Coherent Cooperation

Robert Nau
Fuqua School of Business
Duke University

Presentation for the Philosophy of Science Association Meeting
Vancouver, November 3, 2006

The problem: how to achieve a "cooperative" solution in a game such as the prisoner's dilemma?

	Left	*Right*
Top	4, 4	0, 5
Bottom	5, 0	1, 1

What dilemma?

- **Fact 1:** Each player is 1 utile better off by playing the *second* strategy rather than the first, regardless of the other's move.
- **Fact 2:** Each player is 4 utiles better off if her *opponent* plays the *first* strategy rather than the second, regardless of her own move.
- Noncooperative solution concepts (correlated equilibrium, Nash equilibrium…) only "hear" Fact 1.
- They perceive no dilemma in playing BR, because they *don't know* it yields (1,1) instead of (4,4).

"What we have here, is a failure to communicate." *--Strother Martin in "Cool Hand Luke"*

- The inefficient noncooperative solution is not so much due to a failure to cooperate, but rather due to a failure to *communicate*.
- In the usual noncooperative framework, there is no way for the players to credibly communicate their preferences for *others'* moves via *unilateral* commitments such as offers to accept gambles.
- Hence, full information about payoff functions is not really common knowledge.

A non-cooperative approach to the modeling of cooperation

- Following Myerson (1991), the usual noncooperative approach will be extended to allow a limited form of *bilateral* commitment.
- Players' strategy sets are enlarged to include conditional agreements to accept a cooperative solution, provided that other players also agree.
- The focus here is on the epistemic implications of this approach for the *communication* of preferences and for the modeling of strategic rationality in terms of *coherence*.

Key modeling assumptions

- De Finetti's concept of coherence is extended to a game-theoretic framework (Nau and McCardle 1990).
- For convenience, game outcomes are assumed to be monetary, and utility for money is assumed to be linear and state-independent.
- These assumptions could be relaxed to allow non-monetary outcomes and nonlinear and/or state-dependent utility for money.
- However, it is necessary for money to be *one* attribute of utility, with *local* linearity (i.e., linear, possibly state-dependent utility for "small" gambles).

Review: communication of unilateral preferences via acceptance of gambles

- True payoff matrix ("***F*** ") with entries $f_1(s_1, s_2)$, $f_2(s_1, s_2)$

	L	R
T	a_1, a_2	b_1, b_2
B	c_1, c_2	d_1, d_2

- Gamble matrix (" ***G*** ") whose rows are unilaterally acceptable gambles that reveal individual preferences

Interpretation of gamble *1TB*: in the event that player 1 plays *T* when *B* is available, she would accept a gamble proportional to payoff differences between *T* and *B*

	TL	TR	BL	BR
$1TB$	$a_1 - c_1$	$b_1 - d_1$		
$1BT$			$c_1 - a_1$	$d_1 - b_1$
$2LR$	$a_2 - b_2$		$c_2 - d_2$	
$2RL$		$b_2 - a_2$		$d_2 - c_2$

The "real" rules of the game

- For the purposes of noncooperative game theory, the real rules of the game consist of the revealed gamble matrix $\boldsymbol{G}$, not the true payoff matrix $\boldsymbol{F}$.
- $\boldsymbol{G}$ suffices to determine the set of noncooperative equilibria (correlated, Nash, etc.): a correlated equilibrium distribution π is defined by the constraints $\boldsymbol{G}\pi \geq \mathbf{0}$, and a Nash equilibrium additionally satisfies between-player independence.
- However, $\boldsymbol{G}$ contains *no* information about players' preferences for *each others'* moves: any additive effect on player 1 of player 2's choice is subtracted out of the terms $a_1 - c_1$, $b_1 - d_1$, etc.

The real rules of the PD game

- True payoff matrix:

	L	*R*
T	4, 4	0, 5
B	5, 0	1, 1

- Gamble matrix: players are collectively willing to throw money away if *BR* is not played.

A sum of gambles ***1TB*** and ***2LR*** yields an aggregate loss for the players in outcomes ***TL***, ***TR***, and ***BL***, with zero gain in ***BR***.

	TL	*TR*	*BL*	*BR*
1TB	- 1	- 1		
1BT			1	1
2LR	-1		- 1	
2RL		1		1

Strategic rationality as "joint coherence"

- Let $\boldsymbol{x}$ denote a vector of non-negative multipliers for acceptable gambles (rows of $\boldsymbol{G}$) chosen by an outside observer.
- $\boldsymbol{x}^T\boldsymbol{G}$ is the vector of aggregate payoffs from the observer to the players, whose jk^{th} element $[\boldsymbol{x}^T\boldsymbol{G}]_{jk}$ is the payoff when Row plays j and Column plays k
- Outcome jk is *jointly coherent* if it does not present an opportunity for *ex post arbitrage*, i.e., if there does not exist $\boldsymbol{x} \geq \boldsymbol{0}$ s.t. $\boldsymbol{x}^T\boldsymbol{G} \leq \boldsymbol{0}$, $[\boldsymbol{x}^T\boldsymbol{G}]_{jk} < 0$, i.e., if observer cannot lose and wins a finite amount if jk is played.
- Only *BR* is jointly coherent in the PD game.

Fundamental duality theorem of noncooperative games

- Joint coherence is the appropriate extension of the concept of Bayesian rationality (â la de Finetti) to a noncooperative game of strategy.
- **Theorem** (Nau and McCardle 1990): An outcome is jointly coherent if and only if it has positive probability in some correlated equilibrium of the game.
- Hence, correlated – not Nash – equilibrium is the expression of Bayesian rationality in noncooperative games, as originally claimed by Aumann (1974, 1987).

How to extend this duality result to provide a way out of the prisoner's dilemma?

- Evidently it is necessary for the players to accept additional gambles that express their preferences for their *opponents'* actions.
- This requires a limited form of *bilateral* commitment, enforced by a mediator or contracting mechanism.

The model: an "extended" game

- After some cheap talk, suppose a cooperative solution is proposed, with payoff vector $\boldsymbol{e} = (e_1, e_2)$, e.g., a pure or objectively randomized outcome of original game.
- Let $\boldsymbol{s} = (s_1, s_2)$ denote a vector of strategies in the original game, and let $\boldsymbol{c} = (c_1, c_2)$ denote a vector of cooperate/defect indicators (1=cooperate, 0=defect)
- Player i's strategy now consists of the pair (s_i, c_i), which *doubles* her set of available strategies.
- If all players choose to cooperate ($\boldsymbol{c} = \mathbf{1}$), they receive the cooperative payoff, otherwise they receive the payoffs of strategy $\boldsymbol{s}$ in the original game:
  $f_i(\boldsymbol{s}, \boldsymbol{c}) = e_i$ if $\boldsymbol{c} = \mathbf{1}$, otherwise $f_i(\boldsymbol{s}, \boldsymbol{c}) = f_i(\boldsymbol{s})$

Degree of bilateral commitment required

- All players must allow the cooperative outcome to be implemented on their behalf, yielding the joint payoff vector $\boldsymbol{e}$, if and only if they *all* cooperate ($\boldsymbol{c} = \mathbf{1}$).
- Otherwise they end up at the "disagreement point" $\boldsymbol{s}$ in the original game.
- The game is still simultaneous-move, as if player i submits the move (s_i, c_i) in a sealed envelope before seeing the others' moves.
- The disagreement point may be to some extent uncertain: $\boldsymbol{s}$ may have been obtained by randomization, either independent or correlated.

The extended game (2x2 case)

- Let *TC, TD* denote (*Top, Cooperate*), (*Top, Defect*), etc.
- The extended game is a 4x4 noncooperative game with the following true payoff functions:

	LC	*RC*	*LD*	*RD*
TC	e_1, e_2	e_1, e_2	a_1, a_2	b_1, b_2
BC	e_1, e_2	e_1, e_2	c_1, c_2	d_1, d_2
TD	a_1, a_2	b_1, b_2	a_1, a_2	b_1, b_2
BD	c_1, c_2	d_1, d_2	c_1, c_2	d_1, d_2

Implications of extension of the game

- Noncooperative equilibria of the original game are still equilibria of the extended game (lower right quadrant).
- However, noncooperative equilibria of the extended game may also include the *cooperative* outcomes (upper left quadrant), depending on $\boldsymbol{e}$.
- Key difference: in the extended game, acceptable gambles involving comparisons between cooperation and defection reveal additional information about preferences for others' strategies.

Gamble matrix for extended 2x2 game (player 1's rows only)

	TC LC	*TC RC*	*TC LD*	*TC RD*	*BC LC*	*BC RC*	*BC LD*	*BC RD*	*TD LC*	*TD RC*	*TD LD*	*TD RD*	*BD LC*	*BD RC*	*BD LD*	*BD RD*
1*TCBC*	0	0	*a-c*	*b-d*												
1*TCTD*	*e-a*	*e-b*	0	0												
1*TCBD*	*e-c*	*e-d*	*a-c*	*b-d*												
1*BCTC*					0	0	*c-a*	*d-b*								
1*BCTD*					*e-a*	*e-b*	*c-a*	*d-b*								
1*BCBD*					*e-c*	*e-d*	0	0								
1*TDTC*									*a-e*	*b-e*	0	0				
1*TDBC*									*a-e*	*b-e*	*a-c*	*b-d*				
1*TDBD*									*a-c*	*b-d*	*a-c*	*b-d*				
1*BDTC*													*c-e*	*d-e*	*c-a*	*d-b*
1*BDBC*													*c-e*	*d-e*	0	0
1*BDTD*													*c-a*	*d-b*	*c-a*	*d-b*

Green cells involve new comparisons of e vs. a,b,c,d
Purple cells correspond to comparisons in original game.

Prisoner's dilemma revisited

- Suppose that the proposed cooperative payoff vector is the efficient pure outcome $\boldsymbol{e} = (4, 4)$.
- Then (*BC, RC*) is a jointly coherent strategy: players cooperate while threatening to play *BR*.

	LC	*RC*	*LD*	*RD*
TC	4, 4	4, 4	4, 4	0, 5
BC	4, 4	4, 4	5, 0	1, 1
TD	4, 4	0, 5	4, 4	0, 5
BD	5, 0	1, 1	5, 0	1, 1

General results

- **Theorem 1:** Every jointly coherent outcome in the original game remains a jointly coherent outcome of the extended game.
- **Proof:** If it is recommended that every player defect to some correlated equilibrium strategy, then no *single* player is better off *either* by choosing to cooperate or by choosing some non-recommended strategy in the original game.

General results, continued

- **Theorem 2:** The proposed cooperative solution in the extended game is jointly coherent *if* it weakly dominates the conditional expected payoffs of some correlated equilibrium in the original game.
- **Proof:** Let the weakly dominated CE be the "disagreement point" that enforces cooperation.
- "Only if" does not hold, as shown by a later example (prisoner's non-dilemma).

PD game without cooperation

	L	*R*
T	4, 4	0, 5
B	5, 0	1, 1

By strict dominance, the unique jointly coherent strategy is the Nash equilibrium, whose payoff vector is the inefficient outcome (1, 1).

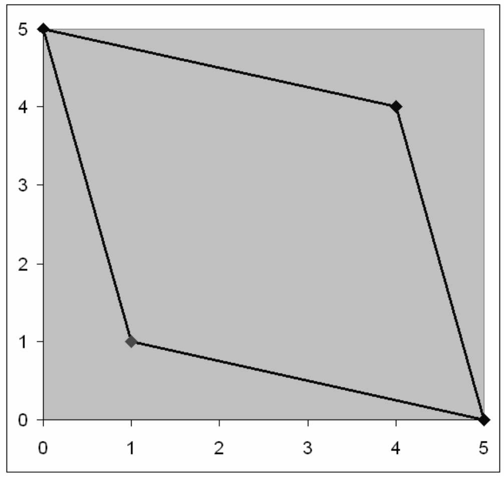

PD game with cooperation

	L	*R*
T	4, 4	0, 5
B	5, 0	1, 1

Jointly coherent cooperative outcomes lie in the shaded region that weakly dominates the NE, including $\boldsymbol{e} = (4, 4)$, and the gamble matrix reveals that *TL* is better than *BR* for both players

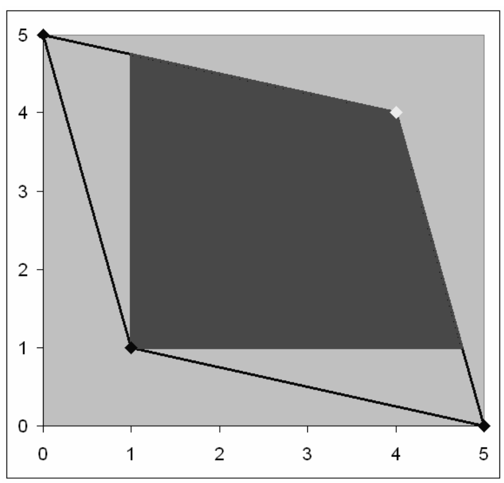

Modified (asymmetric) PD game

	L	*R*
T	4, 2	0, 5
B	5, 0	1, 3

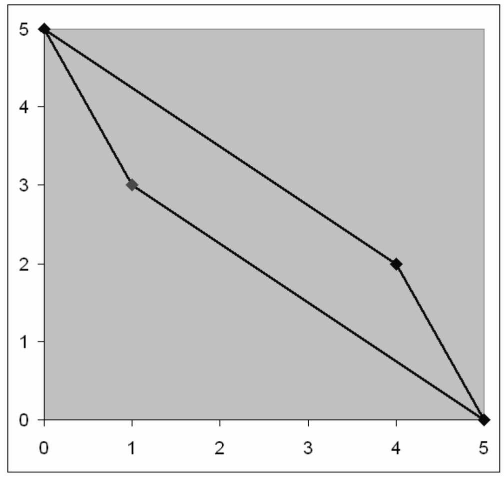

Unique jointly coherent strategy is the Nash equilibrium with payoff vector (1, 3), which is not dominated by any other *pure* strategy.

Modified PD game with cooperation

	L	*R*
T	4, 2	0, 5
B	5, 0	1, 3

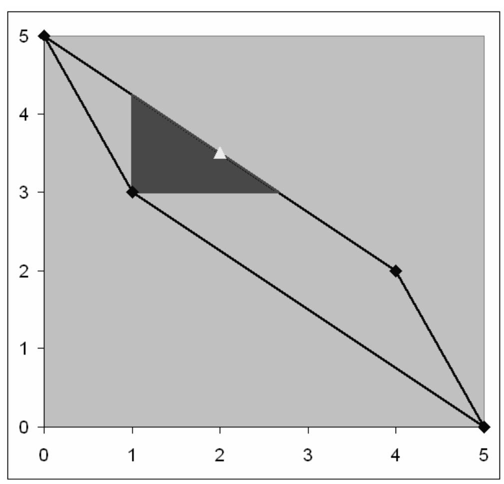

Jointly coherent outcomes lie in shaded region, and the gamble matrix reveals that some *mixtures* of *TL* and *TR* (but no pure strategies) are preferred to *BR*.

Prisoner's non-dilemma

	L	*R*
T	0, 0	1, 1
B	1, 1	2, 2

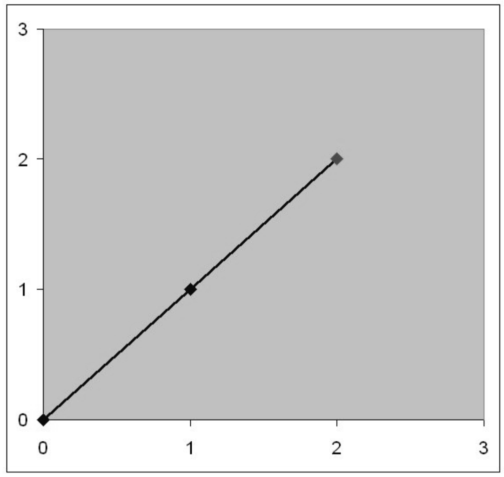

This game is also dominance solvable and strategically equivalent to the original prisoner's dilemma, but the Nash equilibrium is now Pareto efficient, with payoff vector $(2, 2)$.

Prisoner's non-dilemma with cooperation

	L	*R*
T	0, 0	1, 1
B	1, 1	2, 2

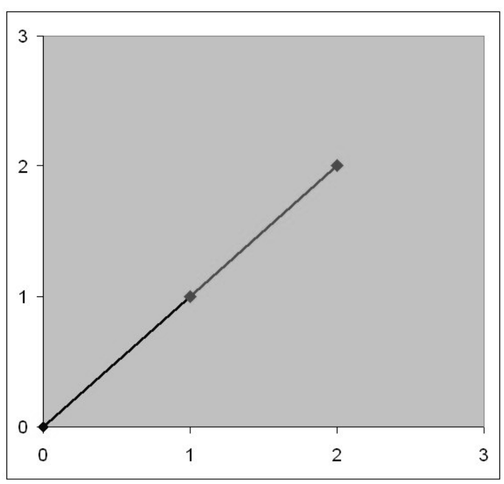

In the extended game, all cooperative outcomes between $(1, 1)$ and $(2, 2)$ are jointly coherent, but more importantly, the gamble matrix reveals that BR is better than TL for both players.

Conclusions

- Noncooperative modeling of cooperation, via conditional contracts, enriches the possibilities for communication as well as rational behavior.
- Joint coherence (no ex post arbitrage) is still an appropriate standard of rationality in this setting.
- Any pure or randomized joint strategy that weakly dominates the conditional payoffs of a correlated equilibrium can be rationalized in this way.