Conditional Game Theory:

A Primer

Andre Hofmeyr[1]
Don Ross[1,2]
[1]University of Cape Town
[2]Georgia State University

Introduction

- Wynn Stirling (2012) formulated a theory of conditional games with the avowed aim of developing a concept of group preference, which is not simply an exogenous aggregation of individual preferences, but which arises endogenously as social influences propagate through a group.
- Stirling's framework is a strict generalisation of orthodox, non-evolutionary game theory that incorporates the influence of social bonds through the technology of *conditional preferences*.
- In this primer, I will outline the basic ingredients of the theory and apply it to the Prisoners' Dilemma and Hi-Lo to show how the framework enriches the analysis of these games.

A Running Example

- To illustrate the intuition of the framework, we employ an example due to Ross (2014).
- Consider a Board of Directors that must decide whether to engage in a risky hostile takeover bid.
- There are at least two ways in which the views of the Board can be elicited:

1. The Chair sends out a detailed risk analysis of the costs and benefits of the proposed takeover prior to the Board meeting.
2. The Chair, citing security concerns, presents the same information to the Board but only after they have assembled in the Board room.

- The question of interest is whether these two processes should be expected to yield the same outcome.

Conditional vs Unconditional

- Process (1) encourages Board members to form unconditional preferences prior to the meeting, which they may defend against the other members' arguments.
- Process (2), by contrast, may induce members to monitor one another while they decide which option is best and may lead them to modulate their preferences on the basis of the preferences of others.
- Under both processes, differing individual preferences are likely to be expressed through non-unanimous votes.
- But the distribution of these preferences might vary across the two scenarios because process (2) encourages revelation of preferences that are influenced by information about the preferences of others, which affords opportunity for preference calibration.

Conditional Preferences and Influence Flows

- The starting point of Stirling's analysis is the distinction between *categorical* and *conditional* preferences.
- Categorical preferences *unconditionally* define an agent's ranking of all possible outcomes, regardless of other agents' preferences.
- Conditional preferences are based on influence flows which propagate through a group and define agents' rankings of alternative outcomes as *conditional* on the preferences of others.
- This propagation of influence flows, which is modelled using graph theory, defines a social model that enables agents to jointly consider individual and group interests, but without requiring us to leave the Nash (i.e., mutual best response) constraint.

The Chair and The Board Member

- Building on the earlier example, but simplifying to the case of the Chair and one Board member, assume that each player has two actions:

1. Support (S) the takeover bid.
2. Do not support (NS) the takeover bid.

- Thus, the outcome space for this game is (S, S), (S, NS), (NS, S) and (NS, NS), where the Chair's action is listed first and the Board member's action is listed second.
- Assume further that the Chair has categorical preferences over the action profiles but the Board member's preferences may be influenced by the Chair's.
- Specifically, suppose that if (S, S) is the Chair's optimal profile, the Board member will define his ranking of the alternatives on the basis of this hypothesis whereas if (NS, NS) is the Chair's optimal profile then the Board member may define a different preference ordering.
- Given the four possible outcomes of this game, the Board member can define different preference orderings which are conditional on his conjecture concerning the preference ordering of the Chair.

Concordance

- As social influence propagates through a group and players modulate their preferences on the basis of other players' preferences, a group preference may emerge.
- This notion does not directly provide the basis for action, but rather serves as a social model which incorporates all of the relationships and interdependencies that exist among the agents.
- Stirling refers to this concept as *concordance* and it captures the extent to which a conjectured set of (categorical or conditional) preferences yield controversy within a group.
- Crucially, concordance does not refer to the goals of a group nor to the goals of the agents who comprise it, but rather to the level of discord that hypothetical propositions concerning players' preferences engender among members of the group.

Conjectures and Concordance

- For example, consider the following joint conjectures for the Chair and Board member:
- $\mathbf{a_1}$ = {(S, S), (S, NS)} and $\mathbf{a_2}$ = {(S, S), (NS, NS)}
- Assume that under $\mathbf{a_1}$, the Chair's conjecture (S, S) is best for her and next-best for the Board member, while the Board member's conjecture (S, NS) is best for him but next-best for the Chair.
- By contrast, assume that under $\mathbf{a_2}$, the Chair's conjecture is, once again, best for her and next-best for the Board member while the Board member's conjecture (NS, NS) is worst for both players.
- Which conjecture is likely to entail a great level of controversy among the players?

Conjectures and Concordance

- $\mathbf{a_1}$ = {(S, S), (S, NS)} and $\mathbf{a_2}$ = {(S, S), (NS, NS)}
- The joint conjecture $\mathbf{a_1}$, involves different conjectures by the two players but they do not include the players' worst outcome.
- The joint conjecture $\mathbf{a_2}$, by contrast, incorporates a conjecture (S, S) that might be satisfactory to either player but one (NS, NS) which is the worst for both players.
- Consequently, we might expect $\mathbf{a_2}$ to produce more severe dispute among the players than $\mathbf{a_1}$, and an ordering over these joint conjectures that is sensitive to these varying levels of controversy encodes the concept of concordance.
- The level of concordance varies with the specific strategic interaction under study.
- In games where players' interests are perfectly aligned, the extent of controversy will be minimised when players conjecture identical action profiles.
- In zero-sum games, by contrast, a low degree of controversy is more likely when conjectures are diametrically opposed.
- As Stirling (2012, p. 40) notes, "… even antagonists can behave concordantly."

Endogeny

- While the concept of concordance may provide the basis for an emergent notion of group preference, its value derives from the extent to which it is determined by the individuals who make up a group.
- In other words, concordance is not imposed exogenously on a group from the outside but is determined by the social linkages and influence flows among members of the group.
- Stirling refers to this principle as *endogeny*.
- It is among the building blocks of his *aggregation theorem*, which in turn provides a model of the social relationships and interdependencies of members of a group, and a device for simultaneously representing individual and group agency.

A Praxeological Framework

- To develop a concordant ordering which respects the principles of conditioning (i.e., that players' preferences may be conditional on the preferences of others) and endogeny, Stirling employs the logic of multivariate probability theory in a *praxeological context*.
- He urges us to understand praxeology on the basis of an analogy with epistemology.
- Where epistemology is concerned with the nature and scope of knowledge and classifies propositions on the basis of their veracity, praxeology classifies propositions on the basis of their efficacy and efficiency.

Probability Theory

- In probability theory, given a set of two discrete random variables {*X*, *Y*}, the conditional probability mass function

  $p_{Y|X}(y \mid x) = P(Y = y \mid X = x)$ is a measure of the likelihood that the random variable $Y = y$ given that, or conditional on, the random variable $X = x$.
- The conditional probability mass function is defined as the ratio of the joint probability of *X* and *Y* and the marginal probability of *X* or: $p_{Y|X}(y \mid x) = \mathrm{p}_{X\,Y}(x, y) / \mathrm{p}_X(x)$.
- Solving this expression for $\mathrm{p}_{X\,Y}(x, y)$ as the subject of the formula (i.e., $p_{X\,Y}(x, y) = p_{Y|X}(y \mid x) \times \mathrm{p}_X(x)$) it is clear that the joint probability of *X* and *Y* can be derived from the conditional probability of *Y* given *X* and the marginal probability of *X*.
- In other words, probability theory provides a framework for combining information from different sources - in this instance, the conditional probability of *Y* given *X* and the marginal probability of *X* - to determine the joint likelihood of an event.

A Concordant Ordering for the Group

- In a praxeological framework, Stirling's goal is to derive a concordant ordering for the group which combines the conditional and categorical preferences of its members, in much the same way as the joint probability of an event is determined by conditional and marginal probabilities.
- Working directly with preference orderings quickly becomes cumbersome, so Stirling seeks to derive utility functions that represent the players' categorical and conditional preferences and the groups concordant ordering.
- The existence theorem for a utility function which represents categorical preferences is well known so we will focus on the derivation of a conditional utility function and the principles which must hold so as to permit aggregation of categorical and conditional preferences to derive a concordant utility function.

Categorical and Conditional Utility Functions

- Let $\{X_1, \ldots, X_n\}$, $n \geq 2$, represent a set of n players, and let A_i denote a finite set of actions available to player i from which he or she must choose one element to instantiate.
- An action or strategy *profile* is an array:
  $\mathbf{a} = (a_1, \ldots, a_n) \in A_1 \times \ldots \times A_n$.
- Under classical game theory, players have categorical utility or payoff functions defined over strategy profiles:
  $u_i : A_1 \times \ldots \times A_n \rightarrow \mathbb{R}$.
- In the context of conditional preferences it is useful to define the parent set $pa(X_i) = \{X_{i1}, \ldots, X_{in}\}$ as the n_i-element subset of players whose preferences influence X_i's preferences.
- Assume that X_{ij}, the j^{th} parent of X_i, forms the hypothetical proposition that profile $\mathbf{a}_{ij}$ will occur.
- This hypothetical proposition is termed a *conjecture*.

The Conditional Utility Function

- Thus, let $\mathbf{a}_i = \{\mathbf{a}_{i1}, \ldots, \mathbf{a}_{in}\}$ represent the *joint conjecture* of $pa(X_i)$.
- Then there exists a function which maps action profiles, *conditional* on the joint conjecture of $pa(X_i)$, to the real line $\mathbb{R}$, which represents X_i's preferences:

  $u_{X_i \mid pa(X_i)}(\cdot \mid \mathbf{a}_i) : A_1 \times \ldots \times A_n \rightarrow \mathbb{R}$.
- Note that if $pa(X_i) = \varnothing$ then the conditional utility $u_{X_i \mid pa(X_i)}$ becomes the categorical utility u_i.
- Given the existence of a conditional utility function which represents players' conditional preferences, the collection $\{X_i, A_i, u_{X_i \mid pa(X_i)}, i = 1, \ldots, n\}$ constitutes a finite, normal form, noncooperative *conditional game*.
- Returning to our earlier example of the Chair (C) and the Board member (B), the conditional game consists of two players $\{X_C, X_B\}$, each with two actions $A_i = \{S, NS\}$, and the utility functions $u_C(\mathbf{a}_C)$ and $u_{B|C}(\mathbf{a}_B \mid \mathbf{a}_C)$, for the Chair and Board member, respectively.

Praxeology-Epistemology Analogy

- Note that through appropriate normalisation one can ensure that all utilities (i.e., categorical and conditional) are non-negative and sum to unity, which implies that the utilities have all of the characteristics of probability mass functions.
- As discussed earlier, in an epistemological framework marginal and conditional probabilities can be combined to determine a joint probability: $p_{XY}(x, y) = p_{Y|X}(y \mid x) \times p_X(x)$.
- Consequently, if the praxeology-epistemology analogy is appropriate, it may be possible to aggregate the conditional and categorical utilities to define a group utility function that incorporates the social linkages and interdependencies of members of a group and thereby represents the level of concordance of the group.
- The value of this analogy is that it lets us apply concepts from multivariate probability theory, such as Bayes' rule and marginalisation, in a praxeological context and derive game-theoretic solution concepts that incorporate both individual and group interests.

Three Further Principles

- Returning to our earlier example, the goal is to combine the categorical preferences of the Chair with the conditional preferences of the Board member to produce an emergent preference ordering for the group.
- The requirement is to prove that the group or concordant utility $U_{CB}(\mathbf{a}_C, \mathbf{a}_B) = u_{B|C}(\mathbf{a}_B \mid \mathbf{a}_C) \times u_C(\mathbf{a}_C)$.
- In words, the concordant utility *U* is the product of the Board member's conditional utility and the Chair's categorical utility.
- In assembling the basis for such a proof, Stirling adopts three further assumptions or *principles* in addition to endogeny and conditioning:

1. Acyclicity
2. Exchangeability (framing invariance)
3. Monotonicity

Acyclicity

- Acyclicity means that no cycles can occur in the social influence relationships among the players.
- In other words, if the Chair influences the Board member, then the Board member cannot influence the Chair.
- The problem with cyclical influence relationships is that they raise the possibility of indirect self-influence: the Chair influences the Board member, who in turn influences the Chair, which leads to a non-terminating cycle.
- An implication of acyclicity is that influence relationships are hierarchical and that at least one player in a strategic interaction must posses categorical preferences.
- Another implication is that social influence relationships can be represented using a directed acyclic graph (DAG)

Exchangeability

- Exchangeability, which Stirling and Felin (2013) refer to as *framing invariance*, requires that if a strategic interaction can be framed in different ways but there is no loss of information under different framings, then all framings must produce an identical concordant ordering.
- What this principle implies is that players must be willing to take into consideration the preferences of others when defining their own preferences, even if only to a small degree, and that the same information is available to players under alternative framings.
- This is a natural restriction in an epistemological context but for framing invariance to hold in a praxeological context, the concordant utility must satisfy the following conditions:

  $U_{CB}(\mathbf{a}_C, \mathbf{a}_B) = u_{B|C}(\mathbf{a}_B \mid \mathbf{a}_C) \times u_C(\mathbf{a}_C) = u_B(\mathbf{a}_B) \times u_{C|B}(\mathbf{a}_C \mid \mathbf{a}_B) = U_{BC}(\mathbf{a}_B, \mathbf{a}_C)$.
- In words, the concordant utility U_{CB}, which combines the conditional preferences of the Board member and categorical preferences of the Chair, must be the same as the concordant utility U_{BC}, which combines the categorical preferences of the Board member and the conditional preferences of the Chair.

Monotonicity

- The final principle that is required to derive a concordant utility function that has all of the characteristics of a joint probability mass function is *monotonicity.*
- This is a natural restriction on the concordant utility function, which ensures that no individual's preferences will be arbitrarily subjugated by the group.
- Specifically, if an individual or subgroup prefers option A to B and other players are indifferent among them, then the group must not prefer B to A.
- Thus, if the Chair prefers S to NS and the Board member is indifferent, the group must not prefer NS to S.

The Aggregation Theorem

- Stirling (2012, p. 59-60) proves that if the principles of *conditioning*, *endogeny*, *acyclicity*, *exchangeability*, and *monotonicity* hold, then a concordant utility function exists that represents the social relationships of the group, and is derived from the conditional and categorical utility functions of its members.
- The most general form of the concordant utility function is:
  $U_{X_1 \ldots X_n}(\mathbf{a}_1, \ldots, \mathbf{a}_n) = \prod^{n}_{i=1} u_{X_i \mid pa(X_i)}(\mathbf{a}_i \mid \mathbf{a}_i)$
- This expression shows that the concordant utility function, which combines information in a praxeological domain, has the same syntax as a joint probability mass function that combines information in an epistemological domain.
- Consequently, the full power of multivariate probability theory (particularly Bayes' rule and marginalisation) can be applied in a praxeological context to determine effective and efficient action when social influences propagate through a group.

Marginalisation

- Marginalisation is an important operation in a praxeological domain because it allows the analyst to extract players' *ex post* preferences once social influence has permeated the group.
- A player's *ex post unconditional* preferences are extracted in the following manner:

  $u_{X_i}(\mathbf{a}_i) = \Sigma_{\sim \mathbf{a}_i} U_{X_1 \ldots X_n}(\mathbf{a}_1, \ldots, \mathbf{a}_n)$,

  where $\Sigma_{\sim \mathbf{a}_i}$ means that the sum is taken over all arguments except $\mathbf{a}_i$.
- These *ex post* categorical utilities represent the players' preferences after taking into account the social relationships and interdependencies that exist in the group.
- As the preferences are unconditional, standard solutions concepts such as dominance and Nash equilibrium (NE) can be applied to them.

Meso → Macro → Micro

- The preceding discussion is summarised in the figure below.
- As social influences propagate through a group, players define their conditional preferences.
- Through the process of aggregation these social linkages and interdependencies lead to an emergent notion of group preference: concordance.
- Finally, through the process of marginalisation, the analyst extracts the players' *ex post* categorical preferences.

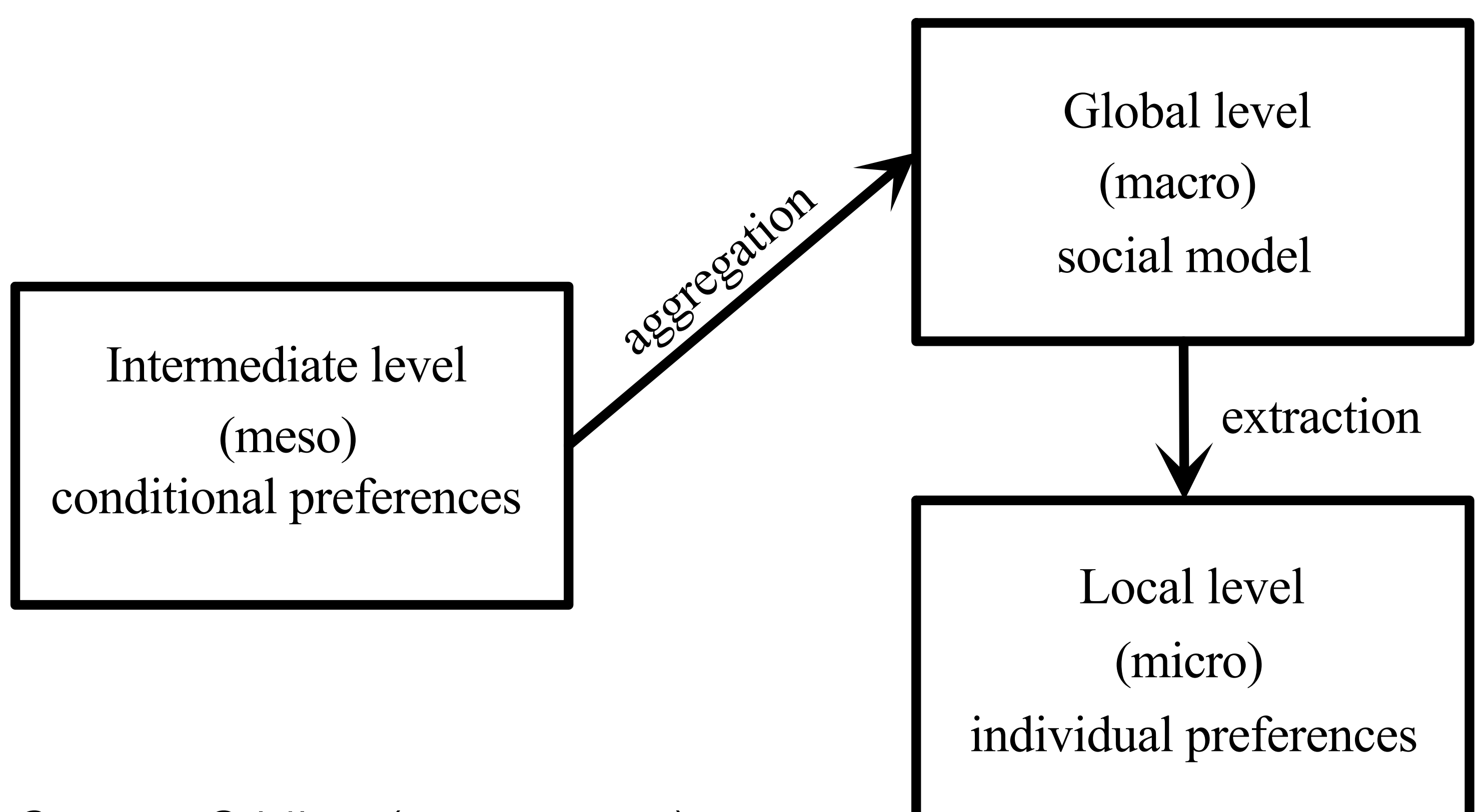

Source: Stirling (2012, p. 19)

Directed Acyclic Graph (DAG)

- Acyclicity implies that social influence relationships in conditional games can be modelled using a DAG.
- A DAG is a *graph* made up of *vertices* or *nodes*, which in a praxeological context represent the players, and *directed edges* or *links*, which capture the influence relationships between the players.
- If one player, C, influences another player, B, we write C → B, where C is referred to as the *parent* of B and B as the *child* of C.
- The set of parents of B is denoted *pa*(B) and the set of children of B is denoted *ch*(B).
- If a vertex has no parents *pa*(C) = ∅ then it is called a root vertex.

Chair and Board Member DAGs

- In the two-player DAG below, the Chair influences the Board member but, given acyclicity, the Board member does not influence the Chair.
- The Board member's conditional utility $u_{B|C}$ is represented by the edge between nodes C and B.
- In the three-player DAG, the Chair influences Board member B1 and Board member B2, and Board member B1 influences Board member B2.
- The influence flow between C and B1 is captured by the conditional utility $u_{B1|C}$ and the influence flows between C and B1 toward B2 are captured by the conditional utility $u_{B2|B1C}$.

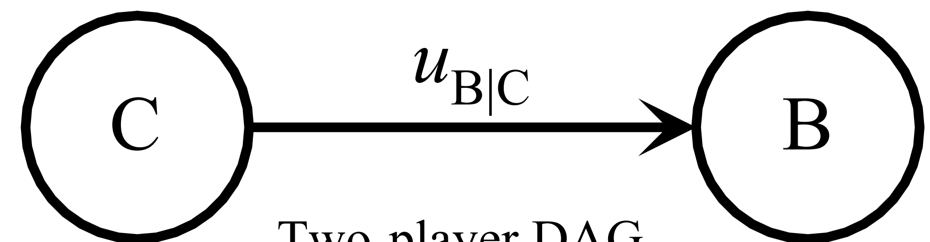

Two-player DAG

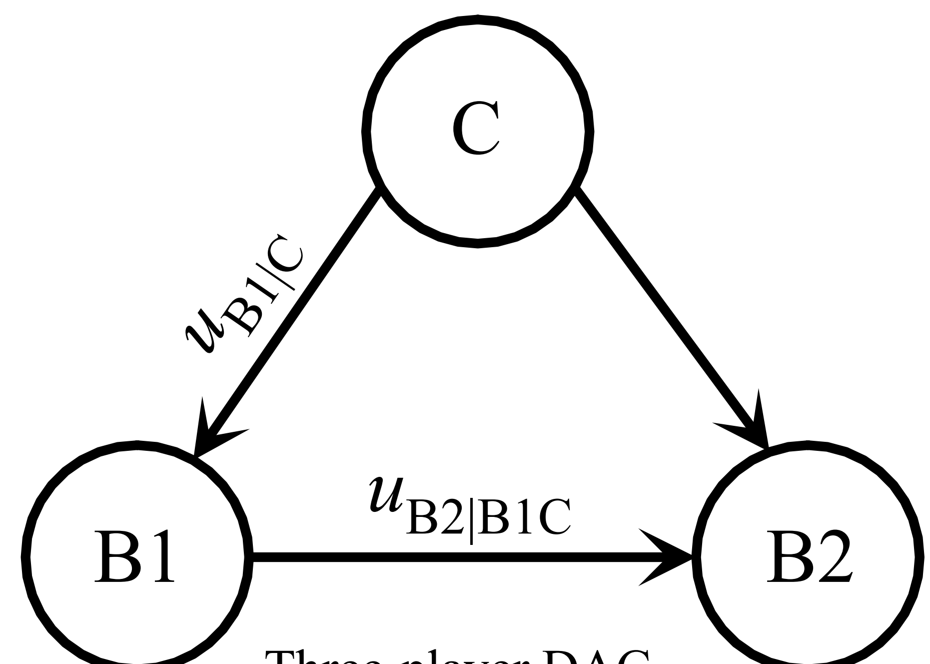

Three-player DAG

3-Player DAG: Concordant Utility

- The associated concordant utility for the three-player DAG is: $U_{CB1B2}(\mathbf{c}, \mathbf{b}_1, \mathbf{b}_2) = u_C(\mathbf{c}) \times u_{B1|C}(\mathbf{b}_1, \mathbf{c}) \times u_{B2|B1C}(\mathbf{b}_2 \mid \mathbf{c}, \mathbf{b}_1)$
- This expression combines information from the categorical and conditional utilities to define the concordant utility in much the same way that a Bayesian network, which can also be represented in a DAG, combines information from marginal and conditional probabilities to determine a joint probability.
- Thus, a DAG provides a graphical method to represent the influence flows, and associated conditional utilities, of a conditional game.
- The three-player DAG shows that B2 does not directly influence B1 and that neither B1 nor B2 directly influence C. However, this does not imply that B1 and B2 have no indirect influence on C.
- Recall that the exchangeability constraint means that a social model should be invariant to the way in which the information about linkages and influence flows is aggregated.

Exchangeability

- The exchangeability constraint implies that once the concordant utility has been defined, we can apply Bayes' rule to extract reciprocal influence relationships.
- Specifically, suppose that B1 conjectures $\mathbf{b}_1$ and we want to determine the influence of this conjecture on the Chair's preference for $\mathbf{c}$: $u_{C|B1}(\mathbf{c} \mid \mathbf{b}_1)$.
- The answer follows directly from Bayes' rule: $u_{C|B1}(\mathbf{c} \mid \mathbf{b}_1) = [u_{B1|C}(\mathbf{b}_1 \mid \mathbf{c}) \times u_C(\mathbf{c})] \,/\, u_{B1}(\mathbf{b}_1)$, where $u_{B1}(\mathbf{b}_1)$ is derived by marginalising the concordant utility.
- We can also determine the influence that B1 and B2 exert on C and the influence that B2 exerts on B1 by computing the appropriate conditional and categorical utilities using Bayes' rule and marginalisation.
- The crucial idea here is that once the concordant utility has been defined, exchangeability implies that many hierarchical structures are compatible with the social model of the group; in other words, the social model is framing invariant.

Group Utility Function

- Stirling then extends – as opposed to refines – the standard solution concepts of dominance and NE, to apply over group-level preference orderings.
- His approach is to extract a marginal utility for the group in much the same way as a marginal utility for each player was extracted from the concordant utility.
- A crucial assumption behind the procedure is that, given that players can only control their own actions, each player will make conjectures over her own action sets and not those of other players.
- Thus, let a_{ij} denote the j^{th} element of $\mathbf{a}_i$, where $\mathbf{a}_i$ is X_i's conjecture profile.
- Now form the action profile $(a_{11}, \ldots, a_{nn})$ by taking the i^{th} element of each X_i's conjecture profile.
- Finally, sum the concordant utility over all elements of each $\mathbf{a}_i$ except a_{ii} to form the group utility or welfare function:
  $V_{X_1 \ldots X_n}(a_{11}, \ldots, a_{nn}) = \Sigma_{\sim a_{11}} \ldots \Sigma_{\sim a_{nn}} U_{X_1 \ldots X_n}(\mathbf{a}_1, \ldots, \mathbf{a}_n)$

Marginal Welfare Functions

- The group does not act as a single entity and it cannot, therefore, instantiate its own preferred alternative, but the group utility provides a metric by which individual players determine the impact of their choices on the group.
- In the same way as players can extract their marginal utilities from the concordant utility function, they can extract their own individual marginal welfare functions from the group utility.
- Specifically, the marginal individual welfare function v_{X_i} of X_i is the i^{th} marginal of $V_{X_1 \ldots X_n}$:

  $v_{X_i}(a_i) = \Sigma_{\sim a_i} V_{X_1 \ldots X_n}(a_1, \ldots, a_n)$

Solution Concept

- The existence of group and individual welfare functions allows Stirling to derive a solution concept that formally integrates consideration of the interests of the group with consideration of the interests of the individual players.
- This solution concept relies on the maximum individual and group welfare solutions.
- The maximum group welfare solution is:

  $\mathbf{a^*} = \arg\max_{\mathbf{a} \in A_1 \times \ldots \times A_n} V_{X_1 \ldots X_n}(\mathbf{a})$
- The maximum individual welfare solution is:

  $a_i^{\S} = \arg\max_{a_i \in A_i} v_{X_i}(a_i)$
- If $a_i^{\S} = a_i^*$ *for all* $i \in \{1, \ldots, n\}$, the action profile is a *consensus* choice, meaning that group and individual welfare is maximised when **a** is instantiated.
- A consensus choice will often not exist, in which case players might be motivated to enter into negotiation to reach compromise.

The Prisoners' Dilemma (PD)

- For present purposes it suffices to show that conditional game theory generalises classical game theory in cases where consensus choice applies.
- To show this we begin with the Prisoners' Dilemma (PD).
- The game is dominance-solvable because both players have dominant strategies to Defect but the outcome that arises when both players use their dominant strategies (D, D) is worse for both players than the outcome that would arise (C, C) if they used their dominated strategies instead.
- As established by Binmore (1994), if the preference structure of the PD describes *all* of the relevant preference information pertinent to the interaction, then general defection is the only outcome that a game theoretic model of it can predict.

The Prisoners' Dilemma (PD)

- Binmore further insists that if the model does *not* incorporate all such information in the specification of preferences, then the game should not be characterised as a PD in the first place.
- However, to admit the *possibility* of social influences is to admit that more than one game structure might be relevant to modelling an interaction.
- This situation is hardly unprecedented in economics: we are used to the idea, for example, that a plurality of models are useful for foregrounding different aspects of international trade, oligopoly, national production, and other phenomena.
- Put another way: there may be more information that is relevant to the strategic interaction than the (full) information about categorical preferences.

Cooperation and Exploitation Indices

- Stirling (2012, p. 80) draws a distinction between simple reciprocal altruism that transforms PDs into coordination games, and background representations of interaction structures in which players' models of their own and others' preferences are consistent with the PD structure but they are also aware of preferences they *would* have *conditional* on the implementation of some degree of socially mediated agency.
- This is indeed the basis on which Stirling's general framework is given its name.
- The machinery by which Stirling represents genuine PD structure simultaneously with scope for team agency representation are *cooperation* and *exploitation indices*.
- Each player X_i is endowed with a cooperation index $\alpha_i \in [0, 1]$ and an exploitation index $\beta_i \in [0, 1]$, where α_i represents the extent to which a player is conditionally willing to cooperate, and β_i represents the extent to which a player is conditionally willing to exploit his or her partner.

PD Categorical Preferences

- As a minimal consistency requirement, assume that $\alpha + \beta < 1$.
- To respect acyclicity, assume further that X_1 has categorical preferences and that X_2's preferences are conditional on X_1's.
- Given the cooperation and exploitation indices, X_1's categorical utility is defined as follows:

$$u_{X_1}(C, C) = \alpha_1 \quad u_{X_1}(C, D) = 0$$

$$u_{X_1}(D, C) = \beta_1 \quad u_{X_1}(D, D) = 1 - \alpha_1 - \beta_1$$

- In the PD representation of the interaction,

  $\beta_1 > \alpha_1 > 1 - \alpha_1 - \beta_1 > 0$, and X_2 has a categorical utility function such that $u_{X_2}(C, D) > u_{X_2}(C, C) > u_{X_2}(D, D) > u_{X_2}(D, C)$.

PD Conditional Preferences

- For the conditional representation, we calculate $u_{X_2|X_1}(a_{21}, a_{22} \mid a_{11}, a_{12})$ by computing utilities for every possible conjecture that player X_1 can make.
- Assume that if X_1 conjectures either (C, C) or (D, D) then X_2 will place all of her conditional utility mass on the same action profile.
- In other words, if X_1 conjectures cooperation then X_2 finds it optimal to cooperate but if X_1 conjectures defection then X_2 finds it optimal to defect.
- If X_1 conjectures (C, D), then X_2's utility mass will be apportioned according to her cooperation and exploitation indices.
- Specifically, X_2 will assign α_2 to (C, C), β_2 to (C, D), $1 - \alpha_2 - \beta_2$ to (D, D) and zero utility mass to (D, C) because this is the worst possible outcome for X_2.
- Finally, if X_1 conjectures (D, C), the worst outcome for X_2, X_2 should place zero utility mass on (D, C), α_2 and (C, C), β_2 on (C, D) and $1 - \alpha_2 - \beta_2$ on (D, D).

PD Conditional Utilities

- The conditional utilities associated with each conjecture of X_1, represented in the columns, and every action profile which can be instantiated by the two players, represented in the rows, are given in the table below:

	(a_{11}, a_{12})			
(a_{21}, a_{22})	(C, C)	(C, D)	(D, C)	(D, D)
(C, C)	1	α_2	α_2	0
(C, D)	0	β_2	β_2	0
(D, C)	0	0	0	0
(D, D)	0	$1 - \alpha_2 - \beta_2$	$1 - \alpha_2 - \beta_2$	1

- Having defined the categorical utility of X_1 and the conditional utility of X_2, we can now calculate the concordant utility.

PD Concordant Utility

- To compute the concordant utility we combine X_1's categorical utility with X_2's conditional utility:

  $U_{X_1X_2}(\mathbf{a}_1, \mathbf{a}_2) = u_{X_2|X_1}(a_{21}, a_{22} \mid a_{11}, a_{12}) \times u_{X_1}(a_{11}, a_{12})$
- The result is shown in the table below where the rows index X_1's conjecture and the columns index X_2's conjecture:

	(a_{21}, a_{22})			
(a_{11}, a_{12})	(C, C)	(C, D)	(D, C)	(D, D)
(C, C)	α_1	0	0	0
(C, D)	0	0	0	0
(D, C)	$\alpha_2\beta_1$	$\beta_1\beta_2$	0	$\beta_1 - \alpha_2\beta_1 - \beta_1\beta_2$
(D, D)	0	0	0	$1 - \alpha_1 - \beta_1$

- The concordant utility can now be used to extract the *ex post* marginal utilities, the group welfare function, and the individual welfare function.

Marginal Utilities

	(a_{21}, a_{22})			
(a_{11}, a_{12})	(C, C)	(C, D)	(D, C)	(D, D)
(C, C)	α_1	0	0	0
(C, D)	0	0	0	0
(D, C)	$\alpha_2\beta_1$	$\beta_1\beta_2$	0	$\beta_1 - \alpha_2\beta_1 - \beta_1\beta_2$
(D, D)	0	0	0	$1 - \alpha_1 - \beta_1$

- X_1's *ex post* utilities are equivalent to her categorical utilities, whereas X_2's *ex post* utilities must be derived through marginalisation: $u_{X_2}(\mathbf{a}_2) = \Sigma_{\sim \mathbf{a}_2} U_{X_1 X_2}(\mathbf{a}_1, \mathbf{a}_2)$
- For example, $u_{X_2}(\text{C, C}) = \alpha_1 + 0 + \alpha_2\beta_1 + 0 = \alpha_1 + \alpha_2\beta_1$
- The *ex post* payoff matrix for the PD is show below:

		X_2	
		C	D
X_1	C	$\alpha_1,\ \alpha_1 + \alpha_2\beta_1$	$0,\ \beta_1\beta_2$
	D	$\beta_1,\ 0$	$1 - \alpha_1 - \beta_1,\ 1 - \alpha_1 - \alpha_2\beta_1$

Group Welfare Function

	(a_{21}, a_{22})			
(a_{11}, a_{12})	(C, C)	(C, D)	(D, C)	(D, D)
(C, C)	α_1	0	0	0
(C, D)	0	0	0	0
(D, C)	$\alpha_2\beta_1$	$\beta_1\beta_2$	0	$\beta_1 - \alpha_2\beta_1 - \beta_1\beta_2$
(D, D)	0	0	0	$1 - \alpha_1 - \beta_1$

- The group welfare function for this two-player game is derived using the following expression:

  $V_{X_1X_2}(a_{11}, a_{22}) = \Sigma_{\sim a_{11}} \Sigma_{\sim a_{22}} U_{X_1X_2}(\mathbf{a}_1, \mathbf{a}_2)$
- For example, $V_{X_1X_2}$ (D, D) =

  $\beta_1\beta_2 + \beta_1 - \alpha_2\beta_1 - \beta_1\beta_2 + 0 + 1 - \alpha_1 - \beta_1 = 1 - \alpha_1 - \alpha_2\beta_1$
- The full group welfare function is:

$$V_{X_1X_2}(C, C) = \alpha_1$$
$$V_{X_1X_2}(C, D) = 0$$
$$V_{X_1X_2}(D, C) = \alpha_2\beta_1$$
$$V_{X_1X_2}(D, D) = 1 - \alpha_1 - \alpha_2\beta_1$$

Individual Welfare Functions

$$V_{X_1X_2}(C, C) = \alpha_1$$
$$V_{X_1X_2}(C, D) = 0$$
$$V_{X_1X_2}(D, C) = \alpha_2\beta_1$$
$$V_{X_1X_2}(D, D) = 1 - \alpha_1 - \alpha_2\beta_1$$

- Finally, the individual welfare functions are extracted from the group welfare function using marginalisation:
  $v_{X_i}(a_i) = \Sigma_{\sim a_i} V_{X_1X_2}(a_1, a_2)$
- For example, $v_{X_2}(C) = \alpha_1 + \alpha_2\beta_1$
- Thus, the individual welfare functions are:

$$v_{X_1}(C) = \alpha_1 \qquad v_{X_1}(D) = 1 - \alpha_1$$
$$v_{X_2}(C) = \alpha_1 + \alpha_2\beta_1 \qquad v_{X_2}(D) = 1 - \alpha_1 - \alpha_2\beta_1$$

NE with Social Influences

- To find the NE of this game after incorporating the social influence flows between X_1 and X_2, we work directly with the conditional and categorical utilities (yielding *conditioned NE*) or the *ex post* marginal utilities (yielding *ex post NE*); the two approaches yield identical solutions.
- The table below shows that (D, D) is a NE for all admissible values of α_i and β_i.
- Unlike the unconditional PD, (C, C) is a NE when $\alpha_i > \beta_i$.
- Furthermore, when $\alpha_i > \beta_i$ and $\alpha_1 > 0.5$, (C, C) is a consensus choice because it maximises both group and individual welfare.

		X_2	
		C	D
X_1	C	$\alpha_1, \alpha_1 + \alpha_2\beta_1$	$0, \beta_1\beta_2$
	D	$\beta_1, 0$	$1 - \alpha_1 - \beta_1, 1 - \alpha_1 - \alpha_2\beta_1$

The Role that Social Influences Play

- It is intuitive that if both players prefer cooperation to exploitation then (C, C) will be a conditioned or *ex post* NE but this result fails to highlight the role that social influences can play in this game.
- To see this, assume that $\alpha_1 = 0.6$ and $\beta_1 = 0.3$ and that $\alpha_2 = 0.3$ and $\beta_2 = 0.6$.
- Thus, X_1's cooperation index is twice as large as her exploitation index but X_2's cooperation index is half as large as her exploitation index.
- So, in the absence of influence flows, X_1 is a cooperator and X_2 is an exploiter.
- But, after X_2 takes into account X_1's preferences, X_2's penchant for exploitation is tempered by X_1's desire for cooperation and (C, C) is a conditioned NE.

		X_2	
		C	D
X_1	C	0.6, 0.69	0, 0.18
	D	0.3, 0	0.1, 0.13

Hi-Lo

- Conditional game theory's explanation of cooperation in a one-shot PD is an important accomplishment, but Michael Bacharach (2006) argued that the litmus test for an effort to represent group or team agency is that it furnish an explanation for High play in the game Hi-Lo.
- From the game table below it is clear that Hi-Lo has two pure strategy NE, (High, High) and (Low, Low), where the former Pareto dominates the latter.
- Hi-Lo raises the same kind of equilibrium selection problem as a pure coordination game in that it does not prescribe play of one pair of equilibrium strategies over the other.
- But the indeterminacy in Hi-Lo seems particularly troubling because in actual applications, people have no problem at all in coordinating on the (High, High) equilibrium.

		X_2	
		High	Low
X_1	High	2, 2	0, 0
	Low	0, 0	1, 1

Conditional Hi-Lo

- Bacharach (2006) proposes a solution to Hi-Lo which relies on team reasoning; this will be discussed by other workshop participants so we will focus on the conditional approach to the game.
- To allow social influences to affect the analysis of Hi-Lo, we endow each player X_i with a *High play index* $\alpha_i \in [0, 1]$ and a *Low play index* $\beta_i \in [0, 1]$, where $\alpha_i + \beta_i = 1$, because the players will assign zero utility mass to mis-matches (i.e., (H, L) and (L, H)).
- Assume again that X_1's preferences are categorical and that X_2's preferences are conditional on X_1's.
- Given the High play and Low play indices, X_1's categorical utility is defined as follows:

$$u_{X_1}(H, H) = \alpha_1 \quad u_{X_1}(H, L) = 0$$

$$u_{X_1}(L, H) = 0 \quad u_{X_1}(L, L) = \beta_1$$

Hi-Lo Conditional Preferences

- To calculate $u_{X_2|X_1}$ (a_{21}, a_{22} | a_{11}, a_{12}) it is necessary to compute utilities for every possible conjecture that player X_1 can make.
- Assume that if X_1 conjectures either (H, H) or (L, L) then X_2 will place all of her conditional utility mass on the same action profile.
- That is, if X_1 conjectures High then X_2 finds it optimal to play High but if X_1 conjectures Low then X_2 finds it optimal to play Low.
- If X_1 conjectures (H, L) or (L, H), then X_2's utility mass will be apportioned according to her High play and Low play indices.
- Specifically, X_2 will assign α_2 to (H, H) and β_2 to (L, L), and zero utility mass to (H, L) and (L, H) because these are the worst outcomes for X_2.

Hi-Lo Conditional Utilities

- The conditional utilities associated with each conjecture of X_1, represented in the columns, and every action profile which can be instantiated by the two players, represented in the rows, are given in the table below:

	(a_{11}, a_{12})			
(a_{21}, a_{22})	(H, H)	(H, L)	(L, H)	(L, L)
(H, H)	1	α_2	α_2	0
(H, L)	0	0	0	0
(L, H)	0	0	0	0
(L, L)	0	β_2	β_2	1

- Having defined the categorical utility of X_1 and the conditional utility of X_2, we can now calculate the concordant utility.

Hi-Lo Concordant Utility

- To compute the concordant utility we combine X_1's categorical utility with X_2's conditional utility:

  $U_{X_1X_2}(\mathbf{a}_1, \mathbf{a}_2) = u_{X_2|X_1}(a_{21}, a_{22} \mid a_{11}, a_{12}) \times u_{X_1}(a_{11}, a_{12})$
- The result is shown in the table below where the rows index X_1's conjecture and the columns index X_2's conjecture:

	(a_{21}, a_{22})			
(a_{11}, a_{12})	(H, H)	(H, L)	(L, H)	(L, L)
(H, H)	α_1	0	0	0
(H, L)	0	0	0	0
(L, H)	0	0	0	0
(L, L)	0	0	0	β_1

- As we did earlier, the concordant utility can now be used to extract the *ex post* marginal utilities, the group welfare function, and the individual welfare function.

Ex Post Payoff Matrix and Welfare Functions

- The *ex post* payoff matrix for Hi-Lo is given below:

		X_2	
		High	Low
X_1	High	α_1, α_1	0, 0
	Low	0, 0	β_1, β_1

- The group welfare function is:

$$V_{X_1X_2}(\mathrm{H}, \mathrm{H}) = \alpha_1$$
$$V_{X_1X_2}(\mathrm{H}, \mathrm{L}) = 0$$
$$V_{X_1X_2}(\mathrm{L}, \mathrm{H}) = 0$$
$$V_{X_1X_2}(\mathrm{L}, \mathrm{L}) = \beta_1$$

- The individual welfare functions are:

$$v_{X_1}(\mathrm{H}) = \alpha_1 \qquad v_{X_1}(\mathrm{L}) = \beta_1$$
$$v_{X_2}(\mathrm{H}) = \alpha_1 \qquad v_{X_2}(\mathrm{L}) = \beta_1$$

- Clearly (H, H) and (L,L) are *ex post* NE.

Hi-Lo Consensus

Group Welfare Function:

$V_{X_1X_2}$ (H, H) = α_1

$V_{X_1X_2}$ (H, L) = 0

$V_{X_1X_2}$ (L, H) = 0

$V_{X_1X_2}$ (L, L) = β_1

Individual Welfare Functions

v_{X_1} (H) = α_1 v_{X_1} (L) = β_1

v_{X_2} (H) = α_1 v_{X_2} (L) = β_1

- Looking at the group and individual welfare functions we see that group and individual welfare is maximised through the profile (H, H) when $\alpha_1 > \beta_1$.
- As this is the assumption in Hi-Lo, the profile that caters for the interests of the individuals and the group is (H, H) and this is a consensus choice.
- Consequently, we would expect this profile to be instantiated when players take into account their own individual interests and the interests of the group, as encoded in the social linkages among the players and expressed through the group welfare function.

Conclusion

- Conditional game theory is a powerful framework that incorporates notions of group or team agency in classical game theory.
- It allows for influence flows which propagate through a group and define agents' rankings of alternative outcomes as conditional on the preferences of others.
- This propagation of influence flows defines a social model that enables agents to jointly consider individual and group interests, but without requiring us to leave the Nash constraint.
- Thus, conditional game theory has full power to represent team agency using only resources that can be defined within standard game-theoretic formalism.