

ENGINEERING COOPERATION IN TWO-PLAYER STRATEGIC GAMES

DRAFT MARCH 2010

ADAM TAUMAN KALAI* AND EHUD KALAI^{†,§}

ABSTRACT. Selfish, strategic players may benefit from cooperation, *provided they reach an agreement*. It is therefore important to construct mechanisms that enable such cooperation, especially in the case of asymmetric private information. There are two major issues: (1) the determination of a fair and efficient outcome among the many compromises possible in a strategic game, and (2) the establishment of a play protocol under which strategic players will agree to the selected compromise. The paper presents a general solution for an important class of two person Bayesian games with monetary payoffs.

The proposed solution builds on earlier concepts in game theory. It coincides with the von Neumann minmax value on the class of zero sum games and with the major solution concepts to the Nash Bargaining Problem. Moreover the solution is based on a simple decomposition of every game into cooperative and competitive components, which is easy to compute.

1. INTRODUCTION

Selfish players in strategic games benefit from cooperation, provided they come to mutually beneficial agreements. In the case of asymmetric private information, the benefits may be even greater, but avoiding strategic manipulations is more subtle. This paper provides a natural focal point for a fair and efficient cooperative play amongst strategic players, in two-person private-information games with monetary payoffs.

The unmodified non cooperative solutions obviously miss this point, as illustrated in the sacrifice game below. The dominant strategy, in which player 1 does not act, yields each player \$0 payoff. This makes sense in many environments, for example when the game is played by prisoners in isolated cells.

sacrifice game	not act
not act	\$0, 0
act	-1,101

But in environments where side payments and agreements are permitted the outcome that player 1 acts in exchange for a \$51 payoff from player 2, so that they each net \$50, is more natural and agrees with real life and experimental observations.

However, when we consider more substantial games, in which players possess many possible strategies and asymmetric private information, the determination and implementation of optimal

Key words and phrases. cooperative game theory, non-cooperative game theory, bargaining, minmax value.

*Microsoft Research New England.

[†]Kellogg School of Management, Northwestern University.

[§]Much of this work was done while the author was visiting Microsoft Research New England.

The authors thank Geoffrey De Clippel, Françoise Forges, Dov Monderer, Phil Reny, Rakesh Vohra, Robert Wilson and other seminar participants at Northwestern, Stanford, The conference of Public Economic Theory in Gallway, and the International Conference on Game Theory in Stony Brook for helpful comments.

This research is partly supported by National Science Foundation Grant No. SES-0527656 in Economics/Computer Science,

cooperative play and associated payoff transfers is less obvious. The main purpose of this paper is to offer a solution to this problem in restricted but important classes of games.

In particular, we restrict ourselves to two-player cooperative Bayesian games with transferable utility (TU),¹ in which players can communicate and sign binding agreements about actions and payoff transfers.² An important subclass are the zero-sum games studied by von Neumann.

The solution offered for the above class of games is described by a *cooperative-competitive value* (alternatively *coco value*, or just *value* for short) that has the following properties: (1) It is Pareto efficient, fair, robust and easy to compute. (2) It generalizes the minmax value from zero-sum to general-sum games. (3) It coincides with, and extends the major variable-threat bargaining solutions to the case of incomplete information. (4) It is justified by natural axioms imposed directly on Bayesian games. And (5) it is implementable by incentive-compatible protocols that resemble real-life partnerships.

The analysis is centered around a cooperative/competitive (coco) decomposition of a strategic game into two component games with orthogonal incentives. For a complete information game with payoff matrices (X, Y) , the decomposition is

$$(X, Y) = \left(\frac{X+Y}{2}, \frac{X+Y}{2}\right) + \left(\frac{X-Y}{2}, \frac{Y-X}{2}\right),$$

and the

$$\text{coco-value } (X, Y) \equiv \left(\max_{i,j} \frac{X+Y}{2}, \max_{i,j} \frac{X+Y}{2}\right) + \min \max \left(\frac{X-Y}{2}, \frac{Y-X}{2}\right).^3$$

The first term is the highest possible pair of payoffs that the players can jointly arrange under an agreement to share their payoffs equally. The second term is a compensating zero-sum transfer from the player with the weaker strategic position to the stronger one. The body of this paper presents the Bayesian (incomplete information) version of the decomposition and definition above.

While the implementation of the coco value is trivial for complete information games, the implementation of any first best efficient payoffs in a Bayesian game is difficult, as illustrated for example by the Myerson-Satterthwaite impossibility result. However, the coco Bayesian decomposition enables the construction of *partnership protocols* that overcome these difficulties in important classes of games.

In rough terms, one can replace a given Bayesian game by two strategically-independent games, one cooperative and one competitive. The play of the cooperative component determines the actual play of the given game, and it leads to an actual pair of equal payoffs. The equal sharing of payoffs gives the two players (partners) the incentives to act cooperatively: truthfully reveal all relevant information and act optimally relative to this information.

The competitive component is played fictitiously. Its purpose is to determine a compensating zero-sum payoff transfer that reflects the strategic and information asymmetries in the given game. Since the play of the competitive component results in a zero-sum transfer, it does not destroy the efficiency obtained through the actual play of the cooperative component, it just corrects for the imposed equal division in the cooperative component.

We are able to overcome the Myerson-Satterthwaite impossibility by restricting ourselves to games in which players actual abilities and/or payoffs can be assessed.

¹A simple way to think of this assumption is that payoffs are monetary or monetarily measurable, and that each player's utility is $u_i(\$x) = x$.

²A more restricted earlier use of payoff transfers in strategic games with complete information is presented in Jackson and Wilkie (2005).

³The decomposition has a straightforward extension to n -person games with $n > 2$. But the definition of the coco value is more complex there, since there is no simple notion of a minmax value for n -person zero-sum games.

1.1. Related literature. Given the fundamental nature of the questions studied here, it is not surprising that closely related concepts have been studied in both the cooperative and strategic theories, and in the smaller earlier semi-cooperative literature. Indeed, the solution presented here may be viewed as a synthesis and generalization of several earlier works.

Starting with two-person zero-sum games, where there is no possibility to cooperate, the coco value generalizes von Neumann's (1928) minmax value to general-sum games. First, in a formal sense, the coco value of a zero-sum game is the minmax value of the game. But in addition, our axiomatic characterization of the coco value (on the class of general-sum games) can be fully carried out on the restricted class of zero-sum games to yield an axiomatic characterization of the minmax value. Thus, the coco value may be viewed as an expansion of the "rational" of the minmax value to the class of general-sum games.

Moving to non zero-sum games, both the strategic and cooperative game theory deal with the issue of cooperation in different ways, but neither is fully satisfactory. The Nash equilibria of a one-shot strategic game in which all binding agreements are possible span all the outcomes described by the folk theorem of the repeated game,⁴ thus "everything is possible." The cooperative Shapley value of a game is unique, but in order to apply it one needs to find the proper way of a "bridge," transforming the strategic game into a cooperative bargaining game. In addition, it seems that both, the folk theorem and the Shapley value, miss some subtle strategic considerations as discussed in the concluding section of this paper.

The semi cooperative approach, taken this paper, defines a unique cooperative solution directly for any strategic game. Motivated by axiomatic bargaining solutions from cooperative game theory, (Nash, 1950b; Kalai-Smorodinsky, 1975; and Kalai, 1977) Nash (1953), Raiffa (1953) and Kalai-Rosenthal (1978), defined (without axiomatizations) efficient arbitration methods for general (non TU) two-person normal-form games with complete information. In the case of TU games all their solutions coincide with each other and with the coco value presented here. Thus, the coco value may be viewed as a TU generalization of their solutions to the case of private information.⁵

Selten (1960, 1963) presented an axiomatic characterization of a cooperative value, defined on the class of TU extensive form games of complete information. While his work predate the definition of incomplete information, the natural extension of Selten's work would give a different value for games of incomplete information.⁶ Different axiomatic solution of complete information strategic games, discussed later in this paper, were presented by Carpenete et al (2005,2006).

Little work was done on semi-cooperative solutions for games with incomplete information. Staying in the purely cooperative model, Myerson (1984) offers an extension of the Nash (1950b, 1953) cooperative bargaining solution to the case of private information, but this solution has not been studied directly for strategic games. Additional important directions can be found in Forges, Mertens, and Vohra(2002), de Clippel and Minelli (2004), Ichiishi and Yamazaki (2006), Biran and Forges (2009) and references therein.

Finally, observations from experimental game theory suggest that players have a tendency to play fair, sometimes even if it is against their selfish material interest. Based on these findings it seems that the coco equilibrium, which is fair in monetary value, may be focal among all the equilibria of a bargaining game. For detailed discussion on related experimental papers we refer the reader to Roth (1979), Rabin(1993) Binmore (1994) Fehr and Schmidt (1999), Camerer (2003), Chaudhuri(2008) and references therein.

⁴see for example XXXXX

⁵Among other reasons, the TU assumption in this paper circumvent the need to take positions on competing bargaining axioms.

⁶In particular, the extension of Selten's value would give a solution which does not maximize expected payoffs sum conditional on the joint information, whereas the coco value does satisfy this *first-best* notion of efficiency.

1.2. Illustrative example. Consider two hot-dog (dog) sellers, called $P1$ and $P2$, located in a town with two selling venues: the airport, A , and the beach, B . The demand at A is for 80 dogs, but the demand at B depends on the weather: If sunny, which has probability 0.1, it is for 155 dogs; and if cloudy, which has probability 0.9, it is only for 30 dogs. Each seller has to choose a location without any knowledge of the opponent's choice. If they choose different locations, they each sell the quantity demanded at their respective locations, and if they choose the same location, they split the local demand equally.

The situation is not symmetric; $P2$ is more productive and $P1$ is better informed. Specifically, $P2$ nets \$2/dog sold while $P1$ only nets \$0.20/dog sold. On the other hand $P1$ has a perfectly accurate weather forecast while $P2$ has no information about the weather other than the probabilities above.

This situation may be described by a Bayesian game with the payoff tables below:

Sunny prob .1			Cloudy prob .9		
	A	B		A	B
A	8,80	16,310	A	8,80	16,60
B	31,160	15.5,155	B	6,160	3,30

The expected payoffs, obtained under three different computational schemes, are summarized in the table below:⁷

	P1	P2	Total
Non-cooperative, Bayesian equilibrium	10.3	88	\$98.3
Purely cooperative play	7	175	182
The coco value: cooperate & transfer	55	127	182

At the unique Bayesian-equilibrium of this game, $P1$ chooses B when he knows the weather to be sunny and A when he knows it to be cloudy; and having no such knowledge, $P2$ simply chooses A . The expected payoffs in the table are computed to be: $(10.3, 88) = 0.1(31, 160) + 0.9(8, 80)$.

The purely-cooperative payoffs are obtained by making coordinated optimal use of their combined information-production resources, in order to maximize the total payoffs: when it is sunny $P2$ goes to B and $P1$ goes to A , and when it is cloudy they reverse the locations. The expected payoffs in the table are computed to be: $(7, 175) = 0.1(16, 310) + 0.9(6, 160)$.

Clearly, they would like to obtain the total cooperative profit of \$182 rather than the non-cooperative total payoff of \$98.3. But the cooperative solution calls for substantial sacrifices on $P1$'s part: always disclose his forecast truthfully and then choose the inferior location.

A natural resolution is to amend the efficient solution with payoff transfers. In the example above, the coco value prescribes that on sunny days, when $P2$ goes to B and $P1$ to A , $P2$ pays him \$111 out of her \$310 payoff. And on cloudy days, with him at B and her at A , $P2$ pays him \$41 out of her \$160 payoff. The expected payoffs in the table are computed to be: $(55, 127) = 0.1(16 + 111, 310 - 111) + 0.9(6 + 41, 160 - 41)$.

In what follows we present the following: (1) A definition and a formula for computing the coco value, as done above. (2) Axioms of fairness and efficiency that justify this value. And (3) an implementing *partnership game*, in which the players have the incentives to disclose their private information and act optimally in order to bring about the payoffs above.

⁷But unlike this example, as we discuss later, there are games where the coco payoff of an individual is lower than her equilibrium payoff.

2. PRELIMINARIES

Unless otherwise specified, we consider games with a fixed set of two players, $N = \{1, 2\}$. A Bayesian game $G = (A = \times_{i \in N} A_i, \Omega, T = \times_{i \in N} T_i, \tau, u)$, consists sets of *actions* A_i and *types* T_i for each player i , a set of *states of nature* Ω , a *prior* probability distribution τ over the set of *states of the world* $\Theta \equiv \Omega \times T$, and a payoff function $u : A \times \Theta \rightarrow \mathbf{R}^N$. We consider exclusively finite games, where each of the sets above are finite, and for brevity, we sometimes denote the game by $G = (A, \Theta, \tau, u)$.

Following standard formulations, we assume that the game and the prior distribution are commonly known to the players. Game play is as follows. First, the state of the world, $\theta = (\omega, t_1, t_2) \in \Theta$, is drawn from τ . Each player i then observes her own type t_i , based upon which she (simultaneously with her opponent) chooses an action $a_i \in A_i$. The payoff to player i is $u_i(a, \theta)$ where the *action profile* a is (a_1, a_2) .

As standard, a *mixed action* for player i is a probability distribution $\alpha_i \in \Delta(A_i)$ over the set of actions; a (pure) *strategy* for player i , $s_i : T_i \rightarrow A_i$, is a function that specifies what player i would do if her type is $t_i \in T_i$; and a (behavioral) *mixed strategy* for player i , $\sigma_i : T_i \rightarrow \Delta(A_i)$ similarly specifies a mixed action to play based upon knowledge of own type. Payoffs are extended by the standard use of expected values.

For information profile $t \in T$ and action profile $a \in A$, we define the conditional expected payoff of player i to be

$$u_i(a|t) = \mathbf{E}[u_i(a, \theta)|t] = \frac{\sum_{\omega \in \Omega} u_i(a, (\omega, t)) \tau(\omega, t)}{\sum_{\omega \in \Omega} \tau(\omega, t)}.$$

It is also common to refer to t_i as the *private information* of player i . A game is *zero-sum* if $u_1 = -u_2$ and it is a *team game* if $u_1 = u_2$. Finally, we use the standard convention that a_{-i} and u_{-i} represent the actions and payoffs of player i 's opponent.

As mentioned, we implicitly assume that the players have linear transferable utility for money (TU), i.e., they can make arbitrary monetary side payments (or their equivalent) from one to another at a one-to-one rate.

The game is said to have *revealed payoffs* if, after play of the game, the entire payoff vector, $u(a, \theta)$, is revealed to all players. The game is said to have a *revealed payoff table*, if, after the play of the game, the entire realized payoff function, $u(\cdot, \theta)$ is revealed to all players. That is, each player knows how much each player would have received had they played any other action profile.

3. THE COCO VALUE: A FORMULA FOR FAIR AND EFFICIENT EXPECTED PAYOFFS

The coco value is a unique pair of numbers for each game $G = (A, \Theta, \tau, u)$. It is Pareto efficient, which in such a TU game means that it maximizes the sum of payoffs, and it fairly reflects the strategic positions and contributions of the players. In the case of complete information games, it coincides with earlier variable-threat bargaining solutions (Nash, Kalai-Smorodinsky, and Egalitarian), as described in the introduction.

For the extension to Bayesian games and for the construction of non-cooperative protocols that follow, it is better to use a different definition of this value (even in the complete information case) than the ones used by the earlier authors. This definition uses a natural decomposition of a strategic game into a cooperative and a competitive components.

3.1. Complete information. For clarity of exposition, we first give a definition for the case of a complete-information bimatrix game (X, Y) , where matrices $X, Y \in \mathbf{R}^{m \times n}$ represent the payoffs, i.e., $u(i, j) = (x_{ij}, y_{ij})$ for $(i, j) \in \{1, 2, \dots, m\} \times \{1, 2, \dots, n\} = A_1 \times A_2$. As discussed in the

introduction, (X, Y) can be uniquely decomposed as the sum of an equal-payoff team game E and a zero-sum payoff-advantage game Z . Specifically,

$$(X, Y) = (E, E) + (Z, -Z) = \left(\frac{X+Y}{2}, \frac{X+Y}{2} \right) + \left(\frac{X-Y}{2}, \frac{Y-X}{2} \right).$$

The coco value of (X, Y) is defined by,

$$\kappa(X, Y) = (e^*, e^*) + (z^*, -z^*),$$

where $e^* = \max_{ij} e_{ij}$ is the natural *team value* for a game of the form (E, E) , and z^* is the min-max value of $(Z, -Z)$, the classical von Neumann solution of a zero sum game.

Since (e^*, e^*) is efficient and $(z^*, -z^*)$ is a zero-sum transfer, the coco value is efficient. The $(z^*, -z^*)$ transfer is a (positive or negative) correction transfer that reflects the asymmetries in the original game, ignored by the equal payoff component (e^*, e^*) .

The following is a direct consequence of the above definition.

Observation 1. *For any game of complete information,*

- (1) *The coco value is feasible in that there is always a simple agreement consisting of a pair of (pure) actions and monetary transfer, which yields net payoffs equal to the coco value.*
- (2) *For any zero sum game the coco value $(A, -A) = \min \max$ value $(A, -A)$.*
- (3) *For any team game the coco value $(A, A) = \text{team value}(A, A) = (\max_{ij} a_{ij}, \max_{ij} a_{ij})$.*

3.2. Incomplete information. Proceeding to the general case of Bayesian games, we first define two auxiliary payoff functions.

Definition 1. (1) *The equal (share), or average payoff of player i is $u_i^{\text{eq}}(a, \theta) \equiv \frac{u_1(a, \theta) + u_2(a, \theta)}{2}$, and*
 (2) *the payoff advantage of player i is $u_i^{\text{ad}}(a, \theta) \equiv u_i(a, \theta) - u_i^{\text{eq}}(a, \theta) = \frac{u_i(a, \theta) - u_{-i}(a, \theta)}{2}$.*

The decomposition presented above for complete information games is now described by the following immediate observation.

Observation 2. *For every (a, θ) , $u_i(a, \theta) = u_i^{\text{eq}}(a, \theta) + u_i^{\text{ad}}(a, \theta)$, $u_1^{\text{eq}}(a, \theta) = u_2^{\text{eq}}(a, \theta)$ and $u_1^{\text{ad}}(a, \theta) + u_2^{\text{ad}}(a, \theta) = 0$.*

Because $u_1^{\text{eq}}(a, \theta) = u_2^{\text{eq}}(a, \theta)$, we may simply omit the player subscript and write $u^{\text{eq}}(a, \theta)$.

But unlike the complete information case, now the players may also improve the sum (or average) of their expected payoffs by sharing information. To this end, we use the following notions.

Definition 2. *An optimal coordinated (pure) strategy is a rule $c : T \rightarrow A$ such that for every pairs t and a ,*

$$\mathbf{E}_\omega [u^{\text{eq}}(c(t), (\omega, t))] \geq \mathbf{E}_\omega [u^{\text{eq}}(a, (\omega, t))].$$

The team optimum of G , $\text{team-opt}(G) = E_{\theta=(\omega, t)} [u^{\text{eq}}(c(t), \theta)]$, where c is any optimal coordinated strategy.

The conditional team optimum for a $t \in T$ is defined by $\text{team-opt}(G|t) = \max_{a \in A} u^{\text{eq}}(a|t)$. Note that $\text{team-opt}(G) = \mathbf{E}_t[\text{team-opt}(G|t)]$.

In words, an optimal coordinated strategy is a rule c that the players may use to select, for every pair of types t , a pair of actions $c(t)$ in order to maximize their (equal) portions of the total expected payoff in G . The team optimum is the maximal expected payoffs that may be generated by such a rule.

The *relative advantage* of player i is defined to be her minmax value in the zero-sum game $G^{\text{ad}} = (A, T, \Omega, \tau, u^{\text{ad}})$ obtained by replacing each player j 's payoff function, u_j , by her advantage

payoff function, u_j^{ad} , with all other parameters unchanged. The game G^{ad} is a zero-sum modification of G , which preserves the differences between the two players' payoffs. Each player is simply trying to maximize the difference between her payoff and that of her opponent. Since the advantage game is a zero-sum Bayesian game, it has a unique minmax expected value, which we denote by $\min\max_i(G^{\text{ad}})$. We refer to this value as the player's *competitive advantage, relative advantage* or just *advantage*.

Definition 3. *The coco value of G to player i , denoted by $\kappa_i(G)$, is defined by,*

$$\kappa_i(G) = \text{team-opt}(G) + \min\max_i(G^{\text{ad}}).$$

In parallel to the complete information case above, one may define a cooperative component of G , G^{eq} , in which the players share both, the information they have coming to the game and the payoffs resulting from any play.⁸ The $\text{team-opt}(G)$ equals the highest possible (common) expected payoff that may result from any pure strategy of G^{eq} , which may be thought of as the *team-value*(G^{eq}). Thus, in parallel to the complete information game one may think of the coco value of a private information game as the sum: $\kappa(G) = \text{team-value}(G^{\text{eq}}) + \min\max\text{-value}(G^{\text{ad}})$.

Note that the coco value is feasible and Pareto optimal for every t (in conditional expectations over Ω), i.e., the sum of the payoffs is the maximum (expected) sum the players can achieve with coordination and sharing of information. We now argue that it is the “right” value by the axiomatic approach.

4. AXIOMATIC CHARACTERIZATION OF THE COCO VALUE

A *value* for games of incomplete information is a function from the set of all finite two-person games of incomplete information to \mathbf{R}^2 , i.e., $v(G) \in \mathbf{R}^2$ where $v_i(G)$ is the value to player i . We will argue that a small number of axioms uniquely imply the coco value. The strongest such axiomatization has as few axioms that are as weak as possible. However, the coco value satisfies a great number of properties inherited from those properties of zero-sum games. In the concluding sections, we discuss several further appealing properties which the coco value satisfies. The following properties are sufficient for our theorem.

- (1) **Pareto efficiency.** Players achieve the maximum total payoff possible with shared information (first-best). $v_1(G) + v_2(G) = \max_{c:T \rightarrow A} \mathbf{E}_{\theta=(\omega,t)}[u_1(c(t), \theta) + u_2(c(t), \theta)]$.
- (2) **Shift invariance.** Shifting payoffs by constants in every cell leads to a corresponding shift in the value. Let $w \in \mathbf{R}^2$ and $G' = (A, \Theta, \tau, u')$ where $u'(a, \theta) = u(a, \theta) + w$ for all $a \in A, \theta \in \Theta$. Then $v(G') = v(G) + w$.
- (3) **Monotonicity in actions.** Removing an action of a player cannot increase her value. Formally, let $G' = (A'_1 \times A_2, \Theta, \tau, u)$ where $A'_1 \subseteq A_1$. Then $v_1(G') \leq v_1(G)$. Similarly, for player 2.
- (4) **Payoff dominance.** If, for any coordinated pure strategy, a player's expected payoff is strictly larger than her opponent's, then her value should be at least as large as the opponent's. In particular, if $\max_{c:T \rightarrow A} \mathbf{E}_{\theta=(\omega,t)}[u_1(c(t), \theta) - u_2(c(t), \theta)] > 0$, then $v_1(G) \geq v_2(G)$. Similarly, for player 2.
- (5) **Invariance to replicated strategies.** Let $\sigma_1 : T_1 \rightarrow \Delta(A_1)$ be a mixed strategy for player 1. Then we can represent σ_1 as an explicit pure action for player 1 without changing the value of the game. Formally, let $a'_1 \notin A_1$ and suppose $G' = ((A_1 \cup \{a'_1\}) \times A_2, \Theta, \tau, u')$ where u' is an extension of u such that $u'(a'_1, a_2, \omega, t) = u(\sigma_1(t_1), a_2, \theta)$ for all $a_2 \in A_2, \theta \in \Theta$. Then $v(G') = v(G)$. Similarly for player 2.

⁸ $G^{\text{eq}} = (A, \Omega, \hat{T}_1 \times \hat{T}_2, \hat{\tau}, \hat{u})$ defined by $\hat{T}_1 = \hat{T}_2 = T_1 \times T_2$, $\hat{\tau}(\omega, \hat{t}_1, \hat{t}_2) = \tau(\omega, \hat{t}_1)$ if $\hat{t}_1 = \hat{t}_2$ ($= \mathbf{0}$ otherwise); and $\hat{u}_i(a, \omega, \hat{t}) = u_i^{\text{eq}}(a, \omega, \hat{t}_i)$.

- (6) **Monotonicity in information.** Giving player i strictly less information cannot increase her value. Formally, let $f : T_i \rightarrow T_i$ be any function and replace the type t_i by $t'_i = f(t_i)$. Then, in the new game $G' = (A, T, \Omega, \tau', u)$, $v_i(G') \leq v_i(G)$.

Theorem 1. *The coco value is the only value that satisfies axioms 1-6 above.*

In order to prove this theorem, it is helpful to consider what we call the *Raiffa game*, $R(G)$, in the spirit of Raiffa (1953). In the Raiffa game, the players choose actions for G based on their information. In addition to choosing a , with mutual consent they may improve both of their expected payoffs. In particular, they may improve their expected total to $\text{team-opt}(G|t)$, where the increase is done equally to both players. Formally, the Raiffa game is $R(G) = ((A_1 \times \{0, 1\}) \times (A_2 \times \{0, 1\}), T, \Omega, \tau, u')$ for u' defined as follows:

$$u'_i[(a_1, b_1), (a_2, b_2), \theta] = \begin{cases} u_i((a_1, a_2)|t) + r((a_1, a_2), t) & \text{if } b_1 = b_2 = 1 \\ u_i(a_1, a_2, \theta) & \text{otherwise} \end{cases}.$$

The function $r(a, t)$ is such that $u'_1(a|t) + u'_2(a|t) = \text{team-opt}(G|t)$. Now, the Raiffa game is a hypothetical game we consider. It is perhaps interesting to note that while the two players cannot quite play the Raiffa game even upon agreement, through side payments (and agreeing upon optimal actions for each $t \in T$) they can play a game which has equivalent payoffs in expectation. We first observe that the Raiffa game has the same value as G .

Lemma 1. *For any finite two-person game G , and any value v that satisfies axioms 1-6 above, $v(G) = v(R(G))$.*

Proof. Take the Raiffa game and force player 1 to play $b_1 = 0$, i.e., remove player 1's ability to agree to change the payoffs. This leaves us in a game G' with $v_1(G') \leq v_1(R(G))$. Moreover, G' is equivalent to G , but with each action for player 2 being duplicated⁹ – it is irrelevant whether player 2 chooses $b_2 = 0$ or $b_2 = 1$. So using invariance to strategic replications and Pareto optimality, $v_1(G) = v_1(G')$, and combining this with the observation just above, we conclude that $v_1(G) \leq v_1(R(G))$. A similar argument shows that $v_2(G) \leq v_2(R(G))$. Finally, it is easy to see that the team-optima of G and $R(G)$ are the same, so $v_1(G) + v_2(G) = v_1(R(G)) + v_2(R(G))$, and hence $v(G) = v(R(G))$. \square

We are now ready to prove Theorem 1.

Proof. First, we argue that the coco value satisfies axioms 1-6. Pareto efficiency is trivially guaranteed by the fact that the advantage game is zero-sum and the team game achieves maximum expected payoff sum. Second, a payoff shift of (w_1, w_2) corresponds to a shift of $(\frac{w_1 - w_2}{2}, \frac{w_2 - w_1}{2})$ in the advantage game and a shift of $\frac{w_1 + w_2}{2}$ in the team-opt. Since the value of zero-sum Bayesian games satisfies shift invariance, this corresponds to a shift of (w_1, w_2) in the coco value. Monotonicity in actions and information clearly holds for zero-sum games and the team-opt value, hence also for the coco value. Similarly, adding an explicit mixed action for i in G corresponds to adding an explicit mixed action in the advantage game, which we know does not change the value (for zero-sum games), and adding an explicit mixed action also clearly does not change the team-opt.

It remains to show that the value must agree with the coco value. For any G we design a sequence of games, G^0, G^1, G^2, G^3, G^4 such that, $v_1(G) = v_1(G^0) + \kappa_1(G)$ and $v_1(G^0) = v_1(G^1) \geq v_1(G^2) \geq v_1(G^3) \geq v_1(G^4) \geq 0$ and an entirely similar argument shows that $v_2(G) \geq \kappa_2(G)$. Hence, by Pareto optimality, we must have $v(G) = \kappa$.

⁹This is not exactly correct since it is equivalent to a game that has action sets $A_1 \times \{1\}$ and $A_2 \times \{1, 2\}$, which is not formally G with A_2 duplicated since the actions have been *renamed* for both players. However, a double-application of axiom 3 implies that renaming any action does not change the value of a game for either player.

- (1) The game G^0 is the game G with the payoffs translated by $-\kappa(G)$, so that $\kappa(G^0) = (0, 0)$ and hence $\text{val}(Z(G^0)) = (0, 0)$ and the team optimum is 0. By translation invariance, we have $v_1(G^0) = v_1(G) - \kappa_1(G)$, so it suffices to show that $v_1(G^0) \geq 0$.
- (2) Let σ_1 be any minmax optimal strategy for player 1 in the zero-sum game $(G^0)^{\text{ad}}$. The game G^1 is the same as G^0 except that we create an explicit pure action a'_1 for σ_1 , i.e., $u^1(a'_1, a_2, \omega, t) = u^0(\sigma_1(t_1), a_2, \theta)$ for all $a_2 \in A_2, \theta \in \Theta$. Since we know σ_1 guarantees player 1 an expected nonnegative payoff in the game where we look at the difference between the two players' payoffs in G^0 , a'_1 must guarantee player 1 an expected payoff in G^1 at least as large as that of player 2. Since we have only made a mixed strategy explicit, we have $v_1(G^1) = v_1(G^0)$.
- (3) The game G^2 is $R(G^1)$, the Raiffa version of game G^1 . We have already argued above that $v(G^2) = v(G^1)$.
- (4) In the game G^3 , we remove all of player 1's actions except for $(a'_1, 1)$, an action in the Raiffa game $G^2 = R(G^1)$. This can only decrease player 1's value, i.e., $v_1(G^3) \leq v_1(G^2)$.
- (5) In the game G^4 , we remove all of player 1's information, i.e., we modify τ by making t_1 take a single value with probability 1. Hence G^4 is a game in which player 1 has only one action and less information, which again can only decrease player 1's value, i.e., $v_1(G^4) \leq v_1(G^3)$.

We now observe that payoff dominance in combination with the translation invariance axiom imply the following slightly stronger version of payoff dominance: "If for all $c : T \rightarrow A$, $\mathbf{E}_\theta[u_1(c(t), \theta) - u_2(c(t), \theta)] \geq 0$, then $v_1(G) \geq v_2(G)$." (The difference compared to aforementioned axiom is the weak rather than strict inequality.) This is because the payoffs could be shifted by $(\epsilon, 0)$ and then we would have that the value to player 1 is at least as large as that of player 2 minus ϵ . Since this holds for any $\epsilon > 0$, it must also hold for $\epsilon = 0$.

It remains to show that $v_1(G^4) \geq 0$, which we will do using this stronger version of payoff dominance. Since a'_1 guarantees player 1 an expected amount at least as large as player 2 in G^1 , regardless of player 2's strategy, $(a'_1, 1)$ also guarantees player 1 an expected amount at least as large as player 2 in G^2 , regardless of player 2's strategy, because the Raiffa game only modifies payoffs by increasing expected payoffs equally. However, in order to use the payoff comparison axiom, we need something stronger. We need that, even if the two players coordinate based on their joint information, they cannot make player 1's expected payoff less than that of player 2. This stronger condition holds as well since there is no way to coordinate with player 1, who has no useful information (only one type) and only one action. Hence, by payoff comparison, $v_1(G^4) = 0$. \square

5. NON COOPERATIVE IMPLEMENTATION OF THE COCO VALUE

The coco value has many desirable properties, but can it be implemented when players act strategically? In the case of complete information games, it is simple to implement coco through a binding agreement, such as a contract in which the two players commit to play in any particular cell with maximum payoff sum and then make an appropriate side payment. Moreover, as pointed out below, the coco value in this case is *individually rational*, i.e., it lies above each player's minmax safety value.

In the case of incomplete information, however, players must reveal private information in order to achieve efficiency. For example, consider a buyer and a seller negotiating over the price of a valuable piece of art, for which they have private monetary valuations. It may be impossible for one player to ever verify the valuation of the other.

To overcome such difficulties, we adopt some of the methods used in real-life formations of partnerships. When two partners agree to share equally the total (net) realized profits of a joint venture, they create individually monotonic payoff functions: the payoff of each increases if the

total realized profit increases.¹⁰ This monotonicity property gives each partner the incentive to truthfully share information and to take actions that are optimal for the success of the project.

But if the situation is not symmetric, for example there are differences in information, resources and opportunity costs, the partners may agree up-front to make a compensating payoff transfer. If the amount of the payoff transfers is determined prior to the actual participation in the joint venture, then the incentives to cooperate in the joint project are not affected by the outcome of the bargaining over payoff transfers.

We use these ideas below, to construct incentives-compatible protocols that implement the coco value in some restricted but important classes of games. We first give a simple mechanism that achieves the coco value in expectation, under a revealed payoff assumption, which is incentive compatible as long as the players make the binding agreement before observing their types. We then give a mechanism that is more realistic in that it is incentive compatible *interim*, even after the players have observed their private information. However the latter mechanism requires a stronger revealed payoff table assumption.

In the rest of this section, $G = (A, \Theta, \tau, u)$ is assumed to be a fixed arbitrary two-player finite Bayesian game as discussed above.

5.1. *Ex-ante* incentive compatibility: individual rationality of the coco value. We make two assumptions in this section: (1) both realized payoffs are revealed to both players after the game is played, and (2) that the players commit to the protocol before observing their types. The protocol is simple. The players form a partnership in which they split the total payoffs (positive or negative) equally. This can always be achieved by a side payment, and it incentivizes them to coordinate by revealing information and playing actions that maximizes the total payoff. However, separate from this partnership, a second corrective side payment is made which serves the purposes of (1) achieving the coco value, and (2) making cooperating individually rational. Of course these two side payments may be combined into one.

Ex-ante partnership protocols for an arbitrary finite two player Bayesian game $G = (A, \Theta, \tau, u)$. Fix any *optimal coordinated strategy* $c : T \rightarrow A$.¹¹

- (1) Players simultaneously choose whether to commit to participate or not.
 - If either does not agree to participate, then they play G unmodified, they collect their respective G -payoffs, and the protocol ends.
 - Otherwise, they have made a binding agreement to continue as follows:
 - (2) A triple (ω, t_1, t_2) is drawn by the prior distribution τ and each player i is informed of her *realized type* t_i .
 - (3) Both players, $i = 1, 2$, simultaneously declare their own types \tilde{t}_i .
 - (4) The players are forced to play the pair of actions $a = c(\tilde{t})$, after which the pair of payoffs $u(a, \theta)$ is revealed.
 - (5) A side-payment is made so that the net payoff to player i is $u^{cq}(a, \theta) + \text{val}_i(G^{\text{ad}})$. In other words, he is paid one half of the total payoffs obtained through the actual play in stage 4, plus his minmax value (positive or negative) of the advantage component-game of G , computed without knowledge of the types.
-

Theorem 2. *The pair of coco values $\kappa(G)$ of any finite two-player Bayesian game $G = (A, \Theta, \tau, u)$ are the expected payoffs of a Nash equilibrium in any ex-ante partnership protocol of the game.*

¹⁰The use of such monotonicity conditions is common in cooperative game theory, see for example Kalai (1997) and Myerson and Thomson (1980).

¹¹Recall that for every pair of types t and for pair of actions a , $\mathbf{E}_\omega [u^{cq}(c(t), (\omega, t))] \geq \mathbf{E}_\omega [u^{cq}(a, (\omega, t))]$.

Proof. Consider the following equilibrium strategy for each player i :

- choose to participate,
- if mutual participation fails, play the (mixed) minmax strategy of G^{ad} , i.e., play G as if you were playing G^{ad} , and
- if mutual participation holds, truthfully reveal your realized type, i.e., $\tilde{t}_i = t_i$.

Observe first that no player can benefit by declaring a false type $\tilde{t}_i \neq t_i$ (given that the other player is being honest), because $\tilde{t} = t$ simultaneously maximizes each player's expected payoff (it maximizes $\frac{u_1(a,\theta)+u_2(a,\theta)}{2}$ and has no effect on $\text{val}_i(G^{\text{ad}})$).

Next, observe that player i cannot increase her payoff by not participating. Say that she does not participate, and instead plays a mixed strategy σ_i , while her opponent plays his minmax strategy of the game G^{ad} , σ_{-i} . By the decomposition, i 's expected payoff is the sum of the expected payoffs in G^{ad} and G^{eq} . The expected payoff of σ in G^{ad} is at most the (min-max) value of G^{ad} , and the expected payoff of σ in G^{eq} is at most the (team) value of G^{eq} , hence the total is at most the coco value for i . \square

Corollary 1. *The coco value is individually rational.*

Proof. The payoff that a player can guarantee herself in the ex-ante partnership protocol is at least as high as what she can guarantee herself in the game G . This is so because by choosing not to participate, she forces the protocol game to coincide with G without any restrictions on her G strategies. And since the coco value is an equilibrium payoff of the protocol game, it must yield each player at least her protocol minmax payoff. \square

While the protocol above illustrates the individual rationality of the coco value, it may be unsatisfactory for two reasons:

- The players must not know their types before committing to play. If either player knows some information about their type before step 1, it may no longer be in their best interest to participate. Hence, the protocol is not *interim* incentive-compatible.
- Wilson(1987) advocated that one should design mechanisms that are implementable without dependency on the prior probability distribution. The mechanism above violates the Wilson doctrine in two respects. First, in order to compute the actions which maximize the expected payoffs, one must know the prior. Second, in order to compute the value of G^{ad} , one must know the prior as well.

These deficiencies are overcome in the next section, under a further restriction on the environment.

5.2. *Interim incentive compatibility.* In this section we make the assumption that payoffs and the *payoff table* are revealed after the play of the game, meaning that both the realized payoffs, $u(a, \theta)$, and the entire realized payoff function $u(\cdot, \theta)$, from A to \mathbf{R}^2 , are revealed to all players, for the true state of the world θ . This is a strong assumption, but in many cases, described in Section 5.4, it may be relaxed.

For an example of an environment that fits this assumption, think again of the hot dog sellers game from the introduction, and assume that both sellers receive partial weather reports (their types) before deciding on a location. The payoffs in this example depend on the weather, and not on the reports, and once the weather is observed the profit in each location (whether chosen or not) is known. In other words, the entire payoff table for the realized θ becomes known, even if the full θ itself does not (e.g., the individual reports may remain unknown).

Under the above assumption, and some weaker ones to follow, one can design effective interim protocols to implement the coco value. The idea is simple, and again, it is similar to the formation of some real-life partnerships.

Interim partnership protocol for an arbitrary finite two-player Bayesian game $G = (A, \Theta, \tau, u)$.

- (1) A triple (ω, t_1, t_2) is drawn by the prior distribution τ and each player i is informed of her realized type t_i .
 - (2) Simultaneously each player selects one strategy from the following two choices:
 - DO NOT PARTICIPATE - She declares NO and selects an action $\tilde{a}_i \in A_i$, to serve as her non cooperative action; or,
 - PARTICIPATE- she declares YES and submits a sealed envelope containing a reported type $\tilde{t}_i \in T_i$ and a selected action $\tilde{a}_i \in A_i$, to serve as her non cooperative action.
 The YES/NO declarations are revealed to both players and then:
 - If either player declares NO, then the non cooperative action \tilde{a}_i selected above are played, the game stops and the players collect their respective G -payoffs.
 - But if both declare YES, then the reported types \tilde{t} are revealed to both players, who are committed to continue as follows.
 - (3) Simultaneously, the players choose actions $a_i \in A_i$ and play G using a . Both $u(a, \theta)$ and the realized payoff function $u(\cdot, \theta)$ are then revealed to both players.
 - (4) Using the cooperative actions a_i from stage 3 and the non cooperative actions \tilde{a}_i from stage 2, side payments are made so that the net payoff to each player i is, $u_i^{\text{eq}}(a, \theta) + u_i^{\text{ad}}(\tilde{a}, \theta)$.
-

In stage 3, one might expect the players to play an optimal coordinated strategy. However, we have not given them any explicit means of coordination. It would be natural to incorporate coordination devices such as cheap talk, in the model. However, this is not necessary for the theorem below.

In the equilibrium discussed in this theorem players choose to optimally share information and coordinate, with an optimal threat defined through the relative advantage game.

Definition 4.

- (1) A strategy π_i of the partnership protocol above is participatory, if it declares YES (with probability one) for every t_i ; and it is honest, if $\tilde{t}_i = t_i$ for every t_i .
- (2) For an optimal coordinated (see Definition 2) strategy $c : T \rightarrow A$ of the game G , a profile of strategies π in the partnership protocol is c -coordinated if:
 - A. In stage 2 each player declares YES, uses her minmax strategy of G^{ad} to choose \tilde{a}_i , and truthfully selects $\tilde{t}_i = t_i$.
 - B. In stage 3 each player selects $a_i = c_i(\tilde{t})$, provided that she had reported truthfully $\tilde{t}_i = t_i$, as planned in stage 2. If she failed to report truthfully in stage 2, i.e. $\tilde{t}_i \neq t_i$ (a probability zero event), then she selects a_i which maximizes $u_i^{\text{eq}}(a_i, c_{-i}(\tilde{t}) | (t_i, \tilde{t}_{-i}))$.¹²

Clearly a c -coordinated strategy is participatory, honest, ex post efficient, and have the additional appealing properties described by the following theorem.

Theorem 3. Consider the interim partnership protocol of a given finite two-player Bayesian game $G = (A, \Theta, \tau, u)$:

- (1) Any c -coordinated strategy profile is a sequential Nash equilibrium of the partnership protocol with expected payoffs that equal the coco value of G , $\kappa(G)$.
- (2) For any participatory Nash equilibrium of the partnership protocol the expected payoffs are $\kappa(G) - (x, x)$ for some $x \geq 0$. In other words, all participatory equilibria are pareto dominated by the coco payoffs.

¹²The specification of how to act under such zero-probability events is needed in order to argue, as we do below, that we have a sequential equilibrium.

- (3) *However, the equilibria of G also remain equilibria of the partnership protocol: for any mixed-strategy Nash equilibrium $\sigma = (\sigma_1, \sigma_2)$ of G , it is also Nash equilibrium of partnership protocol in which both players declare NO and select \tilde{a}_i according σ_i .*

Note that in the above protocol, each player has the option to play G as is, without participating. As a result, as mentioned in the above theorem, any Nash equilibrium of G can be converted to a nonparticipatory equilibrium of the interim partnership protocol. While it is tempting to try to “implement away” these equilibria (especially when they are Pareto dominated), we feel that it is also reasonable to model the possibility that players may choose not to participate and to play a (noncooperative) equilibrium of the underlying game. However, in some games such as the prisoner’s dilemma, participation is a dominant strategy.

Proof of Theorem 3. Notice first that by definition, in every one of player 1’s stage-3 information sets (and similarly player 2), player 1 acts optimally, under the assumption that her opponent tells the truth and follows the c -optimal selection.

So, for part 1, it remains to observe that player 1 acts optimally at her first information set, namely in stage 2. Suppose that instead of the above, she chooses not to participate, and to plays \tilde{b}_1 . By the decomposition, her overall payoff is:

$$u_1^{ad}(\tilde{b}_1, \tilde{a}_2|t) + u_1^{eq}(\tilde{b}_1, \tilde{a}_2|t).$$

But switching her strategy to a c -coordinated one, her payoff may be written as:

$$u_1^{ad}(\tilde{a}_1, \tilde{a}_2|t) + u_1^{eq}(c(t)|t),$$

where \tilde{a}_1 is chosen by her advantage-game minmax strategy against his minmax chosen \tilde{a}_2 strategy. In addition, $u_1^{eq}(c(t)|t) \geq u_1^{eq}(\tilde{b}_1, \tilde{a}_2|t)$. So the switch to a c -coordinated strategy can only increase both terms in her payoff.

It is also easy to see, using the same decomposition argument above, that under participatory strategy she cannot obtain a higher payoff than by following a c -coordinated strategy. Thus, part 1 of the theorem holds.

For part 2 observe again by the decomposition that at any equilibrium the players’ first payoff terms must equal their corresponding first payoff terms under the c -coordinated equilibrium. On the other hand, their second (equal) payoff terms can only be smaller than under the c -coordinated equilibrium, and by the same amount.

Part 3 is obvious, since either player can declare NO, forcing the game to be the original game G . \square

Remark 1. *Notice that while the implementation theorem is stated in terms of sequential Nash equilibrium, the implementing strategies rely on more solid solution concepts: To determine the pair \tilde{a} the players use the minmax solution (as opposed to just a Nash equilibrium), and to determine the actual action pair a they use simple one person optimization. Both, the minmax solution and one person optimization, have better stability, robustness and computability properties than that of Nash equilibrium. Some of these issues are discussed in the concluding section.¹³*

Remark 2. *The coco value is individually rational ex ante, but would players who already know their types agree to it? This falls into the general question addressed by Harsanyi(1967): What do you do when you know your own type, but you play against opponents with unknown types? The implementation theorem above follows the standard Harsanyi’s approach: complete the game of incomplete information by describing all possible past events (together with the I_r probabilities) that lead to the current situation, and then choose actions in the current situation based on an overall equilibrium of the complete game.*

¹³We thank Robert Wilson for making this observation.

The interim protocol above is a complete game (that includes all possible types of the player and his opponents) in which the player (after knowing his type) has to choose whether to participate and play the coco value, or play the game as is. Since there is always a participatory equilibrium in the game, the theorem above offers a positive answer to the question of agreeing to play the coco value.

The next example illustrates that while the coco value is unique, the coco value conditioned on realized types may fail to be unique. This is due to the possible multiplicity of equilibria in the complete game.

Example 1. Player 1 may choose up or down while Player 2 has only one action. Player 2 knows whether the state is "the game on the left" or the "game on the right," while Player 1 only knows that the probabilities are .5,.5.

wp .5	<table style="margin: auto; border-collapse: collapse;"> <tr><td style="padding: 2px 10px;">0, 0</td></tr> <tr><td style="padding: 2px 10px;">2, 0</td></tr> </table>	0, 0	2, 0	wp .5	<table style="margin: auto; border-collapse: collapse;"> <tr><td style="padding: 2px 10px;">0, 0</td></tr> <tr><td style="padding: 2px 10px;">0, 2</td></tr> </table>	0, 0	0, 2
0, 0							
2, 0							
0, 0							
0, 2							

We consider two c-coordinated equilibria of the interim partnership protocol for this game, U and D . In both equilibria Player 2 honestly reports the realized game, Player 1 chooses $a_i = \text{down}$ no matter what is reported, and they generate the cooperative payoffs $u_i^{eq}(a, \theta) = 1$ in either of the true states. The remaining question is what threat action \tilde{a}_1 Player 1 uses? But in the relative advantage component game of our example, described below, every strategy is minmax. So let U be the equilibrium in which Player 1 chooses $\tilde{a}_1 = \text{up}$, and D be the equilibrium in which he chooses $\tilde{a}_1 = \text{down}$.

wp .5	<table style="margin: auto; border-collapse: collapse;"> <tr><td style="padding: 2px 10px;">0, 0</td></tr> <tr><td style="padding: 2px 10px;">1, -1</td></tr> </table>	0, 0	1, -1	wp .5	<table style="margin: auto; border-collapse: collapse;"> <tr><td style="padding: 2px 10px;">0, 0</td></tr> <tr><td style="padding: 2px 10px;">-1, 1</td></tr> </table>	0, 0	-1, 1
0, 0							
1, -1							
0, 0							
-1, 1							

Now, consistent with the fact that the coco value is unique, in both equilibria the expected payoffs are $(1, 1)$. But what about the conditional coco payoffs, for example when the true game is on the left? The overall payoffs are defined by $u^{eq}(a, \theta) + \text{val}_i(G^{ad})$, and as stated above $u^{eq}(a, \theta) = (1, 1)$ in both equilibria. But the minmax payoff, conditional on the being in the left game are different: $\text{val}(G^{ad}) = (0, 0)$ under the equilibrium T , and $\text{val}(G^{ad}) = (1, -1)$ under the equilibrium B .

Thus, the overall conditional coco payoffs, given the left-hand-side game, are either $(1, 1)$ or $(2, 0)$, depending on the equilibrium strategy chosen by Player 1. Similarly, conditional on being in the right-hand-side game, the conditional coco payoffs could be $(1, 1)$ or $(0, 2)$. Clearly, as already stated, the overall coco values, no matter what equilibrium is played, are $(1, 1)$.

The discussion above together with the proof of the theorem, leads to the following theorem.

Theorem 4. If G^{ad} has unique minmax strategies $\sigma^* = (\sigma_1^*, \sigma_2^*)$ for both players, and "cooperation can be strictly beneficial" in the sense that $u^{eq}(\sigma^*|t)$ is strictly Pareto inefficient for every t , then:

- (1) Conditional coco payoffs are unique, and
- (2) The only best responses to a c-coordinated opponent's strategy are participatory strategies.

5.3. Joint venture example: efficiency in the Myerson-Satterthwaite model. A manufacturer, M , can produce a certain item at cost $\$C$, and a distributor, D , can sell this item with a return of $\$R$. The pair of parameters (C, R) is generated by a known prior probability distribution π on the integers in $[0, 100]^2$; M knows the realized value of C and D knows the realized value of R . Under the simple monetary function used in this paper, if M manufactures the item, sells it to D at a price P , who in turn sells with the return R , then M nets the payoff $P - C$ and D nets the payoff $R - P$.

The well known impossibility result of Myerson and Satterthwaite implies that, in general, there is no mechanism that guarantees efficient outcomes: under any negotiation procedure, M and D would fail to agree on a price P in some situation with $C < R$. However, under the assumption of ex-ante observability discussed in this paper, the coco value offers an efficient and fair solution that can be implemented in the interim sense discussed above.

For a complete strategic description of the situation above we use a double-auction non-cooperative Bayesian game G defined as follows: Simultaneously, M submits a demanded price, P^{dem} and D submits an offered price P^{ofr} . If $P^{ofr} < P^{dem}$ then there is no deal and each nets zero payoff. But if $P^{ofr} \geq P^{dem}$ then the item is manufactured by M, sold to D at the price $P^{mid} = (P^{ofr} + P^{dem})/2$, M's payoff is $P^{mid} - C$ and D's payoff is $R - P^{mid}$.

To compute the coco value of G , observe that $u^{eq}((C, R), (P^{dem}, P^{ofr})) = (R - C)/2$, if $C < R$, and it is zero otherwise. Thus, the team-opt(G) = $E[\max(R - C, 0)]/2$.

As for the advantage component, consider the constant strategies $P^{dem} = 100$ and $P^{ofr} = 0$ played respectively by M and D regardless of their values of C and R . Under these strategies in the game G each player guarantees two things: (1) his own payoff is at least zero, and (2) the opponent payoff is not greater than zero. This means that in the advantage game they guarantee themselves payoff advantages of zero, and $(0, 0)$ is the minmax value of the advantage component game.

Applying the definition of the coco value, the two paragraphs above imply the following.

Proposition 1. *The coco value of the joint venture game above yields each player the payoff $E[\max(R - C, 0)]/2$.*

To illustrate the interim implementation of the coco payoffs above, consider the following strategies in the partnership game. After learning their true individual parameters, C and R , both players declare YES, submit the non cooperative strategies $P^{dem} = 100$ and $P^{ofr} = 0$ and report their individual parameters truthfully: $\tilde{C} = C$ and $\tilde{R} = R$. If the reported cost is greater than the reported return, $\tilde{C} > \tilde{R}$, the item is not produced, and each nets a zero payoff. But if $\tilde{C} \leq \tilde{R}$ then M produces the item (at a cost C), D sells it (with a return of R), and a transfer is made so they each net $(C - R)/2$. Notice that the computation of the transfer requires ex-post monitoring of the actual production cost of M and the actual return collected by D.

The computation and implementation of the coco solution above is relatively simple, despite the fact that the bayesian joint venture game may be difficult to analyse by standard methods. The simple analysis and further insights can be obtained under more interesting version of joint venture, as in the following special case.


5.3.1. *One-sided outside options.* Assume that M has an option to produce the item and sell it to some outside buyer at an alternative price a . How does this affect the coco value of M and D? Assume, for simplicity, that C , R and a are common knowledge and that $C < R$. The Figure below describe the coco payoffs of M and D change as we increase a .

Figure 1

5.4. **Implementation with other assessment devices.** In some contexts, it may be possible to implement the coco value, even though the payoff table is not revealed. In many contexts, this may involve engaging in a secondary act that enables one to assess players' relative advantage.

Example 2. *Professional wrestling game.*

Two professional wrestlers are about to participate in a match, where \$1000 will be awarded to the winner and nothing to the loser. Moreover, there is an extra \$500 bonus will be awarded to the two players, divided evenly, if the match is "a good show." We refer to this option as *dancing*

Coco payoffs, when M has an outside option to sell at price a 

since the sequence of moves must be carefully choreographed. A high-level approximate model of this game is the following:

	Fight	Dance
Fight	1000p, 1000(1-p)	1000, 0
Dance	0, 1000	750, 750

Here p is the probability that player 1 would win if the two fought, and the above are the expected payoffs. It is clear that it is a dominant strategy to fight. A simple calculation shows that the coco value of the above game is $(1000p + 250, 1000(1 - p) + 250)$. In the case where p is common-knowledge and there is no relevant private information, the players might adopt one of two simple agreements yielding the coco value. (For example, they might agree that $p \approx 1/2$, i.e., they have roughly equal chances of winning, and each agree to dance.) While this may not be a legally binding contract, such an agreement may be enforced through reputation, repeated play, or various threats.

However, in the case of private information, the players may not agree upon p . For example, each player may have slept well the previous night and woken up feeling especially strong. Instead, they could agree to engage in, say, a scrimmage wrestling match beforehand, whose sole purpose would be to determine the side payment in the real match. Presumably the probabilities of winning in the scrimmage and the real match would be the same. The agreement would be that they would both dance in the actual match, but a side payment would be such that the winner of the scrimmage would get a payoff of 1250 and the loser would get a payoff of 250. This has the property that it matches the coco value, in expectation. Note that this equality holds for *any* type space and any distribution over prior information. Moreover, the protocol is simple enough to be understood by professional wrestlers. Also note that they may choose any other means of determining a side payment, as long as they both agree to it. For example, it may be a convention that the two players merely arm wrestle rather than have a full scrimmage match.

Example 3. *Attack or retreat game.*

Imagine two parties considering engaging in a war over a piece of territory. For concreteness, let us say there are two states $\omega = 1W$ meaning player 1 wins and $\omega = 2W$ meaning player 2 wins.

Each player has a choice of attacking (Att) or retreating (Ret). Suppose that a player prefers to achieve the same outcome without fighting, if possible. To be concrete, say the payoffs in the two states are as follows:

$\omega = \mathbf{1W}$	Att	Ret	$\omega = \mathbf{2W}$	Att	Ret	
	Att	2,-1	3,0	Att	-1,2	3,0
	Ret	0,3	0,0	Ret	0,3	0,0

In this example, because attacking is a dominant strategy in the advantage game, the coco value is $(3 - 3p, 3p)$, where $p = \Pr[\omega = \mathbf{1W}]$. Again note that this is the coco value regardless of the type space or joint distribution over states of the world. Now, if one side was clearly superior and would obviously win, then that side would attack, the other would retreat, and they would achieve the socially efficient total payoff of 3. However, especially in the realistic case of asymmetric information, where each side has private information about their fitness, agreeing on a value *a priori* may be difficult.

However, in some cases, the two parties may be able to engage in a test to see who is likely to win. If the test is conclusive, both parties may stand to gain from performing it. For example, the test may involve a small battle (e.g., *David and Goliath*). Alternatively, one side might engage in a display of prowess in order to attempt to convince the other side that their probability of winning is low. The natural protocol is to use some such device to determine the winner, and then have that player attack and the other retreat, again giving in expectation the coco value of $(3 - 3p, 3p)$.

An interesting feature of the two examples above is that the protocols may make sense even if the players *do not have a common prior*. In particular, when two wrestlers fight, the private information is in fact quite involved, knowing what moves they are particularly good at, and also what they believe the opponent can do. The belief that all of these probabilities are derived from a common prior is certainly questionable in such situations. Nonetheless, it seems perfectly plausible to tell two wrestlers to “go wrestle” and try their hardest to win.

6. FURTHER REMARKS

6.1. On security levels, threats and externalities: comparison with other methods. A common indirect approach to determine players’ cooperative values of a strategic game is through a “bridge” that connect the strategic theory with the cooperative theory. To every strategic game G one associate a cooperative game V^G and adopts some appropriate cooperative solution $\varphi(V^G)$ to yield cooperative values for the G players. When dealing with TU games, as we do in this paper, the associated cooperative game is described by a set function $V^G = (V^G(S))$, in which each $V^G(S)$ is a real number that describes the “worth” of the coalition of players S in the game G .

While the coco value offers a direct cooperative values for the players of strategic games, without the need for a bridge, it still can be interpreted as special case of the bridge approach. We proceed to explain how the coco bridge compares to another common bridge method, see Aumann (1961) for an early example, and Carpenete et al (2005,2006) and Forge, Mertens, and Vohra, (2002) for more recent references.

For a two-person bimatrix games $G = (A, B)$, an associated cooperative game is determined by three numbers, V_{12} , V_1 and V_2 . To have a unique value φ_i associated to each player of G , we consider here the Shapley (1953) value: $\varphi_i \equiv V_i + \frac{1}{2}[V_{12} - (V_1 + V_2)] = \frac{1}{2}V_{12} + \frac{1}{2}(V_1 - V_2)$. Thus, the only question is how to determine the worth of the coalitions, V_{12} , V_1 and V_2 .

The worth of the two player coalition, V_{12} , is naturally defined to be the highest total cooperative payoff that the players may be able to obtain in the game G . But how should we define the worth of

singleton coalitions V_i ? The coco value and the alternative method used by the authors mentioned above differ on this issue.

The alternative used by Aumann and the others is to compute the individual-worth quantities to be $V_1^{\text{Alt}} = \min \max(A, -A)$ and $V_2^{\text{Alt}} = \min \max(B^T, -B^T)$, i.e., the highest payoff that a player can secure herself assuming that her opponent's goal is to minimize her payoff. Under coco value one computes $V_1^\kappa = \min \max(\frac{A-B}{2}, \frac{B-A}{2})$ and $V_2^\kappa = \min \max(\frac{(B-A)^T}{2}, \frac{(A-B)^T}{2})$, i.e., the highest relative payoff *advantage* (over her opponent) that she can secure, assuming that her opponent would act to minimize her payoff *advantage*.

Substituting these individual-worth values into the Shapley formula above, for player 1 for example, we can see clearly the contrasts between the two solutions. For the Aumann alternative approach we have $\varphi_1^{\text{Alt}} = \frac{1}{2}V_{12} + \frac{1}{2}(\min \max(A) - \min \max(B^T))$, whereas for the coco value we have $\kappa_1 = \frac{1}{2}V_{12} + \frac{1}{2} \min \max(A - B)$.

The following two 2×1 games illustrate the difference between the two solutions. Notice that the right hand side game (rhs) differs from the lhs only in the boldface entry.

$$\begin{array}{l} \text{coco value} = \begin{array}{|c|} \hline 1, 1 \\ \hline \mathbf{1}, 0 \\ \hline \end{array} \\ \text{alt. value} = (1.5, .5) \end{array} \qquad \begin{array}{l} \text{coco value} = \begin{array}{|c|} \hline 1, 1 \\ \hline \mathbf{-100}, 0 \\ \hline \end{array} \\ \text{alt. value} = (1, 1) \end{array}$$

In both games $V_{12} = 2$, obtained by her playing up. Notice, however, that the two games show substantial differences in her ability to threaten and extract side payments from him. In the lhs game she can bring Player 2's payoff down from 1 to 0 at no cost to herself. She cannot do so in the rhs game. In other words, his punishment has different externality on her own outcomes.

Should such a difference be reflected in the solution of the game? The coco value reflects this difference by giving her a payoff of 1.5 in the lhs game and only 1 in the rhs game, while φ_1^{Alt} treats the two games identically.

The $\min \max(\frac{A-B}{2})$, used by the coco value, reflects the difference in these externalities, whereas the individual minmax values, $\min \max(A)$ and $\min \max(B^T)$, fail to do so.

The two games above show an additional important observation. Since the individual minmax values of the two games are the same, the feasible payoffs described by folk theorems are the same. In other words, considering the conclusions of the folk theorem the lhs and rhs games are the same. Yet, we feel that the strategic threat possibilities are significant. Thus going straight to a folk theorem analysis, and then using any choice method which is based on the folk theorem feasible set, is bound to miss the effect of such externalities.

6.2. On dominant strategies, commitments and fairness. Consider the following 2×1 one player *up/down game*.

$$\begin{array}{|c|} \hline 0, 0 \\ \hline 1, 5 \\ \hline \end{array} \qquad \text{coco value} = (3, 3)$$

At first look, it seems strange that P2 would be willing to settle for the coco payoff of 3, rather than the payoff of 5 that he can get by cutting out communication with P1 and letting her play her dominant strategy. While this intuition is clear in purely strategic environments, where communication, threats, sidepayments and binding agreements are limited, in cooperative environments the outcome (3,3) may be more reasonable.

Example 4. *The sprinkler game.*

Two neighbors, each having to decide on watering or not a shared lawn, face the following payoff table:

	water	not
water	0,0	5,1
not	1,5	0,0
coco value = (3,3)		

In this payoff table, it seems fair and efficient that one of them will water and the other will compensate her with a transfer of 2, to obtain the coco value (3,3).

What if Player 2 says that his sprinkler is broken, and he cannot water? If he is right, we are back in the one-player up/down game above. And if the solution of the up/down game was (1,5), it would present two problems: First, there is a fairness issue, where the player who cannot water the lawn gets a higher benefit than the one who can. Second, there is the issue of incentives, where each of the two neighbors would have the incentive to destroy her sprinkler first, in order to increase her payoff. When deciding on a solution for cooperative play in games, it is desirable for each player to have the incentive to fully reveal their options. This is captured by the “monotonicity in actions” axiom discussed earlier.

Commitment is another issue that may represent a challenge to the (3,3) coco solution of the up/down game above. Wouldn't Player 2 be better off by simply walking away (making himself inaccessible for the purpose of making side payments) to obtain the payoff 5, under the assumption that Player 1's response would be to water?

In this regard, it is important to note that the implicit game with commitments and communication is substantially richer than the one summarized by the up/down payoff table. And in particular, what may be a dominant strategy in the up/down game may not be a dominant strategy in the implicit cooperative game. For example, in the larger game, Player 1 may walk away first, after leaving publicly observed irrevocable instruction to her gardener to water if and only if Player 2 gives the gardener \$2.

	pay \$2	don't pay
	3,3	0,0

By doing so, P1 creates the one row game above, in which Player 2's dominant strategy is to make the payment.

6.3. The value of information. One of the features of the coco value, is that it makes optimal use of information, and compensates players for providing it. A simple illustration of this was described in Example 1, where Player 2 had information, and no choice of actions, yet he was paid 1 for providing this information.

Valuation of information is an older subject of study in game theory, see for example in Kamien, Taumann and Zamir (1990), and the many recent references in De Myer, Lehrer and Rosenberg (2009). A natural measure of information value may be developed through the coco value, by considering how the coco value of various players changes, as you change the information of single, or even coalitions of, players. Such a measure would be fairly sophisticated, since it would take into account interactive aspects of the information: its provisions, its use, and the direct and indirect benefits that it may provide through the coco value.

6.4. Computational complexity. In the case of complete information, the coco value can be computed in polynomial time, that is time polynomial in the size of a natural representation of the game. In the case of incomplete information, where each player has at most m types, the coco value can be computed in time $(\text{size})^m$. More formally, suppose that a game is represented as follows. Let $|S|$ denote the size of finite set S . The sets of types and actions for each player are taken to be the set $A_i = \{1, 2, \dots, |A_i|\}$, and $T_i = \{1, 2, \dots, |T_i|\}$, similarly $\Omega = \{1, 2, \dots, |\Omega|\}$. The prior distribution is represented by a three-dimensional array of probabilities of any state of the world (ω, t_1, t_2) . The utility functions are similarly represented by multi-dimensional payoff arrays. As is standard, we assume that these are all rational numbers encoded as the ratios of binary integers. The size of the game, $|G|$, is simply the total number of bits required to describe the game.

Observation 3. *There is a constant $c > 0$ and an algorithm such that, given any two-player Bayesian game $G = (A, \Omega, T, \tau, u)$, the algorithm computes the coco value in time $|G|^{c|T|}$.*

It is possible that there are faster algorithms.

Proof. Computing the decomposition is algorithmically trivial – constructing the two games requires a few additions and divisions per payoff cell. Computing the value of the team game is also easy, since $u(a|t)$ is straightforward to evaluate, and the team game value is:

$$\sum_{t \in T} \Pr_\tau[t] \max_{a \in A} \frac{1}{2} (u_1(a|t) + u_2(a|t)).$$

Hence, both the decomposition and team game value can be computed in time polynomial in the size of G . For the zero-sum Bayesian game G^{eq} , one first does the standard expansion into a complete information game. Namely, one constructs the $|A_1|^{|T_1|} \times |A_2|^{|T_2|}$ bimatrix game in which each strategy (function from types to actions) in G is an action in the new game, and the payoffs of the actions in the new game are the expected values of the payoffs in G from using the respective strategies. Computing this expected value in any particular cell can be done in time polynomial in $|G|$, but one must perform this computation $|A_1|^{|T_1|} \times |A_2|^{|T_2|}$ times. Finally, once one has constructed such a game, the value of a zero-sum bimatrix game is well-known to be computed by linear programming. Theoretical algorithms for linear programming are known to take time polynomial in the size of the input (see, e.g., Grötschel *et al*, 1988). (Algorithms that work fast in practice are also well-studied.) Hence, the total run-time of the algorithm is $|A_1|^{|T_1|} \times |A_2|^{|T_2|} \text{poly}(|G|)$, which implies the observation. \square

6.5. Communication complexity. The *communication complexity* of a protocol is the number of bits communicated by the two players during the execution of a protocol. Our partnership protocol formally requires each player to announce their entire type. In a real-world game of interest, such as the professional wrestling game, a player’s type may be a complicated entity that incorporates a lot of information, including everything the player learned from the time that she was born. Hence, explicitly communicating one’s type is impractical in many situations. However, the true requirements of the coco protocol are simply that each player communicate enough information so that they can come to agreement about the best joint course of action. This can be illustrated by a simple variation on the hot dog seller example. Suppose that there are now n possible locations for the two vendors, Player 1 is perfectly informed about the weather at all n locations while player 2 remains completely uninformed. Rather than player 1 communicating her entire type, she need only tell player 2 where to go. In particular, since she has all the information, she can determine the optimal locations for the two of them to be located, and then tell player 2 where to go. Hence, instead of communicating all n weather forecasts, which would require a number of bits growing linearly in n , she sends a single integer, indicating the index of which location to go to, requiring communication complexity only $\lceil \log_2 n \rceil$ bits.

6.6. Composability. Composability of protocols has become increasingly recognized as an important topic in computer science. While cryptographic protocols were typically shown to be secure when run in isolation (such as encrypting a single message or signing a document), Canetti (2000) proposed that cryptographic protocols should be universally secure when executed concurrently in an environment with many other protocols running simultaneously. That is, a secure program for encrypting messages and a secure program for signing documents are of limited utility if the two of them are not secure when they are both used. Similarly, an analysis of a single game is arguably of less value if it doesn't apply when the game is played in a larger context. Repeated games are a classic illustration of how behavior in a composed setting differs from behavior in a one-shot setting, e.g., cooperation can be achieved in Nash equilibria of repeated prisoner's dilemma but not in the one-shot form.

Two-person zero-sum games have an appealing universal composability property. First of all, the optimal play in a repeated zero-sum game is simply optimal play in each stage. Moreover, suppose two players are to play m fixed games G_1, G_2, \dots, G_m , either in parallel or serial, or some combination of parallel and serial. This can be viewed as one large extensive-form game G , where moving in G corresponds to moving in some subset of the constituent games, and the payoffs in G are the sum of the payoffs achieved in the constituent games. The min-max value of G , regardless of the particular order in which moves in G_i 's are played, is equal to the sum of the min-max values of the constituent games. Put another way, suppose you were to play a game of tic-tac-toe, a game of chess, and a game of poker, against the same opponent. Ignoring time constraints and concerns of bounded rationality, the order in which you make your moves in the various games is irrelevant – optimal play is simply to play each game optimally.

Similarly, optimal play in composition of team games is simple. The coco value inherits the appealing composability properties of both team and zero-sum games. An easy-to-formalize version of this fact is the following. Suppose a Bayesian game is played repeatedly in a manner so that types are drawn freshly each round.

Observation 4. *For any finite two-player Bayesian game G , and any discounting parameter $\delta \in (0, 1)$, the coco value of the infinitely repeated game is equal to the value of G .*

7. CONCLUSION

As discussed above, the coco value and its implementations offer conceptual and practical frameworks to deal with issues of cooperation in strategic environments. The semi-cooperative approach is appealing, since it combines desirable elements from cooperative game theory, where binding agreements are possible; and from strategic game theory, where strategic and informational details of the environment are taken into consideration. And despite (and perhaps because of) the combined look at both types of issues, its formula and computations are relatively simple.

The justification of the coco solution was axiomatic, based on principles of fairness and efficiency; and strategic, showing that, through the use of partnership type of protocols, it can be implemented as required by the Nash program. It can even be implemented in the interim sense, after players possess private information.

While the axiomatization and the strategic implementation are standard methods of justifying solution in cooperative and strategic game theory, respectively, there are synergies in having both justifications for the same solution concept.

As a reader of the implementation protocols above may notice, one could generate similar implementation protocols that would lead to different values. In this regard, the axiomatic justification may serve as a focal point for the selection of the coco protocol over the others.

Moreover, from a descriptive point of view, it is hard to identify the protocol that players should agree upon. Indeed, as often seen in real life, the entire problem of bargaining over final payoffs

may be turned into bargaining over protocols. The fact that all protocols with outcomes that satisfy the axioms results in the coco value, may be helpful in such a discussion, or even in an argument that this discussion is not needed.

Another common difficulty in studying games of communication and commitments is the large multiplicity of equilibria. With regards to this difficulty, the axioms may serve as a focal point in selecting the "right" equilibrium: one supported by fairness, as observed in many recent experiments.

8. REFERENCES

- Aumann, R.,J., "The Core of a Cooperative Game Without Side Payments," Transactions of the American Mathematical Society, Vol 98, No. 3, pp539-552, 1961.
- Binmore, K., "Game Theory and the Social Contract vol I: Playing Fair," MIT Press, 1994.
- Biran, O. and F. Forges, "Core-stable bidding rings," University of Paris Dauphine Discussion Paper, 2009.
- Camerer, C., "Behavioral Game Theory: Experiments on Strategic Interaction," Princeton University Press, 2003.
- Carpente, L., B. Casas-Méndez, I. García-Jurado, and A. van den Nouweland, "Values for strategic games in which players cooperate," *Int. J. Game Theory* **33**: 397–419, 2005.
- Carpente, L., B. Casas-Méndez, I. García-Jurado, and A. van den Nouweland, "The Shapley valuation function for strategic games in which players cooperate," *Math. Meth. Oper. Res.* **63**:435–442, 2006.
- Chaudhuri, A., "Experiments in Economics: Playing Fair with Money", Routledge, 2008.
- Chen, X., X. Deng, and S. Teng, "Computing Nash Equilibria: Approximation and Smoothed Complexity," *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 603–612, 2006.
- Daskalakis, C., P. Goldberg, C. Papadimitriou, "The complexity of computing a Nash equilibrium," *Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC)*, 71–78, 2006.
- De Clippel, G. and Minelli, E., "Two-person bargaining with verifiable information," *Journal of Mathematical Economics* 40, 799–813, 2004.
- De Meyer, B., E. Lehrer, and D. Rosenberg, "Evaluating information in zero-sum games with incomplete information on both sides," Discussion Paper # 2009.35, CES Paris 1.
- Fehr, E. and Schmidt, K. M., "A Theory Of Fairness, Competition, and Cooperation," *Quarterly Journal of Economics*, Vol. 114, No. 3, 1999, pp 817-868.
- Forges, F., Mertens, J.F. and Vohra, R., "The Ex Ante Incentive Compatible Core in the Absence of Wealth Effects," *Econometrica*, 70, pp 1865-1892, 2002.
- Gilboa, I. and E. Zemel, "Nash and correlated equilibria: Some complexity considerations," *Games and Economic Behavior* 1:80–93, 1989.
- Grötschel, M., L. Lovász, and A. Schrijver, *Geometric Algorithms and Combinatorial Optimization*, Springer Verlag, 1988.
- Harsanyi, J., "Games with Incomplete Information Played by "Bayesian" Players, I-III. Part I. The Basic Model", *Management Science*, Vol. 14, No. 3, 1967.
- Ichiishi, T. and A. Yamazaki (2006), *Cooperative Extensions of the Bayesian Game*. Hackensack, NJ: World Scientific Publishing Co. Pte Ltd.
- Jackson, M. O. "A Crash Course in Implementation Theory," *Social Choice and Welfare*, Vol. 18, No. 4, 2001, pp 655-708.
- Jackson, M. O. and S. Wilkie, "Endogenous Games and Mechanisms: Side Payments among Players," *Review of Economic Studies* 72, Issue 2, 2005, pp 543-566.

- Kalai, E., "Proportional solutions to bargaining situations: Interpersonal utility comparisons," *Econometrica* **45**, 1623-1630, 1977.
- Kalai, E. and R. W. Rosenthal, "Arbitration of two-party disputes under ignorance," *International Journal of Game Theory*, **7**(2), 65-72, 1978.
- Kalai, E. and M. Smorodinsky, "Other solutions to Nash's bargaining problem," *Econometrica* **43**, 513-518, 1975.
- Kamien, M., Y. Tauman and S. Zamir, "On the value of information in a strategic conflict," *Games and Economic Behavior*, **2**, pp 129-153, 1990.
- Mezzetti, C., "Mechanism Design with Interdependent Valuations: Efficiency," *Econometrica* **72**, No. 5, pp 1617-1626, 2004.
- Myerson, R.B., "Two-Person Bargaining Problems with Incomplete Information," *Econometrica*, **52**(2), pp. 461-88, 1984.
- Myerson, R.B. and W.L. Thomson, "Monotonicity and independence axioms," *International Journal of Game Theory*, **9**, pp 37-49 (1980).
- Myerson, R. B., and Satterthwaite, M.A., "Efficient Mechanisms for Bilateral Trading," *Journal of Economic Theory* **29**, 265-281, 1983.
- Nash, J. F., "Equilibrium points in n-person games," *Proceedings of the National Academy of Sciences* **36**(1):48-49, 1950a.
- Nash, J. F., "The Bargaining Problem", *Econometrica* **18**: 155-162, 1950b.
- Nash, J. F., "Two-person cooperative games," *Econometrica*, **21**, 128-140, 1953.
- von Neumann, J., "Zur Theorie der Gesellschaftsspiele," *Mathematische Annalen*, **100**, 295-300, 1928.
- von Neumann, J., and O. Morgenstern, "Theory of Games and Economic Behavior," Princeton University Press, 1944.
- Rabin, M., "Incorporating Fairness into Game Theory and Economics," *The American Economic Review*, Vol. 83, No. 5, 1993, pp. 1281-1302.
- Raiffa, H., "Arbitration schemes for generalized two-person games," in *Contributions to the Theory of Games II*, Kuhn, H., and A.W. Tucker, editors, 361-387, 1953.
- Rosenthal, R. W., "An arbitration model for normal form games," *Math. Oper. Res.* **1**, 82-88, 1976.
- Roth, A. E., "Axiomatic Models of Bargaining, Lecture Notes in Economics and Mathematical Systems #170," Springer Verlag, 1979.
- Selten, R., "Bewertung Strategischer Spiele," *Zeitschrift für die gesamte Staatswissenschaft*, **116**, Band, 2. Heft, pp. 221-282, 1960.
- Selten, R., "Valuation of n-Person Games," in M. Dresher, L. S. Shapley, A. W. Tucker (eds.) *Advances in Game Theory*, Princeton, New Jersey; Princeton University Press, 52:577-626, 1964.
- Shapley, L., "A value for n-person games," *Contributions to the Theory of Games*, Vol. II, Princeton, pp. 307-317, 1953.
- Thomson, W.L. and T. Lensberg, *Axiomatic Theory of Bargaining With a Variable Population*, Cambridge University Press, 1989.
- Thomson, W., and R. Myerson, "Monotonicity and independence axioms," *International Journal of Game Theory* **9**(1):37-49, 1980.
- Wilson, R., "Game Theoretic Analysis of Trading Processes", in *Advances in Economic Theory*, ed. by T. Bewley, Cambridge University Press, 1987.