Game Theory - Analysis of Conflict
Roger B. Myerson.

1. Decision Theoretic Foundations.

A decision maker is rational if he makes decisions consistently in pursuit of his own objectives.

A decision maker is intelligent if he knows everything that we know about the game.

The assumption that players are rational and intelligent may never be satisfied in real life but we should be suspicious of theories that are not consistent with these.

Decisions are often described by a probability model (lotteries are probability distributions over a set of prizes) or a state variable model (lotteries are functions from a set of possible states into a set of prizes). A probability model is better for objective unknowns (such as the roll of a die) while a state variable model is better for subjective unknowns (such as the outcome of a horse race or the stock market).

Axioms

(1.1a) Completeness and (1.1b) Transitivity assert that preferences should always form a complete transitive order over the set of lotteries.
(1.2) Relevance asserts that only the possible states are relevant to the decision maker.
(1.3) Monoticity asserts that a higher probability of getting a better lottery is always better.
(1.4) Continuity asserts that if probability between lotteries increase in a continuous manner then there is some lottery that is equivalent to the combination.

The four substitution axioms (1.5a-b & 1.6a-b) assert that, if the decision maker must choose between two alternatives and if there are two mutually exclusive events, one of which must occur, such that in each event he would prefer the first alternative, then he must prefer the first alternative before he learns which event occurs.

Theorems

Expected-Utility Maximisation Theorem states (1.3) utility functions range between 0-1, (1.4) the conditional probability of one event is the product of its component events, and (1.5) a decision maker always prefers lotteries with a higher expected utility.
Experimental investigation of decision making has revealed some systematic violation of expected utility maximisation. Three well known examples show three different violations....

Utility functions appear inapplicable when people would prefer a prize with zero risk over a larger prize with a minor risk.

Subjective probability seems inapplicable when people prefer to bet on a known situation (toss of a coin) over an unknown situation (the outcome of a competition between unknown players)

Any economic model seems inapplicable when people make a different decision when the loose $40 cash vs when they loose something replaceable at a cost of $40.

One decision-option is strongly dominated by another when it is a inferior outcome regardless. It is weakly dominated when the outcome is inferior in all case bar one when it is equal.

2. Basic Models

Games in extensive form are depicted as a tree with each node as a decision point (either player or chance). The information state is noted at each node to distinguish the information available to the decision maker. If two nodes have the same information state then they cannot be distinguished by the decision maker.

The strategic form of a game specifies N - the set of players, Ci - the set of (pure) strategies open to each player I, and Ui(c) - the expected utility pay-off for each player if c were the set of strategies implemented by other players. This is equivalent to the extensive form and generally considered to be derived from it.

The normal representation is a table with a player on each axis, their options as columns/rows and the pay-offs as the content.

Two games may be said to be equivalent when a change in the utility functions results in the same preference ordering. Another form of equivalence is best-response which is based on the narrower view that a players utility function serves only to characterise how he should choose his strategy once his beliefs about the other players behaviour are specified.

Strategies are pay-off equivalent iff no matter what the other players do, no player may ever care if player I used one or the other.

When pay-off equivalent strategies are merged then the normal representation is called the purely reduced normal representation.

A strategy is randomly redundant iff there is some way that the player can randomly choose among his other strategies to create a pay-off equivalent strategy and, when eliminated, the result is the fully reduced normal representation.

Iteratively removing strongly dominated strategies results in the residual game.

A multi-agent representation places a temporary agent at each node of the normal representation. The strategic form derived from this approach may not show any strongly dominated strategies even though the normal strategic representation does. This places a question over the elimination of strongly dominated strategies.

A fact is common knowledge if every player knows it and every player knows that every player knows it ad infinitum.

A players private information is any that is not common knowledge. A historical chance node is required in extensive form whenever the players begin (determine their strategies) with some private information. The historical node takes the start back to when everything was common knowledge. This may be modelled in extensive form with a historical chance node with a probability distribution that determines the type of the players. But this is awkward because it is often unrealistic to assume people determine their full strategy prior to knowing their type.

A Bayesian Game is a game that specifies a set of possible actions Ci and types (Ti) complete for each player (i) in N. Thus a strategy in a Bayesian game is a function from the players set of types into his set of actions.

All of the Bayesian games presented here are have beliefs which are a common prior. However, inconsistent situations commonly exist (i.e. two soccer coaches who are both aware of each others opinions but still both believe that their team has a 60% chance of winning). In a consistent model, differences in opinion can be explained as differences in information.

Consistency and independence of types can be important when we make a private values assumption where utility is a function of type (private information) or a transferable unity assumption where utility might = money for instance.

A type-agent representation is where there is one player or agent for every possible type of player in a Bayesian game.

However there are fundamental problems with real world modelling which often arise when players beliefs are categorised by subjective probability (i.e. does the opponent know the answer to a trivia question)

3. Equilibria of Strategic-form Games.

No rational intelligent player will use a strategy that is iteratively dominated. But this does not reduce our solution set much.

A randomised-strategy profile is a Nash equilibrium iff no player could increase his expected pay-off by unilaterally deviating from the prediction of the randomised-strategy profile.

Theorem 3.1: There exists at least one equilibrium in any finite game.

Games can have multiple equilibria and inefficient equilibria. A game is weakly Pareto efficient iff there is no other outcome that would make all players better off. Consider the Prisoners dilemma (inefficient) and the Battle of the Sexes (multiple equilibria)

We can solve for the Nash equilibria in two player games but 3+ players becomes non-linear (i.e. good luck)

Rational intelligent players must end up at a Nash equilibria.

The focal point effect is when a minor piece of public information tends to focus the players attention and preference for one particular Nash equilibrium. Someone who makes preplay communication is a focal-arbitrator. Welfare properties of equity and efficiency can determine the focal equilibrium or sometimes even the nature of the strategies themselves (i.e. a preference for pure strategies over randomised strategies)

The Decision-analytic approach is an alternate where a player first accesses some subjective probability distribution to summarise their beliefs about what strategies will be used by other players before deciding their own strategy. It relaxes the assumption of rational intelligent opponents but quickly gets into trouble when player I considers deeply how other players might be thinking.

An evolution approach takes animals which are preprogrammed to use particular strategies and allows them to play successive games which ultimately determines the population of like animals in successive generations. The resistance of an equilibrium against another is the maximal fraction of I programmed animals that could be inserted into a population of j programmed animals such that the i's would have no advantage over the j's. Other evolutionary models define a viscosity which recognises that the i's will tend to end up in a herd together and compete with each other with symmetrical strategies.

Two person zero sum games are when two people are in pure opposition and ones gain is the others loss.

"The theory of randomised equilibria identifies situations where minor private information may be decisive, and focal-point effect identifies situations where minor public information may be decisive." pg 131

Auctions can involve independent private values (where each bidder knows privately the actual value to himself i.e. a house) or common value situations (where the value to all bidders is common but the estimates vary i.e. oil drilling rights). In the common value situation it is important that the bidder estimate the value of the object by its conditionally expected value given his current information and the additional information that could be inferred if this bid won the auction. Good example beginning bottom pg 133.

Infinite strategy sets - uuuugly!

4 Sequential Equilibria of Extensive-Form Games

A mixed-strategy profile of Te is defined to be any randomised-strategy profile for the normal representation of Te (i.e. a probability distribution over a set of overall strategies).

A behavioural-strategy profile of Te is defined to be any randomised-strategy profile for the multi-agent representation of Te (i.e. a probability distribution over a set of possible moves). A behavioural-strategy specifies a move probability following every decision node. A belief probability is the probability that a player assigns to being at a particular node where there are multiple nodes with the same information state.

An information state s and a pure strategy ci are compatible iff there exists at least one combination of pure strategies for the other players such that a node at which player I moves with information state s could occur with positive probability when I implements strategy ci.

Theorem 4.1 If Te is a game with perfect recall then any two mixed strategies in €Ci that are behaviourally equivalent are also pay-off equivalent.
A Nash equilibrium of an extensive-form game, or an equilibrium in behavioural strategies, is defined to be any equilibrium e of the multi-agent representation such that the mixed-representation of e is also an equilibrium of the normal representation.

Theorem 4.2 if Te is an extensive-form game with perfect recall and t is an equilibrium of the normal representation of Te, then any behavioural representation of t is an equilibrium of the multi-agent representation of Te.

Theorem 4.3 For an extensive-form game Te with perfect recall, a Nash equilibrium in behavioural strategies exists.

A strategy for player i is sequentially rational for player i at information state s if i would actually want to do what this strategy specifies for him at s when information state s actually occurred.

A beliefs vector v is weakly consistent with scenario s iff the belief-probability distribution is equal to Bayes's formula. This is only relevant for states which are "on path" i.e. states that occur with positive probability. The extensive diagram adds a move probability in () on each link and a belief probability in <> on each node.

Examining rational behaviour at nodes of zero probability is important because an irrational move may place the game on such a path (for instance a weakly dominated strategy but it is also possible without a weakly dominated strategy).

A weak sequential equilibria of Te is where the behavioural-strategy profile is sequentially rational for every payer at every information state with a weakly consistent beliefs vector. This can exclude some unreasonable equilibria.

However, we also need to cope with information states with zero probability. This is done by giving all nodes a small positive probability and allowing this to approach zero. Small initial doubts can have a major impact on rational players behaviour in multi-stage games.

The concept of sub-game-perfect equilibria for extensive form games is weaker than sequential equilibria but remains useful for infinite pure-strategy sets. A subgame-perfect equilibria is any equilibria in behavioural strategies such that every subgame is also an equilibria in behavioural strategies.

A game with perfect information is any extensive form game in which each information state of each player occurs at exactly one decision node. The concepts of subgame-perfect equilibria and sequential equilibria coincide for these games.

The concept of sequential equilibrium is very sensitive to the inclusion of new low-probability chance events in the structure of the game.

A zero-probability chance event is not an impossible event if a rational individual might infer that this event had positive probability after some players made moves that supposedly had zero probability. This distinction is fundamental to the whole concept of sequential equilibrium.

Sequential rationality and subgame-perfectness are backward induction principles for the analysis of extensive-form games and rely on players anticipating other players moves at the end of the game. Forward induction asserts that the behaviour of rational intelligent players may depend on the options that were available in the earlier part of the game. Unfortunately some natural forward-induction is incompatible with other natural backward-induction arguments. Backward-induction corresponds in the normal representation to iteratively eliminating weakly dominated strategies in one particular order while forward-induction is in another order.

Sophisticated voting in a binary agenda results in any one of the outcomes in the top cycle dependent on the chosen agenda.

5 Refinements of Equilibrium in Strategic Form.

Different games in extensive form with different sequential equilibria can have the same strategic representation. So if sequential equilibrium is accepted as an exact solution then we must use the extensive form for analysis.

However, a sequential equilibria can sometimes use a weakly dominated strategy which is questionable. A pure equilibrium extends a Nash equilibrium (each players equilibrium is a best response to all other players equilibrium) by allowing arbitrarily small perturbations to all players strategies such that each pure strategy gets a strictly positive probability).

A perfect equilibrium of an extensive form game is any perfect equilibrium of the multi-agent representation. By contrast a Nash equilibrium is an equilibrium for both the multi-agent and normal form. Because a perfect equilibria of the normal form may not correspond to a sequential equilibria.

Theorem 5.2 For any finite game in strategic form, there exists at least one equilibrium.

A proper equilibrium is similar to a perfect equilibrium (that assigns an arbitrarily small probability to mistaken pure strategies) except it assigns a much less probability to mistaken pure strategies (by a multiplicative factor arbitrarily close to zero) than to other pure strategies that would be either a best response or less costly mistake.

A perfect equilibria can be viewed as a way of identifying sequential equilibria in the multi agent representation. A proper equilibrium can be used to identify sequential equilibria in the normal representation.

Perfect equilibrium can be identified that do not correspond to a sequential equilibrium.

In general, adding or eliminating a pay-off equivalent pure strategy makes no essential change to the proper equilibrium.

However, adding a pure strategy that is randomly redundant can change the set of proper equilibrium. This is an inevitable consequence of the fact that two extensive form games that have the same fully reduced normal representation may have disjoint sets of sequential equilibrium scenarios.

A persistent equilibria is any equilibrium that is contained in a persistent retract. A persistent retract is a minimal (i.e. it contains only one) absorbing retract. An absorbing retract is a retract where every randomised strategy profile that is in the neighbourhood has a best response in the retract. A retract is a non-empty closed convex subset of strategies.

A Stable Set of equilibria is a minimal pre-stable set of equilibria A pre-stable set of equilibria is a closed subset of equilibria in the neighbourhood of a Nash equilibrium Each player is given an independent probability e of accidentally implementing the randomised strategys in the neighbourhood

Theorem 5.6 - For any finite strategic form game T, there is at least one connected set of equilibria that contains a stable subset.

Theorem 5.7 - If S is any stable set of the game T, then S contains a stable set of any game than can be obtained from T by eliminating any pure strategy that is weakly dominated or that is an inferior response (that is, not a best response) to every equilibrium in S.

Generic games and the associated terminology was summarised on pg 239

Conclusions - we don't have the definitive answer yet!

A refinement of a Nash equilibrium is a solution concept intended to offer a more accurate characterisation of rational intelligent behaviour. A selection criteria is an objective standard used to determine a focal equilibrium that everyone expects.

6. Games with Communication

While all communications can be included as part of the explicit game definition in a general theoretical way it introduces practical difficulties.

A game with contracts includes both rich communication and a contract between players to adopt particular strategies. A correlated strategy for a set of players is any probability distribution over a set of combinations of pure strategies that these players can choose. An allocation is a vector setting out the pay-off for each player. A minimax value for a player is the best expected pay-off that he would get against the worst set of correlated strategies that could be used against him. A correlated strategy is individually rational iff each player gets not less than their minimax value. These are called participation constraints.

Correlated equilibrium are correlated strategies which are self-enforcing because no player can unilaterally gain by deviating from the plan.

The use of a mediator can enable the implementation of even more profitable but otherwise not self-enforcing strategies. In this arrangement, each player advises the mediator of their type and the mediator advises each player the recommended action. Noisy communication can have a similar effect! Also it has been shown that in any strategic-form or and Bayesian game with four or more players, a system of direct unmediated communication between pairs of players can simulate any centralised communication system with a mediator, provided the communication between pairs is private.

It is worthwhile to consider only mediated communication systems where it is rational for all players to obey because such systems can simulate any equilibrium of any game that can be generated from any give strategic-form game by adding any communications system.

The Revelation principle: "any equilibrium of any communication game that can be generated from a strategic-form game T by adding a system for preplay communication must be equivalent to a correlated equilibrium satisfying the strategic incentive constraints.

Babbling equilibria is where every player ignores all communication.

A probability constraint is the conditional probability that the mediator would recommend to each player i that he should use action ci if each player reported his type to be tj.

A mediation plan in incentive compatible iff it is Bayesian equilibrium for all players to report their type honestly. And obey the recommendations.

In the insurance industry the need to give players an incentive to report information (risk profile) honestly can be called adverse selection and the need to give players an incentive to implement recommended actions (protect property) can be called moral hazard.

A Bayesian collective-choice problem is a set of possible outcomes or social options that are jointly feasible for the group (rather than a set of actions for each player separately). A collective-choice plan or mechanism specifies the probability that a social option d would be the chosen outcome if t were the combination of reported types reported by the players. The mechanism is incentive compatible iff honest reporting by all players is a Bayesian equilibrium of the game. Hence the revelation principle applies. A disagreement outcome will occur if the players fail to agree on the mechanism. A Bayenesian bargaining problem is a collective choice problem with specification of a disagreement outcome. A mechanism is feasible if it I incentive compatible and individually rational.

Trading problems with Linear Utility allow us to model trading mechanisms where the players have an incentive to lie about their private value for an item (i.e. a seller will state a higher value and a buyer will state a lower)

There are auction mechanisms which maximise the return by the auctioneer by encouraging a buyer to report their true value up to 1.5 times that of the next highest bidder. See additional reading Meyerson 1981a.

It can be shown that in bilateral trade (only 1 buyer and 1 seller) where the private values are drawn from an interval 0-1 that the expected difference between the buyers value and the sellers value must be at least 1/2. For this situation to be post efficient the expected difference between the values must be 1/3. Therefore there is no ex post efficient incentive-compatible mechanism.

A trading mechanism is post efficient if it always allocates all of the commodities being traded to the individuals who value them the most highly.

In trading models with many buyers and sellers it can be shown that the vast majority of traders will get to trade. In addition, uncertainty in the direction of trade can help make post efficient mechanisms feasible by creating countervailing incentives.

Sender-Receiver games are simple 2 player Bayesian games with moral hazard and adverse selection. Player 1 (sender) has private information but no choice of actions while player 2 (receiver) has the choice of actions but not private information.

No substantive communication can occur in sender-receiver games where they are restricted to perfectly reliable noiseless communication. However, more favourable equilibriums are possible when unreliable or noisy communication is introduced.

In strategic-form games with communication there are Acceptable and Predominant Correlated Equilibria. An e-correlated strategy is the probability that the mediator recommends a strategy but a few players tremble and implement the e-correlated strategy. There are then e-correlated equilibrium and acceptable e-correlated equilibrium which is the limit as the probability of tremble e goes to zero. An acceptable strategy is one that can be rationally used by a player when the probabilities of tremble are arbitrarily small. The acceptable residue is the game that remains after the removal of all unacceptable strategies. However there may be some strategies that were acceptable in the original game but unacceptable in the context of the acceptable residue so we have an iterative process resulting in the predominant residue with the iteratively acceptable or weakly predominant strategies and finally the predominant correlated equilibrium.

"There is a conceptual trade-off between the revelation principle and the generality of the strategic form. If we want to allow communication opportunities to remain implicit at the modelling stage of our analysis, then we get a mathematically simpler solution concept, because (by the revelation principle) communication equilibria can be characterized by a set of linear incentive constraints. However, if we want to study the game in strategic form, then all communication opportunities must be made an explicit part of the structure of the extensive-form game before we construct the normal or multi-agent representation."

7. Repeated Games

In repeated games it is important to have an infinite time horizon, or at the very least for the players not to know which will be the last repeat because each player will know that the other players will treat the last game as they would any other single game instance. So the 2nd last repeat is actually the last in which case each player knows how the other will react etc. etc. until it is merely a sequence of independent games.

Repeated games have bounded pay-offs when the utility pay-offs do not exceed some positive limit. A State is absorbing if once it occurs in round k then all following rounds will have the same state. If there is an absorbing state which has zero pay-off then we can use the sum-of-pay-off criteria, otherwise the d-discounted average pay-off is the appropriate criteria. Another criteria is the limit of average pay-offs (which sometimes do not exist). The liminf is the lowest number that is a limit of some convergent sub-sequence while the limsup is the highest number that is a limit of some sub-sequence. Another option is the overtaking criterion where the mliminf of the difference of corresponding rounds is > 0. The Banach criteria is another more satisfactory but technical approach.

The discounted average of a sequence of pay-offs beginning at round 1 is the weighted average of the pay-off at round 1, with weight (1-d) and the d -discounted average of the sequence of pay-offs beginning at round 2, with weight d.

A repeated game has complete state information iff at every round every player knows the current state of nature. A behavioural strategy is stationary iff the move probabilities only depend on the current state. Many strategies that do not at first appear stationary may become stationary in an equivalent model with a larger state space (e.g. The Big Match).

A repeated game with standard information, or a standard repeated game (super-game) is a repeated game in which there is only one possible state of nature and the players know all of each others past moves. This represents where groups of individuals face the same competitive situation infinitely often and always have complete information about past behaviour. The stationary equilibria are the equilibria of the one-round game but the non-stationary equilibria are generally much larger sets due to the players ability to punish other players past deviation from supposed equilibrium paths. In fact, almost any equilibrium pay-off allocation that gives each player at least his minimax value can be achieved in an equilibrium of the standard game. For example the Chicken game where the tit-for-tat strategy is very effective (but not sub-game perfect) or modifications such as getting even, the grim strategy or mutual punishment see page 327)

The prisoners dilemma is quite unusual in the difference in equilibrium for a single game and for the repeated game. A finitely repeated game is attractive if introducing a small positive probability of the player being a machine that uses this strategy would substantially change the set of sequential equilibria.

When other players moves are imperfectly observable (examples jointly avoiding accidents and jointly pursuing sales) the likelihood ratio measures how informative the new information is about the unknown (i.e. effort of other players)

Generous behaviour which is sustained in small groups where everyone can recall everyone else's past behaviour is typically unsustainable in large decentralised groups. A mediator can resolve this by acting as a memory for the large group. (sounds like the anonymity of the internet)

In a zero sum two person game the convexity or concavity of the equilibrium pay-off function indicates whether the informed player should conceal or reveal their knowledge.

Study of games in continuous time is difficult (compared to discrete intervals or rounds).

In the evolutionary approach, the tit-for-tat strategy has almost zero resistance for prisoners dilemma but 4/11 resistance for chicken.

8. Bargaining and Cooperation in Two-Person Games.

Nash argued that fundamentally cooperation can be analysed using the same basis of each person increasing their utility during some bargaining process. However, to represent every possible contractual option will result in a very large game with very many equilibria The focal-point effect becomes important.

One approach is to assume that people can cooperate effectively when they use preplay communication to coordinate their expectations on focal point equilibria. These depend on arbitration, negotiation and welfare properties.

An impartial arbitrator would try and base their selection on some objective principals. So we are developing a theory of impartial arbitration. The welfare properties of equity and efficiency may determine the focal equilibrium (regardless of the presence of an arbitrator). Preplay communication is called focal negotiation. The focal arbitrator is in a way a player and is called the principal.

The equity hypothesis: (In summary) the outcome of a negotiation where the players can participate equally should not depend on the presence of an arbitrator (with access to common knowledge)

A two person bargaining problem is essential iff there is at least one allocation that is strictly better for both players than the disagreement allocation.

Three ways to determine the disagreement point. The minmax value for each player, the utility of some focal point equilibrium or the result of some rational threat.

Axiom 8.1 Strong Efficiency: The solution should be feasible and efficient. That is, there should be no other feasible allocation that is better for one player and not worse for the other. Something can be strongly and weakly efficient.

Axiom 8.2 Individual Rationality: Neither player should get less in the bargaining solution that they would in disagreement.

Axiom 8.3 Scale covariance: WE can move between decision theory equivalent utility scales and still get the same outcome.

Axiom 8.4 Independence of irrelevant alternatives. Elimination of feasible alternatives (other than the disagreement solution) which would not have been chosen will not effect the outcome.

Axiom 8.5 Symmetry: If the position of the players are symmetric then the solution should be symmetric.
Nash's remarkable result is that there is exactly one bargaining solution, called the Nash bargaining solution.

Interpersonal comparisons of utility can include the egalitarian solution (principal of equal gains) and the utilitarian solution (greatest good). But in any essential game there exists a solution that satisfies both of these requirements and it is the Nash bargaining equilibrium!

Transferable utility allows the players the option to give any amount to the other player or even to simply destroy money (utility). Where there is transferable utility the two-person bargaining problem can be characterised by three numbers, the disagreement pay-off to player 1, the disagreement pay-off to player 2, and the total transferable wealth.

Rational Threats. Notice in the Nash equilibrium that the pay-off to player 1 increased as the disagreement pay-off to player 2 decreases. That is the possibility of hurting player 2 in the event of a disagreement pay-off can help player 1 is a cooperative agreement is reached!

The equilibrium theory of disagreement is appropriate when players could not commit to a particular disagreement action plan until a disagreement actually occurs. The rational threats theory is appropriate when each player can before the arbitration process commit to a single disagreement action plan in the event of a disagreement (which almost certainly will not occur). The minimax- values theory is appropriate when each player can commit to an offensive and a defensive strategy to use in the event of a disagreement.

Other bargaining solutions have been explored. The non-symmetric Nash bargaining solution drops the symmetry axiom (i.e. Two family where utility is not split 50/50 but by the number of children). Another solution satisfies strong efficiency, scale covariance, symmetry and individual monoticity. There are others....

Alternating-offer bargaining games were considered by Stahl and Rubinstein who found that sub-game perfect equilibria are unique when each players cost of time is given by some discount factor δ. Another way of characterising the cost-of-time is by introducing a exogenous positive probability, after each round that the bargaining process may permanently terminate in disagreement The bargainers final shares in this game depend on the ratio of the players power of commitment. This is depicted by the negotiators ability to convince the other that this is their final offer but to keep the other at the table after rejecting their last offer.

Introducing incomplete information to the alternating offer game reduces the importance of the relative power of commitment. If a player can convince the other that they have an irrational commitment to a particular allocation then he can expect to get an agreement close to or better than this allocation.

When the alternating offer game has discrete values (i.e. dollars) then the focal-point effect becomes the determining factor where in pre-play communication a player will suggest a rational allocation and then suggest that any deviation from this rational allocation can only be made by an irrational player.
Renegotiation.

9. Coalitions in Cooperative Games.

We cannot just scale up from two player games to N player games because of the impact of coalitions.

A coalition of players can negotiate effectively when there is a feasible change in the strategy of members of the coalition that would benefit them all then they would agree to make this change (unless is contradicts some prior agreement with other players)

Consider a game with 3 players who all get 0 unless there is some pair of players who can propose and allocation for themselves up to a combined maximum of 300. The third player will always be able to make a coalition offer to one of the existing coalition pair which is better than what they will get in the existing coalition.

The simplifying assumption of transferable utility is often used to simplify analysis of n-player games. Characteristic functions describe the cooperative possibilities of a game and may be built from the minimax representation (the best guarantee against the worst offensive threat), the defensive equilibrium representation (both coalitions play defensively), or the rational threats representation (). If all these three representations coincide then the game has orthogonal conditions.

The minimax representation is by its nature super-additive – meaning that the whole is greater than the sum of its parts. The super-additive cover is the maximum maximum worth that a coalition could achieve by breaking up into a set of smaller disjoint coalitions.

A players power is his ability to help or hurt any set of players by agreeing to cooperate or refusing to do so. A characteristic function is a summary description of the power structure of the game.

An allocation is feasible in a coalition if the utility available to each player is ≥ that which they can get by themselves. A coalition can improve on an allocation if there is a feasible alternate allocation where all players get a strictly higher pay-off An allocation is in the core if it is feasible and cannot be improved.

Unfortunately the core may be empty and may seem extreme at times (e.g. 1,000,000 left glove suppliers and 1,000,001 right glove suppliers ∴ right glove suppliers have 0 value)

The instability of cores in large games can be mitigated by considering ε-cores. IF x is in the ε-core then no coalition could guarantee all members more than ε above x.

The logical appeal of the core is based on assumptions that when another coalition S contemplates blocking an allocation x then 1. they are not prevented by prior commitments 2. there agreement would be final and 3. if they do not agree they really will get x. These assumptions are questionable in small games but more robust in large games.

There is a unique Shapley value which satisfies 3 axioms (carriers of a coalition game add value to a coalition while dummies do not – the joint worth is divided only between carriers, not dummies – the expected pay-off to each player does not depend on the resolution of uncertainty as to the game to be played). The Shapley value of any player is his expected marginal contribution when he enters the coalition (based on random order of entry). Alternatively the Shapley value of each player depends on the differences between the worths of complimentary coalitions. The rational-threats may the the appropriate form to use when using the Shapley value. The Shapley value has also been extended to games with infinite players.

The Shapley value has been extended to recognise that some cooalitions can be more effective than others. Owen proposed a nested coalition structure so that members of a coalition can divide their worth unequally.

Other solution concepts: The excess of a coalition is the utility remaining after each member has been paid the value that they would have got had they not been in the coalition.

Coalition games with non-transferable utility NTU define unpreventable coalitions as those which can guarantee a strictly better pay-off to all members after the non-coalition players announce their strategy.
The inner core of a NTU game introduces randomisation of coalitions. An allocation x is strongly inhibitive if there exists no viable randomised blocking plan against it. An allocation is inhibitive if an arbitrarily small perturbation we would get a strongly inhibitive allocation. The inner core is the set of all inhibitive allocations.

The Shapley value and other solution concepts can be defined for NTU games. The Harsanyi NTU value and Owen values are pretty interesting!

Roth and Shafer have studied games with counter-intuitive Shipley values (due to the smoothness assumption being violated)

10. Cooperation under Uncertainty.

This chapter considers how players not only compromise with each other but also how a single player should compromise between the goals of his true type and the goals of his other possible types, to maintain an inscrutable façade in negotiations.
A mechanism is ex anti efficient / interim efficient / ex post efficient iff there is no other mechanism that is ex anti Pereto superior / interim Pereto superior / ex post Pereto superior respectively.
492

Additional Reading

Myerson, R, B 1981 "Optimal Auction Design" Mathematics of Operations Research 6:58-73.