential equilibrium of Γ_2 , p such that p is also $\Gamma - Nash$. Player i can profitably deviate from p by switching to strategy s_i . Condition 2 guarantees that such a deviation will be profitable and credible. \square

Remark: If we replace in condition 2 the word "any", in both instances, with "one of" then the meta-game Γ_2 will still have some sequential equilibria that are not $\Gamma - Nash$, but possibly also some that are. Whether or not a $\Gamma - Nash$ equilibrium will be selected, will depend on the choice of continuation equilibria for given first-stage strategies. If, however, we leave condition 2 as it is, we know that the phenomena is robust to choice of strategies in the second-stage. (as well as being independent from information costs.)

The following example illustrates the importance of Condition 2:

	L	M	R
T	$1 - 3\omega, -3\omega$	$-2, -1$	$-2 + 3\omega, 3\omega - 3$
B	$0, 3 - 3\omega$	$0, 2$	$0, 3\omega$

Table 3. $\Omega = [0, 1]$

In this third example, ω is uniformly distributed over $\Omega = [0, 1]$.

For $\omega \in [0, \frac{1}{3})$ the game has a unique pure strategy Nash Equilibrium - (T, L)

For $\omega \in (\frac{1}{3}, \frac{2}{3})$ the game has a unique pure strategy Nash Equilibrium - (B, M)

For $\omega \in (\frac{1}{3}, 1]$ the game has a unique pure strategy Nash Equilibrium - (T, R)

If the cost of a partition of size 3 is relatively small, Proposition 1 guarantees that the only Nash equilibrium³ of Γ_1 is obtained when player one plays $\{ \{ [0, \frac{1}{3}], (\frac{1}{3}, \frac{2}{3}), [\frac{2}{3}, 1] \}, (T, B, T) \}$, and two plays $\{ \{ [0, \frac{1}{3}], (\frac{1}{3}, \frac{2}{3}), [\frac{2}{3}, 1] \}, (L, M, R) \}$. This, however, is not true for the meta-game Γ_2 . Even if the cost of information gathering is equal to zero, players will never play their Γ - *Nash* strategy. To see this consider player 2's expected payoff from playing the Nash equilibrium of $g(\omega)$, in all possible states ω :

$$\int_0^{\frac{1}{3}} -3dx + \int_{\frac{1}{3}}^{\frac{2}{3}} 2dx + \int_{\frac{2}{3}}^1 (3x - 3)dx = \frac{1}{3}.$$

Player 2 can profitably deviate by choosing not to obtain any information, therefore forcing player 1 to play to the (unique) continuation equilibrium (B, M) which guarantees player 2 an expected payoff of 2.

In the equilibria of either of the two cases described in this section, players threat not to use in the second stage information they obtained in the first stage. This threats are credible, in contrast to the simultaneous move meta-game. It is not clear what would be the outcomes of a perturbed two-stage model where players are restricted to using

³modulo the behaviour in a, zero measure, sub-set.

their information. It is clear, however, that neither of the two behaviours presented in this section will survive as an equilibria of the perturbed model, while those sequential equilibria that are $\Gamma - Nash$ (if they exist) will not be affected.

3.5 Discussion

In this paper we presented a general approach to modelling endogenous acquisition of information when games are drawn from a continuum. The result can be generalized to a situations where Ω is any general set. Modelling the set where the game might be drawn from, rather than analysing the game for an exact structure of payoff, is appealing. Economists and game theorists have argued for some time that such a theoretical approach is needed, and findings from psychology seems to support the assumption that people do classify. If the meta-game is modelled as a one stage game, and if information costs are small, Proposition 1 says that players will always play as though they had perfect information. Proposition 2 says that any selection of Nash equilibria combined from the underlying games will constitute a Nash equilibrium of the meta-game when information costs are small. However, neither of the Propositions tells us anything about which of these equilibria is likely to be played. We show that the equilibria of the meta-game could be ranked with respect to their robustness to increasing information costs. In chapter 4 we argue that the equilibria that are more robust to increasing costs are also robust to other permutations of the model. If agents are uncertain about the exact cost of information, these equilibria are more "safe". (in the risk-dominance sense.) Economic equilibria are usually thought of as a steady state of a dynamic process. If we adopt this point of view, and if we think of the costs of information as changing over time, then these robust equilibria are more likely to be selected.

Once the costs of information becomes high players will no longer play their $\Gamma - Nash$ strategies. Instead, their optimal behaviour will be the solution of the following constraint maximization problem: Maximise expected payoffs subject to the costs of implementing

strategies. Results in these cases are highly dependent on the payoffs in $g(\omega)$ and the exact structure of the information-costs function.

In the two stage game we prove that sometimes players will not play a $\Gamma - Nash$ strategy even if information is (almost) free. This happens when they can either credibly threaten not to use their information, or when Γ contains a sub-set of games (with positive measure) exhibiting multiplicity of equilibria. However, we show that the existence of these "non efficient" equilibria could be a fragile feature of the two-stage model. To sustain these outcomes players must be very attentive to the choice of information of their opponents. The less attentive they become, the more their equilibrium outcomes resemble that of a $\Gamma - Nash$. On the other hand, a $\Gamma - Nash$ outcome is sustainable even in a one-stage game, where players are unable to pay *any* attention to what their opponents are doing. This raises some doubts about the appropriateness of the two stage model. We believe that the same problem is shared by many other models of bounded rationality, where the game is played in two stages. Generally in those models, players choose in the first stage how to allocate the scarce resource, and in the second they choose optimal strategies conditional on the resources available to them from the first stage. (as in Abreu and Rubinstein (1988), Dow (1993), and Rubinstein (1993).) In the second stage of their equilibrium strategies players always use their information about the choices made in the first stage by the other players. This is used either to "free ride", as in Abreu and Rubinstein (1988), or to make a credible threat, as in the model presented here.

One might ask whether a two-stage model is ever appropriate. To address this question we, once again, return to our interpretation of an economic equilibrium as a steady state of some dynamic process. We believe that the answer is positive in cases where the speed of adjusting strategies in the first stage is significantly slower than that of the adjusting strategies in the second stage. Think of the attention allocation mechanism as determined by evolution. A person is born with a strategy for the first stage, and very rarely even consider changing it. On the other hand, the underlying games are played

often enough for the second-stage strategies to be changed again and again. In such a scenario it is safe to assume that convergence to equilibrium in the second stage is fast. Therefore, any strategy profile that is a sequential equilibrium of the meta-game could be a reasonable outcome of the model. The actual meta-game equilibrium where the process rests will depend on the specifics of the dynamics and, even more, on the initial conditions of the system. However, if players chose strategies for the meta-game with, roughly, the same speed as they do for the second stage, it is unlikely that an equilibrium that is non Γ – *Nash* will be selected. (as long as the costs of acquiring information are small enough.) Many of the models of firms' R&D behaviour, entry, and information gathering fall into this category. In light of our results, it is therefore disappointing that most of the existing literature analyse such situations (as we show in chapter 4) as a two-stage game. Every time a two-stage model is used, we are bound to get strategic effects, and the modeller should be careful in deciding whether or not such behaviours are an adequate description of that particular economic reality.

References

- ABREU D. AND RUBINSTEIN A. (1988), "The Structure of Nash Equilibrium in Repeated Games with Finite Automata", *Econometrica* **56**, pp. 1259-1281.
- AKERLOF G. AND YELLEN J. (1985), "A near-rationality model of business cycle with wage and price inertia", *Quarterly Journal of Economics*, pp. 823-838.
- BECK A.T. AND EMERY G. (1985) *Anxiety disorders and phobias: A cognitive perspective*. New York: Basic Books.
- CARLSSON H. AND VAN DAMME E. (1993), "Global games and equilibrium selection", *Econometrica* **5**, pp. 989-1018.
- DOW, J. (1993), "Search Decisions with Limited Memory", *Review of Economic Studies* **58**, pp. 1-14.

- EYSENCK M. AND KEANE M. (1990), *Cognitive Psychology*. LEA Publishers.
- HARSANYI J.C. (1967/68), "Games with Incomplete Information Played by 'Bayesian' Players, Parts I, II, and III", *Management Science* **14**, pp. 159-182, 320-334, and 486-502.
- KREPS D. M. (1990), *Game Theory and Economic Modelling*. Oxford: Clarendon Press.
- LI CALZI M. (1995), "Fictitious Play by Cases", *Games and Economic Behaviour* **11**, pp. 64-89.
- MANKIW N.G. (1985). "Small menu costs and large business cycles: a macroeconomic model of monopoly", *Quarterly Journal of Economics*, pp. 529-537.
- RUBINSTEN A. (1988), "Similarity and Decision-Making under Risk", *Journal of Economic Theory* **46**, pp. 145-153.
- RUBINSTEIN A. (1993), "On price recognition and Computation Complexity in a Monopolistic Model", *Journal of Political Economy*, **31**, pp. 473-484.
- SCHELLING T. (1960), *The Strategy of Conflict*. Cambridge, MA.: Harvard University Press.

Chapter 4

Information Acquisition and Entry

4.1 Introduction

Consider the decision of a firm whether or not to enter a market. Especially in the case of a new market firms face uncertainty about the market conditions and other factors that affect the potential profits from entry. For example, firms might be uncertain about demand and costs of production. More specifically, the production costs may be considered to be the aggregate of several partial costs, such as the costs of labour, material, distribution and advertising. Similarly, demand for the product may be determined by factors as fashion, the weather and the state of the economy. Of course, the entry decision that *ex ante* maximizes expected payoff need not maximize the payoff *ex post*. That is, firms would behave differently if they had more information and would achieve higher expected payoffs. Firms thus have an incentive to acquire information, but the usual models do not allow for that possibility. In the real world, however, firms often do have the opportunity to gather information, for example by hiring the services of a consulting agency. Firms could then make their entry decision contingent on the information they receive from this agency.

Information acquisition has been studied extensively in the economics literature before. This literature has, however, two important shortcomings. First, it is assumed that there is uncertainty about one parameter only. As argued before there may be uncertainty about several parameters, especially in the context of an entry model. Aggregating

uncertainties about several parameters into a single one is artificial. In order to reduce uncertainty firms will have to acquire information about the underlying parameters (*e.g.* demand or cost). Those parameters should, therefore, be part of the model. If multidimensional uncertainty is considered, there is the possibility that firms "specialise" in obtaining information of a particular nature. That is, they can differentiate themselves endogenously. This particular feature of our model does not appear in models of one-dimensional uncertainty.

A second shortcoming of the existing literature is that it usually assumes that information cannot be acquired secretly. Information acquisition is modelled as a two stage game. When one firm engages in the acquisition of information in the first stage (say about demand), this activity is observed by the other firms. Although this does not mean that they can learn the content of the information (*i.e.* whether demand is low or high), they do learn the type of information. Subsequent decisions in the second stage may therefore depend on these observations. In our opinion, excluding the possibility of obtaining some information secretly (that is, without being observed) goes beyond reality.

In this paper we model information acquisition therefore as a one stage game. Firms face uncertainty about several stochastic variables that influence the profitability of the market. Firms can learn the outcome of each variable by investing some resources in research. Each firm has to decide about which variables it wants to be informed. After having obtained this information firms have to decide whether to enter or not.

These plausible assumptions yield intuitively appealing results. First, when costs of information gathering are small, entry decisions are *as if* firms had perfect information about the state of the world. This does not necessarily mean that all firms gather all information. It does imply, however, that lack of information cannot cause too much or too few entry. When information costs are higher inefficiencies of this type may occur.

Second, we show that when information costs are small, there are multiple equilibria that exhibit the "as if"-behaviour. In some of these equilibria firms specialize. That

is, each firm becomes informed about one variable only, with different firms becoming informed about different variables. In other equilibria all firms learn about all variables. These equilibria are socially undesirable because many firms do the same research and lots of time and money are wasted. The socially desirable equilibrium (the one where firms specialize) is robust to several variations and perturbations of the model whereas the undesirable equilibrium is not. The specialization equilibrium is feasible for a wider range of information cost parameters, complexity constraints cannot destroy it, and firms have no incentives to "free ride" on other firms' research efforts.

Although all these results are very intuitive, they do not hold when information acquisition is modelled as a two stage game. It is shown that even if information is very cheap, lack of information can cause inefficiencies. Moreover, in a two stage game firms can credibly threaten not to use the information. We also show that when information about one particular variable is cheap, it may happen that no firm will learn about it and that, consequently, inefficiencies may arise. This is caused by the multidimensionality of uncertainty. It is, therefore, important to model multidimensional uncertainty explicitly and one should not aggregate all uncertainty into one single economic indicator.

As mentioned before, endogenous information acquisition has received some attention in the past. Mostly it is modelled as a two stage game, where information is acquired in the first stage. Second period decisions can be conditioned on the type of information acquired by all firms. The only exception we are aware of is Matthews' (1984) auction game, where the participants can acquire some information about the object that is for sale before they submit their bids. The participants cannot observe whether the others are gathering information or not. Chang and Lee (1992), Hwang (1993, 1995), Li *et al.* (1987), Ockenfels (1989), Ponsard (1979) and Vives (1988) consider models of quantity competition with uncertainty about demand. Before quantities are chosen simultaneously in the second stage, there is a stage in which the firms can buy information in the form of obtaining a signal correlated with the true demand. The firms can choose the precision

of the signal, with higher precision requiring more expenditures.¹ It is assumed that at the beginning of the second stage the precisions chosen in the first stage are commonly known. That is, it is assumed that each firm observes the precision of information of his opponents before it makes its decision about quantity, and the quantity choice may therefore depend on these precisions. Daughety and Reinganum (1992) consider a timing game. Information acquisition can endogenously generate a signalling game. Milgrom (1981) considers auctions as in Matthews (1984). Here bids may depend on the number of participants that chose to become informed.

A model of entry with multidimensional uncertainty was also analysed by Fershtman and Kalai (1993). They study the behaviour of an incumbent firm that is active in several markets. In each of these markets there is uncertainty about demand. A firm can learn the true demand and, therefore, be able to respond to fluctuations in demand in some markets. The incumbent faces complexity constraints and is unable to be flexible in all markets. Learning about demand in at most k markets is free while learning about more markets is prohibitively costly. It is shown that this constraint forces firms to concentrate their attention on few markets, and that it can serve to deter entry.

The rest of the paper is organized in the following way. In the next section we introduce the formal model of (secret) information acquisition in an entry game. We state and prove the propositions about the limiting case when costs of information tend to zero. In section 3 we characterize the structure of equilibria when information costs are small. Section 4 gives a detailed example of two firms facing uncertainty about two variables. We compare our results with those of a two stage game in section 5. Section 6 concludes.

¹In Chang and Lee (1992), Hwang (1993, 1995) and Vives (1988) precision is chosen from a continuum. In Ockenfels (1989) and Ponsard (1979) the choice of precision is a binary one: Firms either learn the true demand perfectly or they do not learn anything. Li *et al.* (1987) considers both the continuum and discrete approximations of it.

4.2 The Model

There are n symmetric firms that have to decide whether to enter a market. The profitability of this market depends on the state of the world, ω , which in turn depends on N different variables x_1, \dots, x_N . Each variable x_i is the outcome of a discrete random variable X_i and can take two values, *Good* or *Bad*. We let $K = 2^N$ denote the number of states and denote by $\rho : \Omega \rightarrow (0, 1]$ the probability distribution over the states of the world induced by the variables X_i . We do not exclude the possibility that the random variables are correlated, but we impose that each state of the world occurs with positive probability. Now every state corresponds to an entry game $g(\omega) = (A, u^\omega)$, where $A_i = \{e, d\} \forall i$. If a firm chooses the action d (do not enter) he will receive a profit of zero. The profit of a firm that chooses e (enter) depends on the state of the world ω and on the number of firms that also enter. The profits increase when variables change from *Bad* to *Good*, and are decreasing in the number of firms that choose to enter. For each ω there exists a unique number $0 \leq k(\omega) \leq n$ such that if $k(\omega)$ or less firms enter they will make a strictly positive profit, and if more choose to enter they will make a loss. These assumptions imply that the pure equilibria of $g(\omega)$ are those where exactly $k(\omega)$ firms enter. Firms are symmetric, so any subset of $k(\omega)$ firms can enter.

The above payoff structure can be justified by assuming that after entry firms compete in quantities à la Cournot. Suppose that after entry, which may have taken a considerable amount of time (to build a plant for example), all uncertainty will be resolved. That is, the firms that entered know the state of the world and they know how many competitors there are. The payoffs of $g(\omega)$ represent the equilibrium payoffs of this Cournot competition.

A firm can learn for any of the variables whether the realization is *Good* or *Bad*. It will have to decide for which variables it wants to learn the realization. Suppose it chooses to be informed about the variables in some subset $I \subset \{X_1, \dots, X_N\}$. If two states ω and ω' differ only with respect to variables not included in I , then it has to choose the same action in those states. Therefore, a strategy for firm j in the meta game Γ is a pair (I_j, s_j) , where I_j denotes the set of variables to be informed about and where

$s_j : \Omega \rightarrow \{e, d\}$ is such that $s_j(\omega) = s_j(\omega')$ whenever ω and ω' differ only with respect to variables not included in I_j .

The payoff in the meta game is the expected payoff in $g(\omega)$ over all $\omega \in \Omega$ minus the cost of acquiring information, $c(I_j)$. We assume that $c(\emptyset) = 0$ and that $I \subset I'$ implies that $c(I') > c(I)$. That is, learning about more variables is more expensive. Information about one variable may be more expensive than information about another. Note that, since acquiring information is costly and is done secretly, it is a strictly dominated strategy to choose to learn about a variable but never use the information. In an equilibrium strictly dominated strategies are never used. It will be convenient to write s_j^* for the strategy (I_j, s_j) , where $s_j^*(\omega) = s_j(\omega)$ for all $\omega \in \Omega$, and where I_j is the smallest set that provides sufficient information to play s .

Of course, the equilibria of Γ will depend on the costs of information gathering. In particular, if costs are prohibitively large, no player will gather any information. Firms will either enter or not. In a pure equilibrium of Γ there will be some fixed number of firms that enter. However, in some states ω this number will be greater than the optimal number of entrants, $k(\omega)$. In other states ω' this number will be smaller than $k(\omega')$. Information costs then act as a barrier to entry. For intermediate values of the information cost, firms will not gather all information. Consequently too little or too much entry could occur also in this case. The next proposition, however, shows that this cannot happen if information costs are very small.

Proposition 1 Fix Ω and $g(\omega)$. Let Γ^c denote a meta game in which learning about all variables costs $c > 0$. Let s_c^* be a pure equilibrium of Γ^c . Then, if the limit exists, $\lim_{c \rightarrow 0} s_c^*(\omega)$ is an equilibrium of $g(\omega)$ for all $\omega \in \Omega$.

Proof. Suppose on the contrary that there exist a state ω' and a player who is not choosing the optimal action in this state in the limit. For small information costs this

remains true. This player can profitably deviate by choosing to learn about all variables (which is cheap) and then choose the optimal action in case the true state is ω' . This contradicts the presumption that s_c^* is an equilibrium of $g(\omega)$. \square

Proposition 1 says that if information costs are very small, in each of the states of the world the players will act *as if* the state of the world was known to all. This does not imply that all firms actually know the state of the world. In the following section we will show that this behaviour can be sustained even when firms gather very small amounts of information. The proposition only holds if information about *all* variables is cheap. One might think that if information about a certain variable is cheap, there will be at least one firm that will choose to learn about it. This intuition, however, turns out to be wrong as we show by means of an example in section 4.

Proposition 1 shows that an equilibrium of the meta game implies equilibrium in all entry games, if information costs are small. The opposite also holds, *i.e.* any combination of pure equilibria of the entry games constitutes an equilibrium of the meta game.

Proposition 2 For all $\omega \in \Omega$ let $s(\omega)$ denote a pure Nash equilibrium of $g(\omega)$. Let Γ^c denote a meta game in which learning about all variables costs $c > 0$. Then there exists an equilibrium s_c^* of Γ^c such that $\lim_{c \rightarrow 0} s_c^*(\omega) = s(\omega)$ for all $\omega \in \Omega$.

Proof. Simply define $s_c^*(\omega) = s(\omega)$ for all ω . That is, each firm j learns the minimal amount of information that allows it to play in each state ω the same action as $s_j(\omega)$. \square

Remark: The discussion thus far has been restricted to pure strategies. From the proof of proposition 1 it can easily be seen that it can be extended to mixed strategy equilibria. That is, any mixed strategy equilibrium of the meta game implies (in the limit) equilibrium in each entry game. It is not so clear whether proposition 2 can be generalized to mixed strategy equilibria as well. In any case, the proof of such a claim

should be considerably different than the proof for pure strategies. Namely, learning the minimal amount of information to be able to play a mixed equilibrium of the entry game will often mean that all information has to be acquired. (Because the probabilities used in the mixed strategy equilibria generally differ from state to state.) Notice that in a mixed strategy equilibrium of any entry game firms make zero profits. So learning all information and always playing the mixed equilibrium is dominated by always staying out.

4.3 Information Structures

Consider an entry game with n firms, and n variables, which are either *Good* or *Bad*. We assume that in each of the $\binom{n}{k}$ states where exactly k variables are *Good*, there is room for exactly k firms to enter. That is, in the pure equilibria of the corresponding game k firms enter. If costs of information are small, proposition 1 tells us that in a pure equilibrium of the meta game, exactly k firms will enter in these states. However, the proposition does not tell us anything about what information will be acquired. We now take a closer look at the information structures that might prevail.

There are many information structures that will support the behaviour predicted by proposition 1. Namely, let γ be a correspondence that assigns to each state ω a subset of firms F , such that the number of *Good* variables determining ω equals the cardinality of F .² Now γ corresponds to an equilibrium of the game in which firm i learns the minimal amount of information that allows it to enter in state ω if and only if $i \in \gamma(\omega)$. Learning about all variables will give sufficient information, but there may be cheaper ways to implement the above strategy.

Let us consider two “extreme” examples. First, consider the following strategy for firm i : Learn the outcome of all variables and enter if and only if at least i variables are

²Note that the number of such correspondences equals $\prod_{k=0}^n \binom{n}{k} \binom{n}{k}$.

Good. It is easy to verify that these strategies form an equilibrium of the meta game (which we will denote by s^{ineff}) when the costs of becoming informed are sufficiently low. Profits differ from firm to firm. Each firm has to incur the maximal information cost. Firm 1 enters always, except when all variables are *Bad*. Firm n , however, only enters when all variables are *Good*. To sustain this equilibrium the cost of acquiring all information must be less than the expected profit made by the firm that enters only if all variables are *Good*. This means that information costs must be very low.

Next, consider the following strategy for firm i : Learn the outcome of the i -th variable and enter if and only if it is *Good*. It is easy to verify that these strategies form an equilibrium, s^{eff} , when the information costs are low. Each firm enters in 2^{n-1} states of the world while only one variable has to be learned. On the other hand, each firm has to learn about at least one variable in order to have all firms enter when all variables are *Good* and no firm entering when all variables are *Bad*. So gross profits are relatively high while information costs are at a minimum.

We see that when information costs are small, many equilibria exist which differ with respect to the total amount of money spent on research. They also differ with respect to the degree of "shared" knowledge, that is the number of variables that are commonly learned by several firms. When we now increase the information costs gradually, the equilibria where all firms learn everything will disappear. In particular, s^{ineff} , the equilibrium in which one of the firms learns everything but enters only if all states are *Good*, is the first to disappear. On the other hand, s^{eff} , where each firm learns about one variable and enters in half of the states, will be the last to disappear. Hence, the equilibria where firms "specialize" in obtaining information about a particular variable are more robust to increases in information costs.

Above we assumed that firms can, in principle, decide to obtain as much information as possible by hiring more external market research companies. However, if a firm cannot use the services of external agencies and has to use its own research utilities, then the size of its research unit imposes an exogenous limit on the number of variables it can learn.

If that limit is smaller than n , then s^{ineff} is no longer an equilibrium of the meta-game. Of course, s^{eff} will remain an equilibrium as long as firms can do some research. This exogenous limit on the amount of information firms can get is similar to Fershtman and Kalai's (1993) model of bounded complexity. In their model, bounded complexity may lead firms to focus their efforts on some markets and to withdraw from others. In terms of our model of information costs, bounded complexity represents the case where learning about up to some fixed amount of variables is free, and learning about any additional variable is very expensive.

We have argued that s^{eff} is robust to increasing information costs and to complexity restrictions. We now show that it is robust against an additional perturbation of the model. Suppose that firms do not always choose their strategies simultaneously. A firm can wait and observe the behaviour of the other firms. Without being too formal we observe that s^{ineff} cannot be an equilibrium outcome of this game. This is because the firm that is only supposed to enter when all states are *Good* could "free ride" on the other firms. When it observes that some other firms did not enter, it can infer that not all states are *Good*. In this case it will decide to stay out without having to spend resources on research. Only in case all other firms entered, it may have to do some research. On the other hand, s^{eff} will be an equilibrium outcome of this perturbed game. Firms have no incentive to wait in this case because the information obtained by one firm is of no use to any of the other firms. We conclude that equilibria where all firms spend a lot of time and money on research are subject to the risk of "free riding" and therefore are unlikely to appear.

4.4 Two Firms

The discussion thus far has been restricted to the case where information costs are very small. In this section we remove this restriction. We assume that there are only two firms and that only the two (aggregate) variables cost and demand influence the profitability

of the market. For a particular specification of the payoffs we will analyse the equilibria in detail as a function of the cost of obtaining the different types of information. Here we model information acquisition as a one stage game, that is, entry decision cannot be made contingent on other firms' information gathering decisions. In the next section we will use the same example when we compare the one stage game with the two stage game.

Two firms have to decide whether to enter a market. The profitability of this market depends on two parameters, cost and demand, which are not known to the firms at the time they have to decide on entry. Both parameters can only take two values. Demand (cost) is high with probability one half. For convenience, let us assume that the two parameters are independent. The profit for each firm also depends on whether or not the other firm enters. The profits are given in Figure 1.

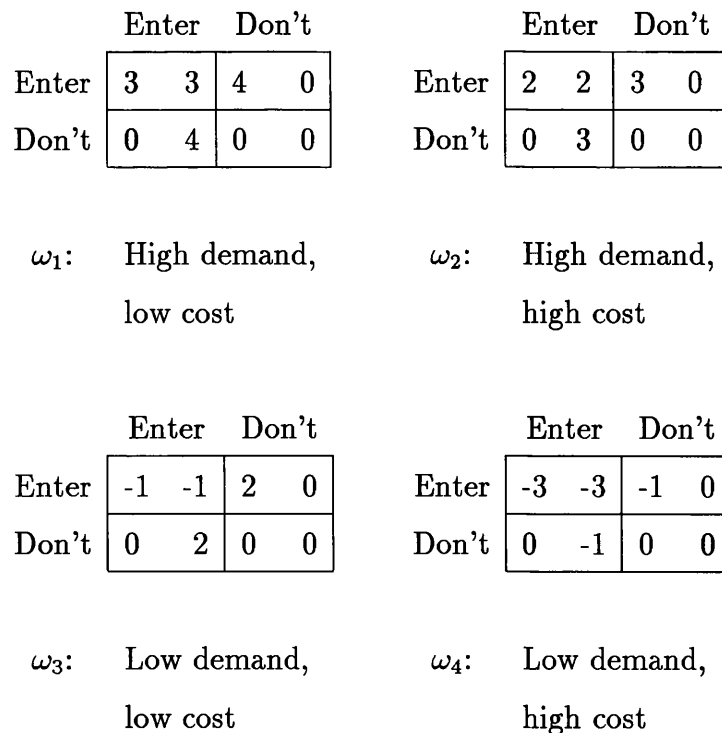


Figure 1.

If demand is high, there is room for both firms to operate profitably, independent of whether the cost is high or low. Even if the other firm decides to enter, it is optimal to enter. Profits are higher when costs are low. However, if demand is low, there is no room for both firms to enter. In fact, when costs are low this market is a natural monopoly and profits can be made only if one firm enters. In case the costs are high, the market is very bad and no profits can be made. In this case no firm wants to enter.

Learning the true value of cost (resp. demand) costs c (resp. d). We assume for convenience that finding out about both costs simply $c + d$. There is obviously a trade-off between becoming better informed and the cost of gathering the information by hiring consulting agencies. Each firm will have to decide first whether he wants to learn about cost, demand, both or nothing. Then it will have to decide whether to enter or not, conditional on the information received.

Since it is strictly dominated not to use all the information obtained, we simplify notation by letting $a_1 a_2 a_3 a_4$ denote the strategy where the firm learns the minimal amount of information that allows it then to choose a_i in state ω_i . For example, $eedd$ denotes the strategy “learn demand and enter if and only if demand is high.” Note that many of the remaining 16 strategies are strictly dominated.³ The analysis of the game can, therefore, be restricted to a 4×4 game. The strategies the players have in this game are:

- $eeee$ Gather no information and enter
- $eded$ Gather information about cost and enter if cost is low
- $eedd$ Gather information about demand and enter if demand is high
- $eedd$ Gather all information and enter unless demand is low and cost is high

Since firms are symmetric we write in Figure 2 only the payoffs for firm 1, the row player. The payoffs depend on the two information cost parameters c and d .

³Mostly, this is because some strategies are strictly dominated in the underlying games.

	<i>eeee</i>	<i>eded</i>	<i>eedd</i>	<i>eed</i>
<i>eeee</i>	$\frac{1}{4}$	1	$\frac{3}{2}$	$\frac{3}{4}$
<i>eded</i>	$\frac{1}{2} - c$	$\frac{1}{2} - c$	$\frac{5}{4} - c$	$\frac{1}{2} - c$
<i>eedd</i>	$\frac{5}{4} - d$	$\frac{3}{2} - d$	$\frac{5}{4} - d$	$\frac{5}{4} - d$
<i>eed</i>	$1 - c - d$	$\frac{5}{4} - c - d$	$\frac{7}{4} - c - d$	$1 - c - d$

Figure 2.

Figure 3 shows the types of equilibria that exist as a function of the costs of acquiring information.




Figure 3: Equilibria of the Meta Game

- A: $(eeee, eeee)$
- B: $(eeee, eded), (eded, eeee)$ and mixed equilibrium with support $\{eeee, eded\}$
- C: $(eeee, eded), (eded, eeee)$ and mixed equilibrium with support $\{eeee, eded, eeed\}$
- D: $(eeee, eedd), (eedd, eeee)$ and mixed equilibrium with support $\{eeee, eded, eeed\}$
- E: $(eeee, eedd), (eedd, eeee)$ and mixed equilibrium with support $\{eeee, eedd\}$
- F: $(eeee, eedd), (eedd, eeee)$ and mixed equilibrium with support $\{eeee, eedd, eeed\}$
- G: $(eedd, eedd), (eedd, eeed)$ and mixed equilibrium with support $\{eedd, eeed\}$

There are some interesting observations to make. First, apart from the equilibrium in area A, where information costs are so high that acquiring any information is not sensible, all pure equilibria are asymmetric: The two firms choose to obtain different information. The mixed equilibria are symmetric, but since firms randomize in these equilibria, it may occur that they end up doing different things. Note that the probabilities that are used in the mixed strategy equilibria may depend on c and d and therefore vary within each of the areas. Consider for example a mixed equilibrium in the interior of area G. In equilibrium firms choose $eedd$ with probability $(1 + 4c)/3$ and $eedd$ with probability $(2 - 4c)/3$. When firms use these strategies they will enter in states ω_1 and ω_2 and stay out in state ω_4 . In state ω_3 each firm will enter with probability $(1 + 4c)/3$. In the limit as c tends to zero, this probability goes to $1/3$, which is exactly the probability of entry in the mixed equilibrium of $g(\omega_3)$. (Compare proposition 1.) The mixed equilibrium outcome of $g(\omega_3)$ can thus be obtained as the limit of equilibria of the meta game as c goes to zero. Note, however, firms randomize between learning and not learning about costs. If they learn that costs are low they will enter for sure. Firms will never deliberately randomize between entering and not entering. (Compare proposition 2 and the remark following it.)

Second, we see from the above results that there is often too much entry, that is, more firms are entering than when there would be perfect information about cost and demand parameters. Also too few entry may occur. For example, in the pure equilibria in areas B and C only one firm enters if the true state is ω_2 . This illustrates that information costs may serve as a barrier to entry. When information about a certain variable is expensive, firms do not learn about it and play an "average" game instead. The payoffs in this game depend on the probabilities of *Good* and *Bad* states. If firms are "optimistic", *i.e.* they believe that the *Good* state will occur with high probability, then we will have too much entry in case this variable happens to be *Bad*. Similarly, if firms are "pessimistic" there exist states with too few entrants.

Finally, consider the part of area F where the costs for learning about production

costs are close to zero. In the pure equilibria, however, no firm will choose to learn about costs. The firm that learns about demand cannot gain from learning costs, given the fact that the other firm will always enter. The other firm could gain from learning about costs, because it could prevent entry in the worst state of the world, ω_4 . But then the firm also needs to know about demand, otherwise entry in the profitable state ω_2 is impossible. Learning about demand, however, is quite costly.

This example shows the importance of having uncertainty about more than one variable. If there was only uncertainty about cost, and it was very cheap to get informed, then (according to proposition 1) firms would choose to become informed such that they would act as if there was perfect information. Hence, at least one firm should learn the outcome of this variable. However, the presence of uncertainty about an additional variable (demand), which is quite costly to learn, destroys the argument and as a result nobody will get informed about cost. The vast majority of models so far investigated the case of one-dimensional uncertainty. Our example shows that results that hold in a world of one-dimensional uncertainty, might not hold in a world where uncertainty exists about more than one variable.

4.5 Observable Information Acquisition

Above we have considered the information acquisition and entry game as a one-stage game. Firms condition their entry decision only on their (privately obtained) information. As we remarked in the introduction, the existing literature on endogenous information acquisition assumes that information acquisition decisions are observed and considers therefore a two-stage game. In this section we examine the two-stage version of the game analysed above. In the first stage firms simultaneously decide about which variables to be informed. Before firms take their decision in the second stage, they are not only informed about the variables chosen, but they also know about which variables their competitors are informed. In this model, each firm can make its entry decision, therefore, contingent

on the other firm's information acquisition decision.

Each firm has four actions in the first stage: become informed about cost, demand, both or neither. We denote these actions by C , D , $C\&D$ and \emptyset , respectively. The second stage has therefore 16 starting points. We will insist that at each of these starting points a continuation equilibrium is played. Formally this means that we take sequential equilibrium as the relevant solution concept.⁴ For some choices of variables in the first stage there exists a unique continuation equilibrium. For example, in case neither firm learns anything entering is a dominant strategy in the second stage. For some other choices of variables, however, multiple continuation equilibria exist. For example, after the choice of $(C\&D, D)$, firms can continue in three ways: (1) firm 1 enters unless demand is low and cost is high, firm 2 enters if demand is high; (2) firm 1 enters if demand is high, firm 2 enters in any case; (3) both firms mix between the strategies used in (1) and (2). Note that the second equilibrium is somewhat peculiar, since firms do not seem to use their information completely. For instance, learning about demand and entering in any state is a dominated strategy in the one stage model. But here it is a credible threat in order to deter the other firm from becoming fully informed.

The diagram in Figure 4 lists all continuation equilibria and the corresponding payoffs (excluding the information costs).

⁴Since the move of Nature precedes the moves of the players, there are not many proper subgames. Subgame perfection does not have much bite here.

	<i>C&D</i>	<i>C</i>	<i>D</i>	\emptyset
<i>C&D</i>	eed,eedd 7/4,5/4 eedd,eed 5/4,7/4 mix 5/4,5/4		eed,eedd 7/4,5/4 eedd,eeee 5/4,6/4 mix 5/4,5/4	
<i>C</i>		eded,eeee 2/4,4/4 eeee,eded 4/4,2/4 mix 2/4,2/4		
<i>D</i>	eedd,eed 5/4,7/4 eeee,eedd 6/4,5/4 mix 5/4,5/4		eeee,eedd 6/4,5/4 eedd,eeee 5/4,6/4 mix 5/4,5/4	
\emptyset	eeee,eedd 6/4,5/4	eeee,eded 4/4,2/4	eeee,eedd 6/4,5/4	eeee,eeee 1/4,1/4

Figure 4

Now suppose that the costs of acquiring information are small but positive. Consider the strategy profile where firm 1 learns about both variables and where firm 2 learns about demand and where they then continue with playing (eed, eed) . This was the equilibrium outcome of the one stage game. But it is obvious that this is not an equilibrium outcome here. Firm 2 can deviate and choose (commit) to not obtain any information. In that case player 1 has to revise his second period strategy, because firm 2 will enter in any case. Therefore it will then choose to enter only if demand is high, and not use its information about costs.

From this diagram it can be easily read that, if the information costs are small but positive, there is essentially a unique pure strategy equilibrium outcome: One firm learns nothing, the other firm learns about demand. To sustain this as an equilibrium outcome, however, firms have to continue with the equilibrium $(eedd, eeee)$ in case the non-learning firm deviates and chooses to become completely informed.

This example clearly shows how perverse the effects can be of assuming that entry decisions can be made contingent on the information acquisition of other firms. Propositions 1 and 2 need not hold in such a context. Hence, even if the costs of acquiring information become very small, inefficiencies (in this case entry in non-profitable markets) arise because of the lack of information. Moreover, in the two-stage game firms can credibly threaten not to use their private information. In our opinion this demonstrates that information acquisition is better modelled by a one stage game than by a two stage game.

4.6 Conclusions

We analysed endogenous information acquisition in the setting of an entry game. In contrast to the existing literature and more in accordance with reality, we assumed that information is acquired secretly. Moreover, we allowed for uncertainty about more than one variable. If information is cheap, firms will buy all the information they need. Lack of information cannot be the cause of inefficiencies in the sense of too much or too few entry. Other types of inefficiencies may arise, however. Some equilibria exhibit socially undesirable amounts of time and money wasted on market research. We showed, however, that these inefficient equilibria are vulnerable with respect to increasing information costs, complexity constraints and free riding.

We find our results intuitive. What is striking, however, is that they do not hold under the assumptions usually made in the literature. For example, if information acquisition is assumed to be observable, firms may gather insufficient information. A firm that is known to be very well informed may namely provoke aggressive behaviour of the other entrants.

We have modelled multidimensional uncertainty explicitly. We find that representing all uncertainty by one single parameter is artificial. In order to reduce uncertainty firms have to gather data about real variables as cost and demand. One can hardly imagine

what it means to gather information about an abstract economic indicator. Still, one might think that analysing models of one-dimensional uncertainty yields also insights in models of multidimensional uncertainty. As we have shown, however, such partial analysis may yield qualitatively different results.

In this paper we have focused on entry games. We believe that the above considerations are valid for a wider range of information acquisition models. The plausibility of the assumptions used in these models can be questioned. The robustness of their results needs to be examined carefully.

References

- CHANG, C.-H., AND LEE C.-W. J. (1992). "Information Acquisition as Business Strategy," *Southern Economic Journal*, **58**, 750-761.
- DAUGHETY, A.F. AND REINGANUM, J.F. (1992). "Asymmetric Information Acquisition and Behaviour in Role Choice Models: An Endogenously Generated Signalling Game," mimeo, University of Iowa.
- FERSHTMAN, C. AND KALAI, E. (1993). "Complexity Considerations and Market Behaviour," *RAND Journal of Economics*, **24**, 224-235.
- HWANG, H. (1993). "Optimal Information Acquisition for Heterogeneous Duopoly Firms," *Journal of Economic Theory*, **59**, 385-402.
- HWANG, H. (1995). "Information Acquisition and Relative Efficiency of Competitive, Oligopoly, and Monopoly Markets," *International Economic Review*, **36**, 325-340.
- LI, L., MCKELVEY, R.D., AND PAGE, T. (1987). "Optimal Research for Cournot Oligopolists," *Journal of Economic Theory*, **42**, 140-166.
- MATTHEWS, S. (1984). "Information Acquisition in Discriminatory Auctions," in:

Bayesian Models in Economic Theory, eds. M. Boyer and R.E. Kihlstrom. Elsevier.

MILGROM, R. (1981). "Rational Expectations, Information Acquisition and Competitive Bidding," *Econometrica*, **49**, 921-944.

OCKENFELS, P. (1989). *Informationsbeschaffung auf homogenen Oligopolmaerkten: eine Spieltheoretische Analyse*. Physica-Verlag Heidelberg.

PONSSARD, J.-P., (1979). "The Strategic Role of Information on the Demand Function in an Oligopolistic Market," *Management Science*, **25**, 243-250.

VIVES, X. (1988). "Aggregation of Information in Large Cournot Markets," *Econometrica*, **56**, 851-876.

Chapter 5

An Economist's Perspective on Probability Matching

5.1 Introduction

The growing popularity of experimental economics provides a new dimension to the understanding of how people learn. In a typical experiment subjects repeatedly play the same game with varying, unknown opponents. Often the environment is such that subjects neglects the effect of their own behaviour on that of the group's. (for a survey of the methodology used by experimental economists see Crawford 1995.) Similar experiments were common in psychology during the 1950s and 1960s. In their settings subjects repeatedly played a game with a random device. At each stage the device would chose one strategy according to some fixed probabilities. The subjects' task was to try and predict that strategy. He or she were informed about the outcome at the end of each trial. In some of the experiments subjects would additionally receive rewards and pay penalties. Subjects' predictions had no effect on the choice of strategy in the next trial. Formally, denote by $A = \{A_1, \dots, A_n\}$ the set of (pure) strategies, and by $\mathcal{P} = \{p_1, \dots, p_n\}$ the probabilities with which those strategies are selected, where $\sum p_i = 1$. \mathcal{P} is independent of the history of outcomes and of the behaviour of subjects.

Psychologists reported that subjects often used a behaviour they called *probability matching*. By matching they meant that subjects asymptotically proportioned their

choices according to \mathcal{P} . Denote by $Q_i^k = \{q_1^{k,l}, \dots, q_n^{k,l}\}$ the frequencies of choices recorded in the last k out of l trials, then for a fixed k (typically between 20 to 80) subjects are using probability matching behaviour if $Q_i^k \rightarrow \mathcal{P}$ as $l \rightarrow \infty$. That is, if the game is repeated often enough then in the last block of trials subjects will choose strategy i with probability approximately equal to p_i . As most of the literature is concerned with binary choice decisions it is useful to refer to the options as *Left* and *Right*, and to the fixed probabilities as P_L and P_R (where $P_L + P_R = 1$).

Matching suggests that subjects learned the probabilities. But once those probabilities had been learnt the strategy that maximizes expected utility is either always choosing *Left*, if $P_L > 0.5$, or otherwise always choosing *Right*. It is therefore not surprising that probability matching is often raised by psychologists in discussions about the psychological foundations in economic modeling. One famous economist (Arrow 1958, p.14) wrote:

We have here an experimental situation which is essentially of an economic nature in the sense of seeking to achieve a maximum of expected reward, and yet the individual does not in fact, at any point, even in a limit, reach the optimal behavior. I suggest that this result points out strongly the importance of learning theory, not only in the greater understanding of the dynamics of economic behavior, but even in suggesting that equilibria maybe be different from those that we have predicted in our usual theory.

Probability matching was also reported in experiments with rats and pigeons, raising an additional problem, this time to theoretical biologists. In biology it is assumed that evolution selects only behaviours that maximize access to critical resources, such as food.

On the other hand, asymptotically matching the probabilities is a behaviour predicted by almost all of the stochastic learning theories, such as Estes' stimulus sampling, Suppes' one-element sampling, and Bush & Mosteller linear learning model. Theoretical economists have been interested in such learning theories starting with Cross (1973), and more recently game theorists, such as Borgers and Sarin (1993), studying the replicator

dynamics are closely related to such models of learning. It is this interest together with the growing popularity of experimental economics that make it important to look back at the parallel work of psychologists. Our goal is to try and clarify what "probability matching" and "optimal behaviour" mean within such experimental context. We review many of the relevant learning models and describe their asymptotics. We survey most of the known results from the 1950s and 1960s. Probability matching was actually reported as early as in 1939, in Brunswik's experiments with rats, and in Humphreys' human verbal conditioning experiments. Unfortunately, it is hard to evaluate these early experiments as many details about the apparatus and experimental design are missing. We take advantage of several earlier surveys of much of the same literature: pp. 184-185 of Edwards (1956), Fiorina (1971), Luce & Suppes (1965), and Brackbill & Bravos (1962). This paper provides a more complete survey, and one that is focus on the matching phenomena. We also do our best to be unbiased in favor of any of the theories. (Edwards, Fiorina, and Brackbill & Bravos also promote their own theories.)

The rest of the paper is organized in the following way: In section 2 we review the theoretical backgrounds for decision theories, and we try to clarify what exactly they predict in repeated, binary choice experiments. Section 3 is a survey of most of the known experiments relevant to our discussion. The five tables in section 3 constitute an important part of the paper. In section 4 we discuss the results and their relationship to the theories of section 2. In section 5 we describe some other behaviours that we believe to be irrelevant to our discussion, but are also known as probability matching. Section 6 concludes.

5.2 Theoretical Background

When psychologists refer to optimal behaviour, they typically have in mind the following naive maximization theory: Denote by p^* the rate of *Left* choices in the limit (and $1 - p^*$ for *Right*). Maximizing the number of correct predictions is mathematically equivalent

to choosing p^* in order to maximize the expression $p^*P_L + (1 - p^*)P_R$. This implies either $p^* = 0$, or $p^* = 1$. This is naive in the sense that it neglects the learning process and any effects it might have on the asymptotic behaviour. Since S's are never informed about the actual values of P_R and P_L , they first have to learn them, the outcome of such a learning process and its length depends on the actual outcomes of the trials. It is important to note that there are theories that describe how rational agents learn and behave in such settings. These are known as the solutions to the *two-armed bandit* problem¹, and they are based on classical Bayesian decision theory (Savage (1972)). The two-armed bandit problem received much attention from both statisticians and economists (Rothschild (1974), Gittins (1989), or for a discussion of a more general case see Banks & Sundaram (1992)). The main idea can be seen from the following intuition; A short period of experimenting is less costly than a longer one which reduces the probability of making the wrong asymptotic choice. Solution for the two-armed bandit problem therefore involves choosing the optimal number of experimental trials. The set of optimal strategies is identified by *dynamic allocation indices*, or *Gittins indices*. These indices depend on the subjects' current beliefs regarding the values of P_i 's (Gittins (1989)). An important property of this rational behaviour model is that, unlike in the naive version mentioned earlier, the decision maker could, with positive probability, choose the "wrong" option forever (Rothschild 1974).

In the models described thus far, subjects are assumed to be able to keep some sort of statistics and to act according to it. A weaker assumption is made by stochastic learning theories. Here the decision maker is characterized in every given moment in time by a probability distribution over the set of possible strategies. This distribution is by the reinforcements received at the end of each trial. The theory makes predictions about subjects' transitory behaviour and, if the distribution converges as time goes to infinity, about the limiting behaviour. This is calculated by solving the following equation:

¹A two-armed bandit is a slot machine with two arms, each operating with a fixed probability.

$$E[P_L(n+1)|P_L(n)] = E[P_L(n)]^2$$

It is a feature of all stochastic learning models that under certain conditions they predict probability matching behaviour in the limit.³ The models differ with respect to the exact specifications of the transition functions and the state spaces. We start with Bush and Mosteller's (1955) linear learning model. The state space in this model is the continuum of all mix strategies. Assuming that an individual *a priori* prefer choosing *Left* (alt. *Right*) when left (alt. right) is chosen by the random device, then the transition function of the model can be described in the following way:

$$P_L(n+1) = (1 - \theta_1)P_L(n) + \theta_1 \text{ with probability } P_L, \text{ or}$$

$$P_L(n+1) = (1 - \theta_2)P_L(n) \text{ with probability } 1 - P_L$$

Solving for $E[P_L(n+1)|P_L(n)] = E[P_L(n)]$ we get:

$$P_L(\infty) = \frac{P_L}{P_L + (1 - P_L)\theta_2/\theta_1}$$

where θ_1 and θ_2 are constants determined by the strengths of reinforcements. The limit behaviour depends on the ratio $\frac{\theta_2}{\theta_1}$. In particular if that ratio is close to 1, the model predicts matching.

In Suppes' *one-element stimulus sampling* model (Suppes & Atkinson (1960)) only pure strategies are considered. (*i.e.* there are only two states in the state space.) If the subject's choice coincided with that of the random device in the n th trial, then her distribution will not change in the $n+1$ th trial. Otherwise she will switch strategies in the next

²This is not true in general. Borgers and Sarin (1993) discuss the conditions for convergence in the continuous time case. It is true for a discrete time version if the state space contains only two states, as in, for example, Suppes (1960) or Suppes and Atkinson (1960).

³The actual conditions under which the model predicts matching vary from one model to the other.

period with a fixed probability ϵ_i ($i=1,2$). If we further assume that in these cases she receives no positive or negative reinforcement, then the model's asymptotics are given by:

$$P_L(\infty) = \frac{P_L}{P_L + (1 - P_L)\epsilon_1/\epsilon_2}$$

Chapter 11 of Suppes & Atkinson (1960) describes several other models with transition function defined over the two state space.

Suppes' (1961) *stochastic expected utility* model is an effort to combine some of the features of utility theory and stochastic learning. Preferences satisfy some of the requirements of utility theory, but they also change over time. Though the asymptotic behaviour of Suppes model are generally unclear, he shows that in a binary choice case, under certain conditions, the model predicts matching in the limit.

Closely related to the principle of maximizing expected utility are Edwards' *Relative Expected Loss Minimization (RELM)* rule (Edwards 1961, 1962) and Simon's principle of *Minimal Regret (MR)* (Simon 1976). Also similar are their predictions for the asymptotic behaviour.⁴ Subjects will choose in the limit with probability 1 the more frequently rewarded option, except in the case where the reward for the less frequent option is significantly greater than the reward for the frequent side. In that case both theories predict that subjects will always choose the more profitable side. As with maximizing expected utility, predictions are made only for the asymptotic behaviour.

Very different are the predictions of some psychologists studying probability learning with pigeons (like Herrnstein (1961, 1974), and Baum (1974)). They predict that in the limit subjects will *always* match. This is different from the predictions of stochastic learning theories where behaviour in the limit is dependent on several parameters. Herrnstein's main idea is that matching the underlying probabilities is the main force operating on the decision maker in a repeated choice situation. The experimental setting here is

⁴Instead of maximizing some expression with parameters taken from the payoff matrix, both models use rules that minimize expressions with parameters taken from the regret matrix (Savage 1972).

different not only with respect to the choice of subjects. For this reason we postpone the discussion of these experiments to section 5.

Finally, there have been several theoretical attempts to explain probability matching as a sort of optimal behaviour. Some experimenters noted that their subjects get bored with always choosing the same option, therefore switching between guessing *Right* and *Left* throughout the trials. Somewhat different is Brackbill & Bravos's (1962) model where S's receive a greater utility by guessing correctly the outcome of the less frequent option. In both models utility-maximizing individuals will not choose one strategy with probability 1 in the limit. Another interesting explanation often suggested (though I could not find any formal model based on this) is that subjects believe in the existence of some sort of regularities, or patterns, in the sequence of outcomes. Such a belief causes them to disregard their own experience and to keep looking for rules and patterns. If there *is* a pattern, it is worth spending a long time trying to find it, because once it is found one can get 100 per cent of the rewards. This idea let Fiorina (1971) to conclude that the gambler's fallacy may not be a fallacy after all. Like in the two previous models, the specified behaviour will prevent subjects from choosing one strategy in the limit.

5.3 Experimental Results

Subjects: In most of the experiments S's are undergraduate psychology students. In some experiments they are undergraduate students recruited throughout the college. Neimark (1956) and Edwards (1956, 1961) used basic airmen trainees as subjects. Derks and Paclisanu (1967), Brackbill et al. (1962) and Brackbill and Bravos' (1962) subjects were childrens (exact details of their age appear in tables 4 and 5). Many, like Estes used rats and pigeons. A separate discussion of the experiments with animals appears in section 5.

Apparatus: In most experiments S's sat in a small room (or a separate area within a bigger room) facing a box with two lights, one at each side. At the end of each

trial one of the lights would illuminate. Otherwise, pre-prepared multiple choice answer sheets were used (as in Edwards 1961). These sheets are covered except for two holes in each line. One is where subjects write their choices, and the other (which is covered in the beginning of each trial) lists one strategy. By removing the cork from the second hole subjects learn if they were correct. The sheets are prepared in advance according to fixed probabilities. In Mores and Randquist (1960) subjects collectively observed a random event, individually predicting its outcome.

Instructions: Subjects were instructed to try and correctly predict the strategy chosen as many times as possible. In many of the experiments they were told that their actions will not effect the choice of strategies in the next trials (this is obvious in the case of pre-prepared sheets). In some cases the instructions mention the word "probabilities". In all of the surveyed experiments subjects did not know in advance that the probabilities are fixed. Experiments where they were told that (like Tversky and Edwards 1961) are excluded from our discussion.

Experimental Design: A fixed groups sizes and number of trials (see tables for specifics). Each group of subjects were supervised by a experimenter. Subjects could only observe the outcomes of their own trials. Note also that if $P_L \approx 0.5$ any asymptotic behaviour is optimal. We therefore exclude from the tables experiments where $0.4 \leq P_L \leq 0.6$.

Tables of Results:

Tables 1 and 2 are adopted from Luce and Suppes (1965). These tables contain a summary of results in repeated, binary choice experiments. The first table summarizes results from experiments where subjects did not receive any payoff . They were still informed about the outcome of the trial. Table 2 lists the results of those experiments with monetary payoffs. The payoffs, in cents, appear in the fourth column where the left most number describes the payoff for making a correct guess, and the right number is the payoff in all other cases. Edwards' (1956) third experiment, is the only exception where asymmetric payoffs were used; Subjects received 12 cents for correctly predicting the right light, 4 cents for correctly predicting the left light, and -2 cents otherwise.

We use π to denote the (fixed) probability of strategy *Left* (P_L).

Estimations of the limit distribution of strategies are surveyed. In the case of a binary choice experiment, this distribution is one dimensional. In the first two tables the estimates of π_∞ are obtained by taking the group's average frequency of choosing *Left* over the last block of trials. In Edwards (1961) he specifies the results for the individual asymptotic behaviour which we include in **table 3** (also the averages of each of these groups appear in table 1). All comments in the first two tables are ours. Further discussion appears in section 4.

Table 4 summarizes the results of Brackbill, Kappy and Starr (1962), and Barckbill & Bravos (1962). The frequencies of choosing the most profitable outcome (*Left*, as $\pi=0.75$) in the n th trial are given as a function of the outcome of the $n-1$ th trial. Subjects here are typically younger than those in the experiments reported before.

Table 5 is taken from Derks and Paclisanu (1967). Their goal was to investigate the relationship between matching-optimization ratio and age. This is a part of a more general research investigating the relationships between cognitive development and decision making. 200 trials were used with $P_L = \pi = 0.75$, for all groups.

The following notes refer to Tables 1 and 2: ¹Trails do not seem to be independent.

²Behaviour changing in the last block of trials. Estimate of P_∞ may be misleading.

Experimenter	Size	Trials	π	Estimate of π_∞
Grant, et al.(51) ¹	37	60	0.25	0.15
			0.75	0.85
Jarvik (51) ¹	29	87	0.60	0.65
			0.67	0.70
			0.75	0.80
Hake and Hyman (53)	10	240	0.75	0.80
Burke, et al.(54)	72	120	0.90	0.87
Estes, Straughan (54) ²	16	240	0.30	0.25
		120	0.85	0.87
Gardner (57) ¹	24	450	0.60	0.62
			0.70	0.72
Engler (58)	20	120	0.25	0.29
			0.75	0.71
Neimark and Shuford (59) ²	36	100	0.67	0.63
Rubinstein (59)	37		0.67	0.78
Anderson, Whalen (60) ¹	18	300	0.65	0.67
			0.80	0.82
Suppes, Atkinson (60)	30	240	0.60	0.59
Edwards (61)	10	1000	0.30	0.11
			0.40	0.31
			0.60	0.70
			0.70	0.83
Myers et al. (63)	20	400	0.60	0.62
			0.70	0.75
			0.80	0.87
Friedman et al. (64)	80	288	0.80	0.81

Table 1

Experimenter	Size	Trials	Payoffs	π	π_{∞}
Goodnow (55) ¹	14	120	(1, -1)	0.70	0.82
				0.90	0.99
Edwards (56)	24	150	(10, -5)	0.30	0.19
				0.80	0.96
Edwards (56)	6	150	(4, -2)	0.70	0.85
				0.80	0.96
Edwards (56)	6	150	(4 or 12, -2)	0.70	0.46
				0.90	0.95
Galanter, Smith (58) ¹	30	200	unclear Higher	0.75	0.78
				0.75	0.90
Nicks (59)	144	380	(1, 0)	0.67	0.71
	72			0.75	0.79
Siegel, Goldstein (59)	4	300	(0, 0)	0.75	0.75
			(5, 0)	0.75	0.86
			(5, -5)	0.75	0.95
Suppes, Atkinson (60)	24	60	(1, 0)	0.60	0.63
	30	240	(5, -5)	0.60	0.64
			(10, -10)	0.60	0.69
Sigel (61)	20	300	(5, -5)	0.65	0.75
				0.75	0.93
Myers et al. (63)	20	400	(1, -1)	0.60	0.65
				0.70	0.87
				0.80	0.93
			(10, -10)	0.60	0.71
				0.70	0.87
				0.80	0.95

Table 2

$\pi = 0.7$	$\pi = 0.6$	$\pi = 0.4$	$\pi = 0.3$
100	91	49	26
100	90	48	26
97	85	47	20
96	81	46	20
95	77	43	17
93	76	43	13
91	74	43	13
88	74	40	13
88	71	35	13
87	70	31	12
85	69	31	11
85	66	29	11
80	64	26	8
80	64	22	8
75	63	21	4
70	61	20	4
65	61	16	0
60	59	15	0
58	56	11	0
56	46	0	0
$\mu = 83$	$\mu = 70$	$\mu = 0.31$	$\mu = 0.11$

Table 3. Percentage of *Left* choices in the last 80 trials (out of 1000) for each subject. Each column contains the results for one of four groups. Each group faced different probability. For each group the results are ordered. For example, in the 0.7 group two subjects chose *Left* in all of the last 80 trials. Five (out of 20) chose *Left* 70 percent of the time or less.

Reward	N	Grade	Trials analyzed	MM	LM	ML	LL
None	4	5	321-400	.77	.71	.56	.68
None	12	3	101-200	.74	.82	.38	.56
1 M or L	12	3	101-200	.80	.83	.65	.85
3 M or L	12	3	101-200	.82	.87	.67	.73
5 M or L	12	3	101-200	.89	.87	.64	.78
1M: 4L	10	4	121-200	.76	.79	.50	.75
1M: 3L	10	4	121-200	.83	.70	.68	.76
2M: 3L	10	4	121-200	.80	.84	.58	.85
1M: 4L	10	12	121-200	.74	.57	.71	.61
1M: 3L	10	12	121-200	.72	.62	.63	.76
2M: 3L	10	12	121-200	.84	.88	.70	.72

Table 4. The left most column describes the ratio between the two rewards: for correctly predicting M (the most frequent event, with $\pi = 0.75$), and L (least frequent, probability 0.25). The number of subjects in each group is given by N , followed by their school grade. The frequencies of choosing the most frequent strategy in the n th trial, given a prediction and the outcome in the $n-1$ trial, are in the appropriate column - if the prediction was M and the outcome L, the appropriate column is ML.

Group	Over-match	PM	Under Match	Total
Nursery	22	3	4	29
Kindergarden	5	3	21	29
First Grade	5	5	10	20
Second Grade	4	8	8	20
Third Grade	3	10	7	20
Fifth Grade	2	13	5	20
Seventh Grade	2	13	3	20
College	4	13	3	20

Table 5. Frequency of S's in choice behaviour categories. Behaviour is measured over the last 100 (out of 200) trials. Payoffs are either candies, toys or small sums of money, depending on S's age.

In Morse and Rundquist (1960) 16 S's were instructed to guess whether a small rod dropped to the floor would intersect with a crack in the floor. The probability of such an event was clearly less than a half. The same subjects went through a standard light guessing experiment, where each side was armed according to their behaviour in the first experiments. Each subject faced a sequence generated by their own behaviour. In the first stage Morse and Rundquist reported that 5 subjects adopt a "maximizing" strategy, and the group average was much higher than predicted by probability matching. Matching was observed, however, in the second stage. Subjects reported that they have been trying to find patterns in the trials of the second experiment.

5.4 Discussion

Quality of Experiments: In some of the experiments surveyed here (see footnote 1 in tables 1 and 2), and in many more that have been left out, the sequence of outcomes did not specify the non-contingency condition. This is partially because of the technology that was available in those years for generating random sequences and partially because some experimenters did not appreciate the importance of this condition. Often randomization took place within small blocks, *e.g.* in every block of ten trials, 7 were *Lefts* and 3 *Rights*. Also common was to exclude from the experiment three or more (sometimes four or more) consecutive *Lefts* or *Rights*. In these cases it is only reasonable to expect that subjects will notice the contingencies. Furthermore, it is optimal to sometimes guess the less frequent option. However, some economists, like Fiorina (1971), went too far in suggesting that the whole literature should be disregarded. There are many experiments with independent trials where matching behaviour was reported.

Comparing the PM hypothesis (PMH) with that of Extreme behaviour: It is important to note probability matching as an asymptotic prediction *a priori* is better than the that of extreme behaviour for several reasons: First, if behaviour still changes in the last block of trials, estimates of π_∞ are more likely to support the matching hypothesis. This was

even stronger in the early experiments of probability learning, such as Estes (1957), where the estimator was obtained by taking the average over *all* trials. Second, the theory of two-armed bandit tells us that utility maximizing subjects might be always choosing the wrong arm, hence taking them into considerations (by taking the group's average) will give results that are biased towards the matching. In fact averaging is only appropriate if the distribution of the individual results is approximately binomial with most of the mass concentrated around the mean. Finally, consider Edwards' (1961) comment on p.392:

Obtaining an estimate for π_∞ and testing the null hypothesis that that estimate is not significantly different from π is widespread in the probability learning literature. Such a procedure constitutes attempting to prove a null hypothesis; the smaller the amount of data or the greater its variability, the more likely it is that such a procedure will "confirm" PMH. This is why the small but consistent disagreements with PMH relevant by most probability learning experiments have not been noticed.

Discussion of results: Results in Tables 1 and 2 mostly support the matching hypothesis, but are not conclusive. When monetary payoffs are introduced (in table 2), π_∞ mostly exceeds π . Generally matching decreased with size of the reward (Edwards (1956), Sigel & Goldstein (1959), Suppes & Atkinson (1960, chapter 10), Atkinson (1962), Myers et al. (1963), and Brackbill, Kappy & Srarr (1962) lines 2-6 in table 4), though the effect is small.

Individual results in Table 3 show that very few S's (7 of 80) always chose *Left*. None of the subjects asymptotically choose the less frequent alternative, as the two armed bandit theory predicts. The large variability of the data suggests that looking at groups' averages may not be appropriate. Notice that only 16 (of 80) subjects choose *Left* with probability smaller or equal to π . This could be related to the large number of trials used in the underlying experiment. Generally it seems that matching is decreasing with the number of trials⁵.

⁵Note that this is a cross-experiments observation and therefore, could be misleading. I could not

Results summarized in Table 4 show that subjects' behaviour is (partially) contingent on what happens in the last trial, as predicted by stochastic learning theories. The results show that the best indicator for behaviour is the outcome of the last trial and not the difference between the prediction and the outcome. This suggests that the structure of reinforcement might be different from the intuitive one.

Table 5 relates probability matching to cognitive development. It shows an inverse relationship between optimal behaviour in a repeated binary choice situation and age. Derks and Paclisanu (1967) relate their findings to developmental changes in human learning. Young children use a very simple maximization strategy. The young children who under matched π , typically choose *Right* throughout the trials. Around the ages of 5-7, children develop skills of learning by associating events and outcomes, and sometimes matched π .

Additional experiments (not in tables) showed that the matching hypothesis does poorly in predicting subjects' behaviour in decision problem with three or more strategies. Instead, subjects tend to choose the most frequent option with asymptotic probability higher than the one determining its success (Gardner 1957, 1958, Cotton & Rechtschaffen, 1958, McCormack 1959). This is still consistent with some stochastic learning models, like Bush and Mosteller's linear learning model. Generally, this supports what we already see in the tables: Probability matching as a prediction of asymptotic behaviour is very unlikely, but any theory that wants to predict asymptotic behaviour should include it.

Finally, the results seems to indicate that S's often condition their behaviour on outcomes of the last trials. It is therefore interesting to study the predictions of theoretical models where such strategies are considered. Subjects in such models should be able to test hypothesis about patterns and repetitions in the sequence of outcomes. Such models, however, would easily become complicated and would not easily lend themselves for empirical testing.

find any experiments where this was controlled.

5.5 Other Matching Behaviours

In this section we briefly describe some of the other behaviours associated in the literature as with probability matching. These experiments use different settings. The findings do not seem to be in conflict with the assumption that animals maximize their access to food. Instead the observed behaviour could be seen as a special case of optimal behaviour.

Experiments with Groups

S's are groups of fish in a tank. Food is being offered from both ends of the tank. The food at one side is given twice as often as food at the other side. After a few seconds a pattern was formed where 2/3 of the fish located themselves in the more frequently rewarded side, while the remaining 1/3 went to the other side. A similar experiment was conducted with ducks. Once more, members of the group proportioned themselves according to the ratio of food supply. S's behaviour here is optimal in the sense that it constitute a Nash equilibrium (in pure strategies) of the $2 \times 2 \times \dots \times 2$ game where each individual choose a side.

Fretwell and Lucas used somewhat similar idea to explain the hunting behaviour of great tits: Observations showed that the hunters would choose a certain hunting area with probability equal to that of succeeding in finding prey there. They showed that the hunters PM behavior is compatible with the mixed strategy Nash equilibrium of the game where hunters' success is determined by the presence of prey and the number of other hunters present. Moreover, they showed that in that game, the mixed Nash equilibrium is the only evolutionary stable strategy.

Experiments with Pigeons

In this famous set of experiments S's are hungry pigeons, at 80% of their normal body weight, located in a box with keys on both side. On each side a VI schedule⁶ is controlling

⁶A concurrent VI schedule is a program that generates a sequence of random time intervals with a given mean. It is widely used tool in operational psychology. Using a VI-24 schedule for operating a

the rate with which food appears, if the pigeon is pecking on that key. The frequencies of food appearing on a given side are determined by the average delay time of the VI schedules. Herrnstein (1960) compared the recorded number of pecks on each key with the ratio of the mean delay times of the two schedules. He formulated his findings in his "direct matching law": Denote by R_i the reinforcement rate determined by schedule i , and by P_i subject's rate of respond to choice i , then:

$$\frac{P_1}{P_1+P_2} \approx \frac{R_1}{R_1+R_2}$$

or the algebraically equivalent condition: $\frac{P_1}{P_2} \approx \frac{R_1}{R_2}$.

Baum and Richlin (1969) found a systematic deviation from Herrnstein's law which later led Baum to his "generalized matching law" (Baum 1974):

$$\frac{P_1}{P_2} \approx a\left(\frac{R_1}{R_2}\right)^b \text{ where } a \approx 1, \text{ and } b \approx 0.9.$$

Neither of the versions of the matching law was successful in predicting pigeons behaviour once the VI schedules were replaced with VR⁷. Instead, after a learning period, pigeons spent almost all of their time on the most profitable side (Herrnstein and Loveland (1975)). An interesting observation is that pigeons faced with two VR schedules with the same mean, chose one of the sides and stayed there. In contrast, when faced with two VI schedules with the same mean time, pigeon spent about half of their time in each of the sides. This result seems odd at first since behaviour is seemingly affected by the particulars of the mechanism governing the delivery rate, despite it being unobservable.

Herrnstein and Loveland (1975), and later Myerson and Miezin (1980), suggested

certain key would mean that, on the average, it will be armed every 24 seconds.

⁷In a concurrent variable-ratio (VR) schedule, the program advances to the next stage as a function of the number of responses made by the subject. It does so using a random sequence of numbers with a given mean. For example, using a VR-45 means that the bird will, on the average, get food every 45 pecks it makes.

models that explain both behaviours. Matching and optimizing are parts of a more general rule of behavior: the matching law for observed reward. The main idea is that pigeons do not learn the relationship between their actions and the appearance of food, but rather maximize somewhat differently the amount of reward they observed. Both models depend on pigeon not knowing what happen in the side they did not go to. In all of the VI experiments pigeons could, by moving back and forth fast enough (with respect to the average delay of the schedules), get more food than it would have by staying at one side. Even if we are to assume that the pigeons had perfectly learned the structure of the experiment, the observed "matching" behaviour could be optimal.

5.6 Conclusions

There is large set of experiments where humans and other animals use probability matching behaviour in a repeated binary choice experiments with fixed probabilities. In many other experiments subjects behaviour was closer to the "optimal" one. In this paper we surveyed the literature and tried to identify some of the theoretical conditions that predict different asymptotic behavioral patterns. The findings are not conclusive in favor of any of the theories mentioned in section 2. Stochastic learning theories seems to make the best predictions overall. Taking estimates for behaviour in the group level only may lead to overlooking some regularities in the data. When individual data is considered it is clear that some S's asymptotically choose the most profitable outcome. More information about S's beliefs about outcomes in different stages of the experiment may prove useful in understanding why they behave as they do. Particularly it is important to see if S's choose behaviour that assumes some form of contingency between experiments. Information about the learning process, rather than just estimates of the asymptotic could also be useful.

Economists often find mathematical decision theories such as Estes' and Bush & Mosteller's hard to accept as those theories assumes too little about the individual's

decision making ability. However, one must accept that even in relatively simple decision situations people tend to behave very differently from what *homo economicus* would have done.

References

- ANDERSON N.H., AND WHALEN R. (1960), "Likelihood judgments and sequential effects in a two choice probability learning situation", *Journal of Experimental Psychology*, **60**, 111-120.
- ARROW K. (1962), "Utilities, attitudes, choices: A review note", *Econometrica*, **26**, 1-23.
- ATKINSON R.C. (1962) "Choice behavior and monetary payoffs", In *Mathematical methods in small group processes*, Stanford University Press.
- BANKS J.S. AND SUNDARAM R.K. (1992), "Denumerable-Armed Bandits", *Econometrica*, **5**.
- BAUM W.M. AND RICHLIN, (1969), "Choice as time allocation", *Journal of the Experimental Analysis of Behavior*, **12**.
- BAUM W.M. (1974), "On Two Types of deviation from the Matching Law: Bias and Undermatching", *Journal of Experimental Analysis of Behavior*, **22**.
- BÖRGERS T. AND SARIN R. (1993), "Learning through Reinforcement and Replicator Dynamics", University College London discussion paper, 1993-19.
- BRACKBILL N., AND BRAVOS A. (1962), "Supplementary report: The utility of correctly predicting infrequent events", *Journal of Experimental Psychology*, **64**, 648-649.

- BRACKBILL N., KAPPY M. S., AND STARR R. H., (1962), "Magnitude of Reward and Probability Learning", *Journal of Experimental Psychology*, **1**, 32-35.
- BRUNSWIK, (1939), "Probability as a determiner of rat behavior", *Journal of Experimental Psychology*, **25**.
- BURKE C.J., ESTES W.K. AND HELLYER S. (1954), "Rate of verbal conditioning in relation to stimulus variability", *Journal of Experimental Psychology*, **48**, 153-161.
- BUSH R. AND MOSTELLER F. (1955), *Stochastic models for learning*. New York: Wiley.
- COTTON J., AND RECHTSCHAFFEN A. (1958), "Replication report: Two-and-three-choice verbal conditioning phenomena", *Journal of Experimental Psychology*, **56**, 96.
- CRAWFORD V. P. (1995), "Theory and Experiment in the Analysis of Strategic Behaviour", Mimeo.
- CROSS J.G. (1973), "Stochastic Learning Model of Economic behaviour", *Quarterly Journal of Economics*, **87**.
- DERKS P. L., AND PACLISANU M. I. (1967), "Simple Strategies in Binary Prediction by Children and Adults", *Journal of Experimental Psychology*, **2**, 278-285.
- EDWARDS W. (1956), "Reward probability, amount, and information as determiners of sequential two-alternative decision", *Journal of Experimental Psychology*, **52**, 177-188.
- EDWARDS W. (1961), "Probability Learning in 1000 trials", *Journal of Experimental Psychology*, **4**, 385-394.
- EDWARDS W. (1962), "Dynamic decision making and probabilistic information processing", *Human Factors*.

- ENGLER J. (1958), "Marginal and conditional stimulus and response probabilities in verbal conditioning", *Journal of Experimental Psychology*, **55**, 303-317.
- ESTES W. (1957), "Theory of learning with constant, variable or contingent probabilities of reinforcement", *Psychometrika*, **22**.
- ESTES W.K. AND STRUGHAN J.H. (1954), "Analysis of a verbal conditioning situation in terms of statistical learning theory", *Journal of Experimental Psychology*, **47**, 225-234.
- ESTES W. (1964), "Probability Learning", In A.W. Melton editor: *Categories of Human Learning*.
- FIORINA P. M. (1971), "Critique and Comments: A Note on Probability Matching and Rational Choice", *Behavioral Science*, **16**, 158-166.
- FRIEDMAN M.P., PADILLA G., AND GELFAND H. (1964), "The learning of choice between bets", *Journal of Mathematical Psychology*.
- GALANTER E.H. AND SMITH A.S. (1958), "Some experiments on a simple thought problem", *American Journal of Psychology*, **71**, 359-366.
- GARDNER R. (1957), "Probability learning with two and three options", *American Journal of Psychology*, **70**, 174-185.
- GARDNER R. (1958), "Multi-choice decision behavior", *American Journal of Psychology*, **71**.
- GITTINS J. (1989), *Allocation Indices for Multi-Armed Bandits*. London: Wiley.
- GOODNOW J. (1955), "Determinants of Choice distribution in two choice situations", *American Journal of Psychology*, **68**, 106-116.

- GRANT D.A., HAKE H.W., AND HORNSETH J.P. (1951), "Acquisition and extinction of verbal expectations in situation analogous to conditioning", *Journal of Experimental Psychology*, **42**, 1-5.
- HAKE H.W. AND HYMAN R. (1953), "Perception of the statistical structure of a random series of binary symbols", *Journal of Experimental Psychology*, **45**, 64-74.
- HERRNSTEIN (1961), "Relative and absolute strength of response as a function of frequency of reinforcement", *Journal of Experimental Analysis of Behavior*, **4**.
- HERRNSTEIN (1974), "Formal Properties of the Matching Law", *Journal of Experimental Analysis of Behavior*, **21**.
- HERRNSTEIN AND LOVELAND (1975), "Maximizing and Matching on concurrent ratio schedules", *Journal of Experimental Analysis of Behavior*, **24**.
- HUMPHREYS L. G. (1939), "Acquisition and extinction of verbal expectations in a situation analogous to conditioning", *Journal of Experimental Psychology*, **25**.
- JARVIK M.E. (1951), "Probability learning and a negative recency effect in a serial anticipation of alternative symbols", *Journal of Experimental Psychology*, **41**, 291-297.
- LUCE D. AND SUPPES P. (1965), *Preference, Utility, and Subjective Probabilities* in Luce D., Bush R., and Galanter E., *Handbook of Mathematical Psychology, Volume III*. New York: Wiley, 1965.
- MCCORMACK P. (1959), "Spatial generalization and probability learning in five-choice situation", *American Journal of Psychology*, **72**.
- MORSE E., AND RUNDQUIST W. (1960), "Probability matching with an unscheduled random sequence", *American Journal of Psychology*, **73**.

- MYERS J.L., FORT J.G., KATZ L., AND SUYDAM M. (1963), "Differential memory gains and losses and event probability in a two-choice situation", *Journal of Experimental Psychology*, **66**, 521-522.
- MYERSON AND MIEZIN (1980), "The kinetics of choice: An operant system analysis", *Psychological review*, **87**.
- NEIMARK, E. (1956), "Effect of type of non reinforcement and number of alternative responses in two verbal conditioning situations", *Journal of Experimental Psychology*, **52**, 209-220.
- NEIMARK E. AND SHUFOD E. (1959), "Comparison of predictions and estimates in a probability learning situation", *Journal of Experimental Psychology*, **57**, 294-298.
- NICKS D.C. (1957), "Prediction of sequential two-choice decisions from event runs", *Journal of Experimental Psychology*, **57**, 105-114.
- ROTSCHILD M. (1974), "A Two-Armed Bandit Theory of Market Pricing", *Journal of Economic Theory*, **9**, pp185-202.
- RUBINSTEIN I. (1959), "Some factors in probability matching", *Journal of Experimental Psychology*, **57**, 413-416.
- SIEGEL S., AND GOLDSTEIN D. A. "Decision making behaviour in a two-choice uncertain outcome situation", *Journal of Experimental Psychology*, **57**, 37-42.
- SIMON H. (1976), "Rational Choice and the Structure of the Environment", *Psychological Review*, **63**,129-138.
- SUPPES P., AND ATKINSON R.C. (1960), *Markov learning models for multi person interactions*. Stanford: Stanford University Press.
- SUPPES P. (1961), "Behavioristic foundations of utility", *Econometrica*, **29**, 186-202.

TVERSKY A. AND EDWARDS W. (1965), "Information versus Reward in Binary Choices",
Journal of Experimental Psychology, 71.