Nash's Nobel

Andrew D Smith, UK
August 2002

Abstract: We've all been to see A Beautiful Mind, the Oscar winning film about John Nash, starring Russell Crowe. Although Nash won a Nobel for economics, the film focuses on the love story and mental health aspects. So if you're the rare actuary whose education has left intact some residue of intellectual curiosity, you might have left the film with a few unanswered questions. In insurance, game theory has been proposed for evaluation of strategic options involving competitor responses, and explanation of phenomena such as premium cycles. This note explains the concept of a Nash equilibrium, outlining the game theory context and the key results for which Nash gained his Nobel.

Keywords: Nash Equilibrium, Non-cooperative game, Nobel Prize

Contact details:
Andrew D Smith, B&W Deloitte, Horizon House, 28 Upper High Street, Epsom, Surrey KT18 4RS, United Kingdom.
email: andrewdsmith8@bw-deloitte.com

Introduction

In 1994, the Bank of Sweden Prize in Economic Sciences in Memory of Alfred Nobel was awarded to John Harsanyi, to John Nash and to Reinhard Selten, for their pioneering analysis of equilibria in the theory of non-cooperative games.

Subsequently, John Nash found himself the focus of some media interest. John Nash's biography, "A beautiful mind", written by Sylvia Nasar became a best seller in 1998. A film of the same title, starring Russell Crowe as John Nash and Jennifer Connelly as his wife Alicia, has won four Oscars.

The popular interest was largely due to John Nash's tragic personal circumstances. In 1958, on the threshold of his career, paranoid schizophrenia struck. Nash lost his job at M.I.T. in 1959 and was virtually incapacitated by the disease for the next two decades.

This note provides a brief summary of the economic contributions recognised in the 1994 Nobel Prize. For the human interest we refer you to the book. If you want a good cry and a hopelessly romantic "love conquers all" view of the world then go see the film.

Game Theory – a Short History

Game theory considers situations where someone's behaviour is influenced by his forecasts of his opponents' behaviour. The classical examples relate to pricing and capacity decisions in oligopoly situations.

An 1838 paper by the French economist Cournot is usually cited as the genesis of game theory. Cournot investigated the situation of two suppliers to a market, who have to make decisions as to how much to produce. The optimal production for each supplier depends on the other supplier's production decision. As the suppliers must choose their production simultaneously, each producer must estimate the other supplier's decision. It is this circularity, which characterises game theory.

The next major contribution to game theory was von Neumann and Morgenstern's 1994 book *The Theory of Games and Economic Behaviour*. Actuaries know this text better for introducing the concept of a utility function, and setting it up on an axiomatic basis. The game theoretic contribution was far more than this though. Particularly groundbreaking was the treatment of two-person zero sum games.

In 1950, Nash proposed what came to be known as "Nash Equilibrium". This extended the von-Neumann and Morgenstern solution to non-zero-sum games and to multiple players. A Nash equilibrium applies when each player's strategy is a payoff-maximising response to the strategies pursued by the other players. This concept has been the point of departure for most economic work in the field of game theory.

Nash also provided a mathematical proof that, under certain conditions, a Nash equilibrium exists. There are a number of complexities to applying this in practice, including the problem of multiple equilibria and the strong economic assumption that a game's participants are aware of each other's constraints and preferences. The other Nobel winners have contributed to resolving these issues.

In 1965, Selten published a paper investigating multi-stage games. In such games, a player's move can be contingent on what she observes from the moves of other players in previous stages of the game. In such games, many Nash equilibria may exist, so an economist seeking to predict the game's outcome must find some way of ranking the equilibria in order of how likely they are to occur. Selten argued that in multi-stage games, many of the Nash equilibria rely on "empty threats"; the remaining "subgame perfect" equilibria are more plausible.

Harsanyi's contribution in 1967 was to extend the concept of Nash equilibrium to situations where players are uncertain about other players' payoffs. He did this by allowing players to start with a Bayesian prior distribution describing other players' payoffs – an idea with which many actuaries will be familiar.

Cournot Equilibrium – An Example

We start with a simple example of a Cournot equilibrium. There are two manufacturers in the widget market. Player 1 has a production cost of €1 per unit, while player 2 has a production cost of €2 per unit. The market price depends on the total production, and is given by the formula:

market price per unit = €6 * (1 – total production / 120)

The Cournot equilibrium applies when player 1 makes 40 units and player 2 makes 20 units. The market price is then €6 * (1 – 60/120) =

€3 per unit. Player 1 therefore makes a profit of €80, and player 2 a profit of €20.
To see why this is a Cournot equilibrium, let us suppose that player 1 already knows player 2's strategy. Player 1's strategy then is to maximise profit, knowing that player 2 will make 20 units. For each possible level of own production, player 1 can add 20 to get the total production, and thus a forecast both of the market price per widget and of player 1 profit. Player 1's profit maximisation is shown in the chart below:

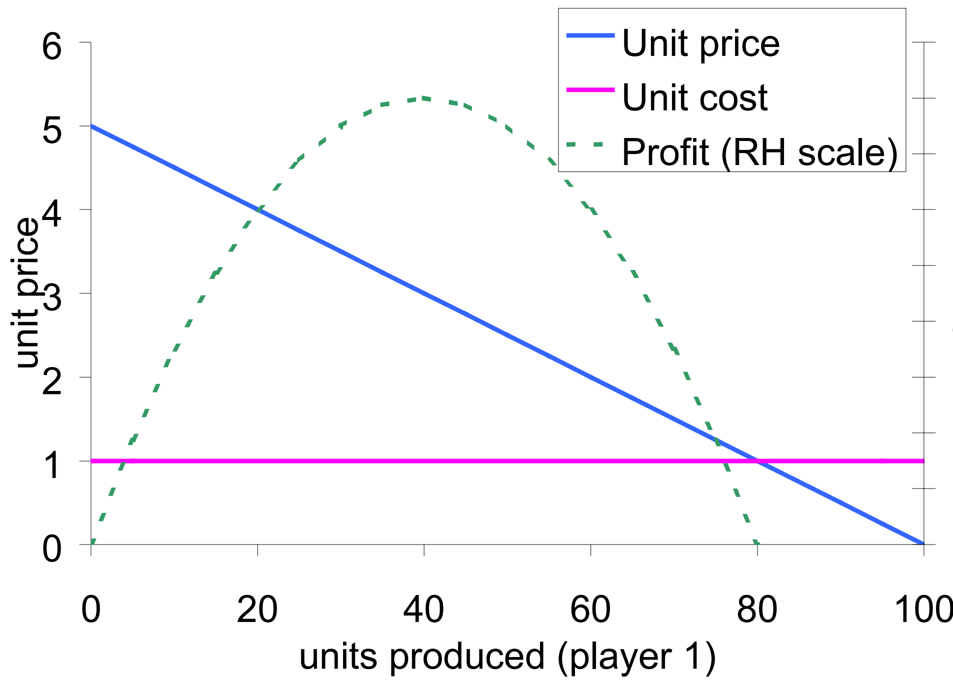

Player 1 maximises his own profit producing 40 units.

For Player 2, the situation is the similar. Given player 1's strategy of producing 40 units, player 2's optimal response is to produce 20 units. In this case, (40, 20) is the unique Cournot equilibrium in which each player responds optimally to the strategy of the other.

Cournot equilbria provide rich hunting grounds for would-be social planners. There are ways in which the Cournot equilibrium could be improved for both players, if they cooperate. For example, player 1 could produce more of the goods, and provide player 2 with a side-payment as an incentive not to produce anything. Both players could then increase their profits. As the goods are now being produced at lower average cost, there may also be a wider benefit to society. Is there then a case for price regulation to tame player 1's monopoly?

von Neumann – MorgenStern Equilibrium – An Example

A simple example of the von Neumann – Morgenstern equilibrium is the two-player children's game "rock, paper, scissors", also known (in the UK) as "scissors, paper, stone". The following rules are an abbreviation of the full set which are published by the World Rock Paper Scissors society:

The two players must each release one fist in any of the following manners:

- **Rock**: represented by a closed fist with the thumb resting at least at the same height as the topmost finger of the hand. The fingers must not conceal the thumb.
- **Scissors**: Is delivered in the same manner as rock with the exception that the index and middle fingers are fully extended toward the opposing player. It is considered good form to angle the topmost finger upwards and the lower finger downwards in order to create a roughly 30–45 degree angle between the two digits and thus mimic a pair of scissors.
- **Paper**: Is also delivered in the same manner as rock with the exception that all fingers including the thumb are fully extended and horizontal with the points of the fingers facing the opposing player.

Each player has the full range of throws to play, as follows:

- Rock wins against scissors, loses to paper and stalemates against itself
- Paper wins against Rock, loses to scissors and stalemates against itself
- Scissors wins against paper, loses to rock and stalemates against itself

These are sometimes summarised as "rock sharpens scissors, scissors cut paper, paper covers rock". In the case of a stalemate, where players reveal the same throw the round must be replayed. There are no limits to the numbers of stalemates, which may occur in any given match. Should players find themselves in a continuous stalemate situation, also known as "Mirror Play", a good approach can be to take a short "timeout" to rethink your strategy.

Plainly no deterministic strategy wins Rock Paper Scissors. For example, if my strategy is to choose "rock" then my opponent will choose "paper". The only strategy, which cannot be beaten in this way, is a strategy to choose the position randomly between the three alternatives. This doesn't give me a winning strategy – on average I break even – which is the best that can be hoped for in a symmetric zero sum game. Such randomised strategies are sometimes called "mixed", as opposed to the "pure" strategies, which involve deterministic play.

This example is important because it indicates how randomness can be introduced into the solution to a game, even when the rules contain no random element.

Although von Neumann and Morganstern applied their work to the theory of games, the mathematical results on which they base their exposition are much more general – and indeed represent the first step in formalising solutions to linear programming. For many years, linear programming problems were regarded as applications of game theory, even when the problem formulation was unrelated to any game. Benjamin (1959) is one of the few attempts to apply this literature in an actuarial context – in this case to a minimax definition of prudent valuation bases.

Nash's Theorem

We now extend the concept of a game to allow many players and possible non-zero sums. An *N*-player game is characterised by its payoff matrix, and *N*+1 dimensional matrix describing the payoff to each player for every combination of plays.

Nash's major contribution to game theory was a proof that such games always have a mixed-strategy equilibrium. Pure strategy (Cournot) equilibria may not exist. The proof itself is a few pages only, but it relies on a deep topological result – the fixed-point theorem of Kakutani (1941). Other authors were quick to generalise the result to compact strategy spaces with continuous payoffs. Debreu (1952) characterised a subset of problems with concave payoffs where a pure strategy equilibrium could be guaranteed.

The Bargaining Problem

We now consider a version of a classic problem, which Nash (1950) solved prior to his more general game theory paper.

Consider a game with two players. Player A has an asset, which he wishes to sell. He knows that he can sell the asset to a third party dealer for €5, but he hopes that by bargaining with player B, he can get a better deal.

Player B has the reverse situation of wanting to buy the asset, but the dealer's asking price is €10, incorporating a €5 dealing spread over the market bid price. Player B is also hopeful that by bargaining with player

A he can negotiate a price better than either can achieve by going to a third party dealer.

The process works as follows. The payers take it in turns to suggest prices to each other. As each price is suggested, the other player may accept or decline. The gain to player A is the extent to which the negotiated price exceeds the dealers bid price of €5; for player B the gain is the extent to which the negotiated price falls short of the €10 asked by the dealer. After declining a price, either player may terminate the game, in which point each player goes to a dealer and there are no gains to either player.

Edgeworth first proposed this problem in 1881. He argued that, without competition, the solution was indeterminate; there is no way of predicting what bargain might be struck. Over the next 70 years many other great economists, including John Hicks and Alfred Marshall, took up this problem but made no headway.

To make some progress, let us assume that both players are also impatient. Let us suppose that for every step of the game (that is, for every price suggestion followed by an accept/decline decision), beyond the first step, player A incurs a cost of €1 and player B incurs a cost of €2. Can we then solve the game?

Even with these cost structures, there are many Nash equilbria. For example:

Nash Equilibrium #1
Player A always asks a price of €5 and accepts any bid of €5 or more
Player B always bids €5 and will pay a maximum asking price of €5

is a Nash equilibrium – indeed one of many. Each strategy is a rational response to the strategy of the other.

However, this equilibrium looks wrong. It would result in player A selling the asset to player B for €5 at the first move. Player B has captured all the gains from private negotiation. This seems odd given that player B has the higher time cost; we would expect that player B would be keener to do a deal. Player A should be able to exploit player B's impatience to extract a higher price than €5.

Selten (1965) provided a more precise diagnosis of the problem with our equilibrium #1. It involves player B making empty threats. For example, if player A offers to sell for €6, then player B should accept. This is because, if player B is to wait, he'll still pay at least €5 anyway, and incurs an additional waiting cost of €2. Player B's threat to reject an offer to sell at €6 is an empty threat.

So let us focus on strategies, which eliminate empty threats. In this case, we find the following unique solution:

Selten Equilibrium
Player A always asks €10 but accepts any bid of €9 or above
Player B always bids €9 but pays any asking price up to €10

In other words, if player A moves first, he succeeds in selling for €10. Player A, with more patience, has gained all the benefits of negotiation. However, if player B moves first, he can expect player A to accept a bid of €9. This situation is described as a first mover advantage. Even in this case, it is clear that most of the gains from bargaining have gone to the more patient player.

How does this fit to the way the world works? In real life, we might commonly observe two players agreeing to "split the difference", and transact at €7.50. A strategy where each player suggests €7.50 and accepts €7.50 is a Nash equilibrium but not a Selten equilibrium. Game theory defines the notion of a Selten Equilibrium. It is an empirical question whether Selten Equilibria or other concepts are best able to explain the outputs of real life games.

Our Selten equilibrium predicts that players will agree a bargain in the first round, but in practice we know that haggling does take place. This may point either to a failure in the theory's predictive power, or alternatively a weakness in our formulation of the problem.

Haggling is sub-optimal for both players if they have complete information about each other's cost structures. On the other hand, it might be more realistic to assume the players have patchy information about each other. In that case, we could see a situation where an initial haggling phase is useful in providing information to each player about the other's cost structure. This would be an example of a Harsanyi (1967) equilibrium. Needless to say, incorporating such a Bayesian updating process into the game makes the analysis far more complicated. These are the same Selten and Harsanyi with whom Nash shared the Nobel prize.

Other softer factors may also be at work in the bargaining process, and some have argued that bargaining must be set in a cultural context of what is fair, given the possibility of further business following a successful deal. The acceptability of haggling in a shop or street market varies greatly from one part of the world to another. Can these differences really be captured by reference to different time costs and information structures?

Does it Work?

The relevance of Nash equilibria to modern economic practice is still controversial. Perhaps the most promising area is the design of auctions. Milgrom (1995) built on the work of Nash, Selten and Harsanyi to consider the optimal design of public auctions. At the time, Milgrom was advising the US government on ways of selling licenses to use airwave bandwidth for cellular telephones. In March 1995, the US government announced that the winning bids totalled more than US$7 000 million. Milgrom described this as "the biggest sale in American history of public assets and one of the most successful (and lucrative) applications of economic theory to public policy ever".

This auction followed a number of less successful previous attempts. According to Nasar (1998), "before 1994 Washington simply gave away licenses for free. Until 1982 it had been up to regulators to decide which companies deserved the licenses. … After 1982 Washington awarded licenses using lotteries". The successful American auction also followed costly flops of less well-designed auctions in Australia and New Zealand.

These auctions can be long and protracted affairs. In April 2000, the UK government announced the results of its own airwave auction, using similar techniques to the US. After 150 rounds, the accepted bids totalled £22 470 m.

This proved to be the high point of the telecom auction frenzy. In April 2001, the BBC reported that " as companies began to balk from paying huge sums for the licences, the only way governments were guaranteed to make money was to set the price in advance and award the licences on merit through a beauty contest. " Amounts raised from these less sophisticated means were as follows:

France:	£ 6 320m
Germany:	£ 30 400m
Italy:	£ 7 500m
Netherlands:	£ 1 680 m
Poland:	£ 1 900 m
Sweden:	£ 26 m
Switzerland:	£ 80 m
Belgium:	£ 300m

Australia:	£ 500m
Spain:	£ 12m

Game theory enthusiasts might explain the low revenues by poor auction design. But were the fluctuations simply reflecting the boom and subsequent collapse in the fortunes of telecom companies? Does game theory justly deserve the credit Milgrom claims?

Historically, the US Military has been one of the biggest spenders on game theory research – and indeed, employed Nash for several years. More recently, Major (2002) provided a description of game theory, and an approach to managing terrorism risk, although the relationship between the two parts of the document is not entirely clear. We can describe some well-known conflicts in game theoretic terms. For example, the cold war strategy of “mutually assured destruction” is an example of a Nash equilibrium, but not a Selten equilibrium. For example, suppose that USA makes an unprovoked nuclear attack on Moscow. The USSR has threatened to retaliate, but given the existing damage to Moscow, and that the USA has already launched its deadly attack, would this be a rational response? There is no sense of teaching a lesson, because the destruction is final. Furthermore, by launching a nuclear attack on the USA, the USSR would deliver to its Asian citizens a far more lethal does of radiation via the trade winds, than they would have suffered as a result of Moscow’s bombardment. Arguably, the USSR threat of retaliatory action is empty in the Selten sense. But few forecasted the end of the cold war – not even game theorists seeking Selten equilibria.

The Financial Times of 25 March 2002 recounted another application of game theory, as applied by a UK firm:

“In 1994, Yorkshire Water, the privatised UK utility, made extensive use of game theory while preparing for a regulatory review that would set prices from 1995 to 2000. Trevor Newton, then managing director, felt confident on the basis of his game theoretic analysis that the company would secure a favourable outcome.

In fact, the industry regulator not only set tougher price controls than the company and its investors had expected; it also opened an inquiry into Yorkshire's operating performance. Following a year of drought and regional water shortages in 1995, Mr Newton resigned from the company. “

Green (2002), a New Zealand based academic, tried the following experiment. Twenty-one game theorists made 99 forecasts of decisions for six conflict situations. The same situations were described to 290 research participants, who made 207 forecasts using unaided judgement, and to 933 participants, who made 158 forecasts in active role-playing. He found that role-play predictions were better than chance and unaided judgement for all situations, and better than game-theory experts' predictions for all but one situation. Game theorists' forecasts, on the other hand, varied more widely in their accuracy than did role-play forecasts.

So overall, the case for using game theory in corporate decision-making is not yet proven. There are success stories and there are failure stories.

Actuarial Applications?

Are Nash equilibria ever going to be of practical use to actuaries? Is this a fast developing area where actuaries must work hard to catch up, or is game theory a peripheral discipline which most of us can leave to the specialists?

The first thing to stress is that applied game theory is not yet a quantitative discipline. I am not aware of a single example where a businessperson has sought to calibrate his own and competitors' payoffs and then successfully forecasted the future by solving for a Nash equilibrium. Instead, most claimed applications of game theory actually involve people using ideas, concepts or insights from game theory in a judgmental fashion. In fact, it is not clear whether game theory is actually being used at all. Claims to use Nash equilibria may turn out to be the application of simpler general reasoning, given a veneer of rigour by the adoption of a Nobel-winning name.

Why has the quantitative progress been so slow? The main reason is the difficulty of formulating even simple problems in a game-theoretic framework. There are many parameters to estimate to formulate the problem, most of which relate to hypothetical payoffs under strategies, which have not in the past been followed – so for which no supporting data is available. On top of this, Nash's equilibrium theorem is merely an existence result, which gives us no guidance on how to characterise numerically the set of Nash equilibria. Such numerical algorithms are not well developed. Furthermore, Nash's result applies to finite games – computations for the real world of continuous time and continuous valued decision variables are uncharted territory. For these reasons, several major leaps forward on the theoretical front are needed before the computation of Nash equilibria could be useful for day-to-day decision support.

One promising avenue of research applies game theory to capital allocation. Denault (2001) treats the allocation of capital in a company as a multi-player game, each line of business representing one player. He derives a capital allocation based on work by Aumann and Shapley (1974), using bargaining theory to establish a formula for distributing capital credits for diversification.

In life assurance, there is an element of game theory in underwriting decisions. How much should an insurer spend on checking health details of assurance applicants? As an insured life, what is the incentive to lie on assurance forms (for example regarding non-smoker status). What are the potential selection effects for example from HIV positive individuals who withhold details of medical tests? What is a rational underwriting response to these moral hazards? Such effects are considered in Cummins *et al* (1982). It is unclear to what extent these techniques are used in practice, but the Society of Actuaries in the US requires some knowledge of Nash equilbria in this context for its specialist life subject.

There are also public policy implications. The neo-classical perspective, initiated by Adam Smith, has long argued that private contracting is likely to give rise to resource allocations, which are in some sense socially optimal. If these arguments were completely accepted, then there would be no role for any form of government intervention in the economy. But game theory gives us a whole host of examples where non-cooperative games can have socially sub-optimal Nash equilibria. This could lend weight to those who would replace Adam Smith's invisible hand with the long arm of the state. Such views remain controversial. Many would argue it is more realistic to model government as yet another class of players, with their own selfish objectives, rather than as a benign coordinator. Even so, on balance, the concept of the Nash equilibrium has given comfort to advocates of interventionist economic policies.

Another possible game theory application relates to market volatility. It has been widely noted (see, for example, Geman and Ané, 2000) that high trading volumes often accompany times of high market volatility. The question here is of cause and effect. One explanation would be that at times when markets are moving fast, investors' preferred asset allocations and dynamic hedges will also move, and this alone should increase trading volume. Another school argues that the trading itself causes the volatility (with the consequence that regulatory or fiscal intervention to reduce trading volume could have a social benefit of reduced market volatility). Where does this volatility come from – a mixed strategy Nash equilibrium perhaps? Disappointingly, most formulations of investment problems result in concave payoffs, in which case equilbria are pure strategy.

The market impact of sales and purchases is an area of much current study. Until recently, the dominant belief was that an attempt to sell a large quantity of a share would automatically produce a large price fall – although the mechanism causing other players to be prepared to sell at a lower price was unclear. Investment managers would suggest strict secrecy and small incremental trades to beat the market impact. Recent literature has suggested that market impact is instead attributable to rational fears of inside information – of which large attempted trades are a signal. The consequence is that a large trade not motivated by inside information should have a smaller price impact – provided this fact is communicated credibly to the market. Arguably, trades by Boots pension fund out of equities into bonds, have had a smaller market impact thanks to clear communication of the matching rationale for these trades (see Alexander, 2002). Many life insurers and pension funds are seeking to reduce equity exposure. Is there a game theoretic angle to the timing and information management decisions involved?

Insurance premium cycles have ruined many insurance companies, inconvenienced customers and confounded regulators. Many popular explanations resort to assuming irrationality on the part of market participants. From an economist's view this is unsatisfactory, as it implausibly supposes the economist has more information about the market than its own participants do. Furthermore, surveys seldom reveal insurers taking the blame for their own participation in the premium cycle. Indeed, insurers are near unanimous in their diagnosis of the problem – its everybody else's strategies that forces an insurer to behave in they way he does. This is essentially the definition of a Nash equilibrium.

Feldblum (1992) provides persuasive explanations of market cycles in terms of rational participants in the presence of principal/agent conflicts and costs of entry / exit. Feldblum's theory has a number of parallels with recent developments in the theory of repeated non-cooperative games, which combine Nash equilibrium concepts with ideas from dynamic programming. Such games can produce cyclical behaviour. It is possible that Feldblum's explanations of market cycles can be expressed in terms of a Nash equilibrium. The important question from a practical perspective is whether re-casting an insurance phenomenon in terms of games theory provides any extra insight into the insurance problem. More work remains in order to investigate whether general theorems from game theory will yield specific insights to insurance phenomena, or might even suggest ways that a well-designed regulatory environment could mitigate adverse social consequences of insurance cycles.

Game theory is still developing rapidly. It holds out a promise – as yet unfulfilled – of explaining puzzling effects in insurance and capital markets. Without some major breakthroughs, game theory is not going to be running our lives next week. However, actuaries should keep abreast of developments, and be ready to adopt game theoretic tools as they become more practical to apply.

Bibliography

Alexander, B (2002). *Gentlemen Prefer Bonds – Boots Pension Fund Case Study.* London Business School.

Aumann R & Shapley, L. (1974) *Values of Non-Atomic Games*. Princeton University Press.

British Broadcasting Corporation (2001) *Scramble for Licenses*. http://news.bbc.co.uk/2/hi/business/1272501.stm

Benjamin, S (1959). *The Theory of Games and its Application to Interest Rates*. JIA 85, 373-413.

Cournot, A (1838) Researches into the Mathematical Principles of the Theory of Wealth (English translation published by MacMillan, 1897).

Cummins, J. David, B.D. Smith, R.N. Vance, J.L. Vanderhei. (1982) *Risk Classification in Life Insurance* , Kluwer.

Debreu (1952) *A social equilibrium existence theorem*. Proceedings of the National Academy of Sciences 38:886-893.

Denault, M. (2001). *Coherent Allocation of Risk Capital*. Journal of Risk vol 4 #1 http://www.riskwaters.com/jcf/freepdfs/jor_v4n1a1.pdf

Edgeworth (1881) *Mathematical Psychics: An Essay on the Application of Mathematics to Moral Sciences* (C Keegan Paul, London)

Financial Times (25 March 2002) *Games or Serious Business?* http://news.ft.com/ft/gx.cgi/ftc?pagename=View&c=Article&cid=FT32FI4K8ZC&live=true

Feldblum, S (1992). *Underwriting cycles and Insurance Solvency* Discussion Paper Program, Casualty Actuarial Society http://www.casact.org/pubs/dpp/dpp92/92dpp383.pdf

Fudenberg, D and Tirole, J. *Game Theory*, MIT Press.

Geman, H and Ané T (2000) *Order flow, transaction clock, and normality of asset returns*, Journal of Finance 55 (2000), no. 5, 2259-2284.

Green, K (2002) *A comparison of Forecasting Decisions in conflict situations: game theory, role-playing, and unaided judgement* International Journal of Forecasting 18(3): 321-344

Harsanyi, J (1967). *Games with Incomplete Information played by Bayesian Players*. Management Science 14:159-182, 320-334, 468-502.

Kakutani, S (1941): *A Generalization of Brouwer's Fixed Point Theorem*, Duke Mathematics Journal., 8, 457-458.

Major, J (2002). *Advanced Techniques for Modelling Terrorism Risk*. Available at http://www.guycarpenter.com/pdf/major_terrorism.pdf

Milgrom, P. (1995). *Auction Theory for Privatisation.* Cambridge University Press.

Nasar, Sylvia (1998) *A Beautiful Mind*. Faber and Faber (published in the US by Simon & Schuster).

Nash, J (1950) *The Bargaining Problem*. Econometrica, vol 18 (1950) 155-162.

Nash, J (1951) *Non-cooperative games*. Annals of Mathematics, vol 54, 286-295

von Neumann, J & Morgenstern, O (1944). *The Theory of Games and Economic Behaviour*. Princeton.

Osborne, M & Rubinstein, A (1994): A course in Game Theory, MIT Press

Rubinstein A (1982). *Perfect equilibrium in a bargaining model*. Econometrica 50:97-110.

Selten, R (1965). *Spieltheorische Behandlung eines Oligopolmodells mit Nachfrageträgheit*. Zeitschrift für die gesamte Staatswissenschaft 12:301-324.

World Rock Paper Stone Society (2002) *The Official Rules of Rock Paper Scissors play*. http://www.worldrps.com/gbasics.html