



Northwestern University

2001 Sheridan Road 580 Leverone Hall Evanston, IL 60208-2014 USA

Discussion Paper #1512
June 16, 2010

A Cooperative Value for Bayesian Games

Key words: cooperative game theory,
non-cooperative game theory,
bargaining, min-max value

JEL classification: C70, C71, C72, C78

Adam Kalai
Microsoft Research

Ehud Kalai
Northwestern University

www.kellogg.northwestern.edu/research/math

The logo for CMS-EMS, consisting of a square frame with a circle inside. The text is centered within the circle.

CMS-EMS
The Center for
Mathematical Studies
in Economics &
Management Sciences

A COOPERATIVE VALUE FOR BAYESIAN GAMES

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

ABSTRACT. Selfish, strategic players may benefit from cooperation, *provided they reach agreement*. It is therefore important to construct mechanisms that facilitate such cooperation, especially in the case of asymmetric private information. The two major issues are: (1) singling out a *fair* and *efficient* outcome among the many individually rational possibilities in a strategic game, and (2) establishing a play protocol under which *strategic* players may achieve this outcome. The paper presents a general solution for two-person Bayesian games with monetary payoffs, under a *strong revealed-payoff* assumption.

The proposed solution builds upon 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 fair and efficient cooperative play among strategic players in two-person private-information games with monetary payoffs.

The unmodified noncooperative solutions obviously miss this point, as illustrated in the sacrifice game below. The dominant strategy, in which player 1 plays “pass” and each player nets a \$0 payoff, is logical, for example, when the players are prisoners in isolated cells.

	sacrifice game	pass
pass		\$0, 0
act		-1,101

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

*Microsoft Research New England.

[†]Kellogg School of Management, Northwestern University.

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

The paper benefited from helpful comments from Geoffrey De Clippel, Françoise Forges, Dov Monderer, Phil Reny, Rakesh Vohra, Robert Wilson, and other seminar participants at Northwestern, Stanford, and Tel Aviv University. Earlier versions of this paper were presented as the keynote lecture of the 2009 Conference of Public Economic Theory in Galway, Ireland, and an invited lecture in the 2009 International Conference on Game Theory in Stony Brook, New York.

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

But in the many important economic environments where communications, side payments, and agreements are permitted, other outcomes are appealing. For example, an outcome that agrees with real life and experimental observations is for player 1 to act in exchange for a \$51 payoff from player 2, so that they each net \$50.

However, when we consider more substantial games, in which players possess many possible strategies and asymmetric private information, the determination and implementation of optimal cooperative play and associated payoff transfers are more challenging. The main purpose of this paper is to offer a solution to this problem in restricted but important classes of such games.

In particular, we focus on two-player cooperative Bayesian games with transferable utility (TU),¹ in which players can communicate and make binding agreements about actions and payoff transfers.² An important subclass consists of the zero-sum games studied by von Neumann (1928).

The solution offered 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 Bayesian 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, under a strong revealed-payoff assumption.

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 coco value is defined by

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

We refer to the first component in the decomposition as the *cooperative component*, or the *team game*. The pair of payoffs associated with this component is the *team value*, which is the highest possible pair of payoffs that the players can jointly arrange under an agreement to share their payoffs equally. We refer to the second component in the decomposition as the *competitive component*, the zero-sum game, or the *advantage game*. The minmax value of the advantage game may be thought of as a compensating zero-sum transfer from the player with the weaker strategic position to the player with the stronger one. The body of this paper presents the Bayesian (incomplete information) version of the decomposition and definition above.

¹For simplicity, we assume that payoffs are in terms of a common currency, such as dollars, and that players' 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 $n > 2$. However, in such games the minmax value is not defined, and thus the definition of the coco value may have to be significantly more complex.

With complete information, the coco value can always be implemented trivially through an agreement on any play which maximizes the payoff sum, along with a appropriate side payment. However, the implementation of first-best efficient payoffs in a Bayesian game is difficult, as illustrated by the Myerson-Satterthwaite (1983) impossibility result. The coco Bayesian decomposition enables the construction of *partnership protocols* that overcome these difficulties in restricted but 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 incentive to cooperate: to truthfully reveal all relevant information and to 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 informational 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 equal division imposed in the cooperative component.

In order to overcome the Myerson-Satterthwaite impossibility result, we restrict ourselves to games that satisfy a *strong revealed-payoff* assumption: After the play of the game each player knows the realized payoffs of both players, as well as the entire payoff function. That is, everyone knows how much everyone would have received had they played any alternative vector of actions.⁴ In Section 8, we show that in many specific games, this assumption may be weakened.

1.1. Related literature. Given the elementary nature of the questions addressed here, it is not surprising that closely related concepts have been studied earlier in both cooperative and noncooperative (or strategic) game theory. Indeed, the solution presented here may be viewed as a synthesis and generalization of several earlier works.

Starting with two-person zero-sum games, which are strictly competitive, 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 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 "rationale" of the minmax value to the class of general two person games.

Moving to games that are not zero-sum, 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,⁵ thus "everything is possible." Under the cooperative approach, one must first decide on an "artificial bridge" that transforms the strategic game into a cooperative bargaining game, to which one may apply one of several unique solutions. It

⁴See Mezzetti (2004) for an earlier use of revealed payoff assumptions in a different context.

⁵See for example Fershtman et al. (1991), Tennenholtz (2004) and Kalai et al. (2010).

seems that both the folk theorem and standard bridge resolutions miss some subtle strategic considerations, which are discussed in Section 9.1.

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.⁶

Selten (1960, 1964) presented an axiomatic characterization of a cooperative value, defined on the class of TU extensive form games of complete information. While his work preceded the definition of incomplete information, the natural extension of Selten's work would give a different value for games of incomplete information.⁷ Different axiomatic solutions of complete information strategic games, discussed later in this paper, were presented by Carpenete et al. (2005, 2006).

Little work has been done on cooperative solutions for Bayesian games. 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 et al. (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 solution, which is fair in material value and compatible with individual incentives, may serve as a focal point 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.

2. ILLUSTRATIVE EXAMPLES

We begin with a symmetric complete-information **hot-dog cart** game. 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 40 dogs, while the demand at B is for 100 dogs. 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. Each seller has to choose a location without any knowledge of the opponent's choice. Suppose they each net \$1/dog sold. The game is given below.

	A	B
A	20,20	40,100
B	100,40	50,50

⁶The TU assumption is made here, among other reasons, to circumvent the need to take positions on competing bargaining axioms.

⁷In 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.

Unsurprisingly, the coco value of this game is $(70, 70)$, while the unique Nash equilibrium achieves $(50, 50)$. Although the coco *value* is unique, this example illustrates that there may be a multiplicity of agreements which all achieve the coco value; for example, they may play (A, B) with a side payment of \$30 from P2 to P1, or the reverse. They may also achieve this value, on average, without any side payment: they may agree to flip a coin to decide between playing (A, B) and (B, A) , and in a repeated setting, they may alternate between (A, B) and (B, A) .

In the more interesting asymmetric version below, P2 nets \$2/dog sold while P1 still nets \$1/dog sold. The game is given below on the left, and the decomposition on the right.

$$\begin{array}{c} \text{A} \\ \text{B} \end{array} \begin{array}{|c|c|} \hline \text{A} & \text{B} \\ \hline 20,40 & 40,200 \\ \hline 100,80 & 50,100 \\ \hline \end{array} = \begin{array}{c} \text{A} \\ \text{B} \end{array} \begin{array}{|c|c|} \hline \text{A} & \text{B} \\ \hline 30,30 & \mathbf{120,120} \\ \hline 90,90 & 75,75 \\ \hline \end{array} + \begin{array}{c} \text{A} \\ \text{B} \end{array} \begin{array}{|c|c|} \hline \text{A} & \text{B} \\ \hline -10,10 & -80,80 \\ \hline 10,-10 & \mathbf{-25,25} \\ \hline \end{array}$$

The team game has value $(120, 120)$, while the zero-sum game has value $(-25, 25)$; hence the coco value is $(95, 145)$. The coco value achieves the maximum total of 240, which the players can obtain by playing (A, B) followed by a \$55 side payment from player 2 to player 1 $((40, 200) + (55, -55) = (95, 145))$.

Whether real-world players would reach agreement at all, and whether such agreements would involve transfers of approximately \$55, are left for further study. Nonetheless, the coco value is a focal point that may aid in reaching such agreements. The coco decomposition, disentangling the cooperative from the competitive incentives, offers a rationale for this focal point.

The main body of this paper extends the above analysis to games with asymmetric private information, as illustrated next.

2.1. Incomplete-information hot-dog cart game. Now, suppose the demand at B depends on the weather: if sunny, which has probability $1/2$, it is for 200 dogs; and if cloudy, which has probability $1/2$, there are no customers at the beach. Furthermore, suppose that player 1 is perfectly informed, a priori, of the weather and player 2 has no information.

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

<p>Sunny prob 1/2</p> <table border="1" style="margin-left: auto; margin-right: auto;"> <tr> <td></td> <td style="text-align: center;">A</td> <td style="text-align: center;">B</td> </tr> <tr> <td style="text-align: center;">A</td> <td style="text-align: center;">20,40</td> <td style="text-align: center;">40,400</td> </tr> <tr> <td style="text-align: center;">B</td> <td style="text-align: center;">200,80</td> <td style="text-align: center;">100,200</td> </tr> </table>		A	B	A	20,40	40,400	B	200,80	100,200	<p>Cloudy prob 1/2</p> <table border="1" style="margin-left: auto; margin-right: auto;"> <tr> <td></td> <td style="text-align: center;">A</td> <td style="text-align: center;">B</td> </tr> <tr> <td style="text-align: center;">A</td> <td style="text-align: center;">20,40</td> <td style="text-align: center;">40,0</td> </tr> <tr> <td style="text-align: center;">B</td> <td style="text-align: center;">0,80</td> <td style="text-align: center;">0,0</td> </tr> </table>		A	B	A	20,40	40,0	B	0,80	0,0
	A	B																	
A	20,40	40,400																	
B	200,80	100,200																	
	A	B																	
A	20,40	40,0																	
B	0,80	0,0																	

The expected payoffs, obtained under three different computational schemes, are summarized in the table below:⁸

	P1	P2	Total
Noncooperative, Bayesian equilibrium	70	100	\$170
Purely cooperative play	20	240	260
The coco value: cooperate & transfer	115	145	260

⁸In contrast to this example, there are games where the coco payoff of an individual is lower than her equilibrium payoff, as we discuss in Section ??.

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; having no such knowledge, P2 simply chooses B . The expected payoffs in the table are computed to be: $(70, 100) = 0.5(100, 200) + 0.5(40, 0)$.

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 , but when it is cloudy, they do the opposite. The expected payoffs in the table are computed to be: $(20, 240) = 0.5(40, 400) + 0.5(0, 80)$.

Clearly, the players would like to obtain the cooperative total of \$260 rather than the total noncooperative total of \$170. But the cooperative solution calls for substantial sacrifices on P1's part: he must always disclose his forecast truthfully and then choose the inferior location.

A natural resolution is to amend the efficient solution with payoff transfers. In our solution to the example above, on sunny days, when P2 goes to B and P1 to A , P2 pays P1 \$130 out of her \$400 payoff. And on cloudy days, when he goes to B and she goes to A , P2 pays him \$60 out of her \$80 payoff. The expected payoffs in the table are computed to be: $(115, 145) = 0.5(170, 270) + 0.5(60, 20)$.

In Section 9.3, we compute the value of information for the two players.

2.2. Overview. The rest of the paper is organized as follows. Section 4 presents the definition of the coco value and the formula for computing it, as done above. Section 5 presents axioms of fairness and efficiency that justify this value. Section 6 studies implementation. It presents partnership protocols in which the players have the incentive to disclose their private information and to act optimally in order to bring about the coco payoffs. Section 7 is devoted to a joint venture example, relating the coco value approach to the Myerson-Satterthwaite bargaining model. Additional elaboration and connections to earlier literature are presented in sections 8 and 9.

3. PRELIMINARIES

Unless otherwise specified, we consider games with a fixed set of two players, $N = \{1, 2\}$. A Bayesian game G is defined by: $G = (A = \times_{i \in N} A_i, T = \times_{i \in N} T_i, U = \times_{i \in N} U_i, \mu)$ where for each player i , A_i denotes the set of *actions*, T_i denotes the set of *types*, and $U_i \subseteq \mathbf{R}^A$ denotes the set of payoff functions (utilities), $u_i : A \rightarrow \mathbf{R}$. All these sets are assumed to be finite and μ is the *prior* probability distribution over $T \times U$. To increase readability, we sometimes write $(t, u) \sim \mu$ to indicate that the pair (t, u) is drawn from the distribution μ . Notice that in addition to *information types*, this formulation also allows for *payoff types* (for every t_i have $\mu(u_i|t_i) = 1$ for some u_i), and for types that combine the two.

As is standard, 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, $(t, u) \in T \times U$, is drawn from μ . Each player i then observes her own type t_i , on the basis of which she chooses (simultaneously with her opponent) an action $a_i \in A_i$. The payoff to player i is $u_i(a)$, where $a = (a_1, a_2)$ is the selected *action profile*.

As is standard, a *mixed action* for player i is a probability distribution $\alpha_i \in \Delta(A_i)$ over the set of actions. We also extend the domain of payoff functions, u_i , to mixed actions by the use of expected values. A (pure) *strategy* for player i , $s_i : T_i \rightarrow A_i$, is a function that specifies the action that player i would choose if her type were

$t_i \in T_i$; 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 her knowledge of own type.

It is also common to refer to t_i as the *private information* of player i . A game is *zero-sum* if, with probability 1, $u_1 = -u_2$, and it is a *team game* if $u_1 = u_2$, with probability 1. Finally, we use the standard convention that a_{-i} and u_{-i} represent, respectively, the actions and payoffs of player i 's opponent. A *coordinated strategy* is a function $c : T \rightarrow A$ from type profiles to action profiles.

As mentioned, we assume that the players have additive 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.

3.1. Revealed-payoff assumptions. The *revealed-payoff assumption* requires that, after play of the game, the payoff vector $u(a)$ is revealed to all players. The *strong revealed-payoff assumption* requires that, after play of the game, the entire payoff function u is revealed to all players. That is, each player then knows how much each player would have received had they played any other action profile.⁹

One example that satisfies the strong revealed-payoff assumption is the hot-dog sellers game. Provided that payoffs resulting from any location choices depend only on the weather, and that the weather is observed by all players after the play, one can compute the payoffs for any hypothetical location choices. Notice that strong revealed-payoff assumption continues to hold in all games in which payoffs depend entirely on a *state of nature* which is observable after the play of the game.

We should point out, however, that the payoff-revelation assumptions do not require that the types of the players be revealed. As an example of where this distinction is meaningful, consider again a game in which the payoffs depend on the weather, and the players' types consist of imperfect individual weather forecasts. If the weather is observed by all players after the play of the game, then the revealed-payoff assumption holds even if the forecasts are not revealed.

But nevertheless, there are important games in which the revealed payoff assumption does not hold. One example is an auction in which the payoff of the winner depends on his private value of the item, which is not verifiable.

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

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

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

⁹Mezzetti (2004) used similar revealed-payoff assumptions in a study of auctions.

4.1. 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 :

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

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 $z^* = \min \max(Z, -Z)$ is the classical von Neumann value for 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, which are 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 the following conditions hold.*

- (1) *For any zero sum game, $\kappa(A, -A) = \min \max \text{value}(A, -A)$.*
- (2) *For any team game, $\kappa(A, A) = \text{team value}(A, A) = (\max_{ij} a_{ij}, \max_{ij} a_{ij})$.*
- (3) *The coco value is feasible and does not require the use of mixed strategies: there is always a simple agreement consisting of a pair of (pure) actions and a monetary transfer, which yields net payoffs equal to the coco value.*

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

Definition 1. *For any $u : A \rightarrow \mathbf{R}^N$, define $u^{eq}, u^{ad} : A \rightarrow \mathbf{R}^N$ as follows:*

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

Although $u^{eq}(a) \in \mathbf{R}^2$, we sometimes use $u^{eq}(a)$ to denote the single equal payoff that it allocates to the players, and thus write $u^{eq}(a) = u_1^{eq}(a) = u_2^{eq}(a)$.

The cooperative-competitive decomposition presented above for complete information extends naturally to incomplete information:

$$u = u^{eq} + u^{ad}.$$

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 define the following notions.

Definition 2. *For $G = (A, T, U, \mu)$, the team optimum of G is defined by:*

$$\text{team-opt}(G) = \max_{c: T \rightarrow A} \mathbf{E} [u_1(c(t)) + u_2(c(t))].$$

A coordinated (pure) strategy $c : T \rightarrow A$ is called optimal if $\mathbf{E} [u_1(c(t)) + u_2(c(t))] = \text{team-opt}(G)$.

In words, an optimal coordinated strategy is a rule c that the players may use to select, for every pair of realized types t , a pair of actions $c(t)$ that maximizes the sum of their expected payoffs in G . The team optimum is the maximal sum of expected payoffs that may be generated by such a rule. Notice that this definition assumes that they truthfully share all their information and then coordinate their actions.

Definition 3. *The relative advantage of player i is defined to be her minmax value in the zero-sum game $G^{ad} = (A, T, V, \mu^{ad})$, where V consists of all payoff functions $v = (v_1, v_2)$ with each $v_i = u_i^{ad}$ for some payoff function u_i of G , and $\mu^{ad}(t, v) = \mu(\{(t, u) : v_i = u_i^{ad}, \text{ for } i = 1, 2\})$.*

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 $\text{minmax}_i(G^{ad})$. We refer to this value as the player's *competitive advantage*, *relative advantage*, or just *advantage*.

Definition 4. *The coco value of G to player i , denoted by $\kappa_i(G)$, is defined by,*

$$\kappa_i(G) = \frac{1}{2} \text{team-opt}(G) + \text{minmax}_i(G^{ad}).$$

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.¹⁰ 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 case, one may think of the coco value of a private-information game as the sum $\kappa(G) = \text{team-value}(G^{\text{eq}}) + \text{minmax-value}(G^{ad})$.

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

5. AXIOMATIC CHARACTERIZATION OF THE COCO VALUE

We define a *value* to be a function from the set of all finite two-person games to \mathbf{R}^2 , i.e., $v(G) \in \mathbf{R}^2$ where $v_i(G)$ is the value to player i . To better understand the axiomatization, the reader may benefit from first examining the axioms in the case of complete-information games, given in Figure 1. These five simple axioms can be shown to be uniquely satisfied by the coco value. We omit the analysis, which is a direct simplification of that of Theorem 1 below.¹¹

Moving to the general case of games of incomplete information, the coco value satisfies a large number of appealing properties inherited from the minmax value of zero-sum games. Some of these are discussed in the concluding sections of this paper. But a justification of a value is more convincing if it is unique even in satisfying a

¹⁰ $G^{\text{eq}} = (A, \hat{T}_1 \times \hat{T}_2, V, \hat{\mu})$ defined by $\hat{T}_1 = \hat{T}_2 = T_1 \times T_2$, $V = \{v : v = u^{\text{eq}} \text{ for some } u \in U\}$ and $\hat{\mu}((t, t), v) = \mu(\{(t, u) : u^{\text{eq}} = v\})$.

¹¹As mentioned earlier, Selten (1960, 1964) gave a more involved axiomatization of the analogous value for extensive-form games with complete information. We are grateful to Moshe Tennenholtz and Dov Monderer for pointing us to Selten's work.

- 1. Pareto efficiency.** Players maximize *total* payoff: $v_1(X, Y) + v_2(X, Y) = \max_{ij} x_{ij} + y_{ij}$.
- 2. Shift invariance.** The shifting of payoffs by the same constants in every cell leads to a corresponding shift in the value: $v(X + c_1, Y + c_2) = v(X, Y) + (c_1, c_2)$.
- 3. Monotonicity in actions.** Removing an action cannot increase a player's value, e.g., for player 1, $v_1(X', Y') \leq v_1(X, Y)$ where (X', Y') are formed by removing a row of (X, Y) .
- 4. Payoff dominance.** If $x_{ij} > y_{ij}$ for all i, j , then $v_1(X, Y) \geq v_2(X, Y)$. Similarly for player 2.
- 5. Invariance to redundant strategies.** An action which is equivalent, in expectation, to a mixed action, may be removed without changing the value: if the i th rows of X and Y equal a convex combination α of other rows, with $\alpha_i = 0$, then the i th rows may be removed without changing the value. Similarly for player 2.

FIGURE 1. The coco value is unique in satisfying the above set of axioms for complete information games $X, Y \in \mathbf{R}^{m \times n}$.

small number of weak reasonable requirements. The following properties, or axioms, are sufficient for our axiomatization theorem.

- (1) **Pareto efficiency.** Players achieve the maximum (first-best) total expected payoff possible, $v_1(G) + v_2(G) = \max_{c: T \rightarrow A} \mathbf{E}[u_1(c(t)) + u_2(c(t))]$, with *shared information*, i.e., $\text{team-opt}(G)$.
- (2) **Shift invariance.** The shifting of payoffs by constants in every cell leads to a corresponding shift in the value. Formally, fix $c = (c_1, c_2) \in \mathbf{R}^2$. For any u , let $u'(a) = u(a) + c$. Then $v(G') = v(G) + c$ where $G' = (A, T, U', \mu')$, with $U' = \{u' : u \in U\}$ and $\mu'(t, u') = \mu(\{(t, u)\})$.
- (3) **Monotonicity in actions.** Removing an action of a player cannot increase her value. Formally, let $A'_1 \subseteq A_1$ and u' be the restriction of any u to $A'_1 \times A_2$. Then $v_1(G') \leq v_1(G)$ where $G' = (A'_1 \times A_2, T, U', \mu')$, in which U' consists of the restricted payoff functions from G , and μ' is the induced distribution over (t, u') (i.e., $\mu'(t, u') = \mu(\{(t, u) : u|_{A_1} = u'\})$). Similarly for Player 2.
- (4) **Payoff dominance.** If, under 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 $\min_{c: T \rightarrow A} \mathbf{E}[u_1(c(t)) - u_2(c(t))] > 0$, then $v_1(G) \geq v_2(G)$. Similarly for player 2.
- (5) **Invariance to redundant strategies.** Let $a_1 \in A_1$ and $A'_1 = A_1 \setminus \{a_1\}$. We say a_1 is *redundant* (in expectation) if there exists $\sigma_1 : T_1 \rightarrow \Delta(A'_1)$ with the property that for every $t \in T$, and every $a_2 \in A_2$ $\mathbf{E}_\mu[u(a_1, a_2) | t] = \mathbf{E}_\mu[u(\sigma_1(t_1), a_2) | t]$. Then removing such a redundant action a_1 does not change the value of the game for either player. Similarly for any redundant action of player 2.
- (6) **Monotonicity in information.** Giving player i strictly less information cannot increase her value. In particular, $v_1(G') \leq v_1(G)$, where G' is defined by replacing player 1's information t_1 by some function $f(t_1)$. Formally, take an arbitrary function $f : T_1 \rightarrow T_1$ and $G' = (A, T, U, \mu')$ with $\mu'((t'_i, t_{-i}), u) = \mu(\{(t_i, t_{-i}), u) : f(t_i) = t'_i\})$. Similarly for player 2.

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

Before turning to the proof of the theorem, we first offer some intuition through a sketch of the proof for complete information games. Specifically, for any value v that satisfies the Axioms of Figure 1, $v(G) = \kappa(G)$ for any complete information game G . Shift invariance implies that it suffices to consider the special case of

games G with $\kappa(G) = (0, 0)$. Moreover, it suffices to show that $v_1(G) \geq 0$, because a similar argument would show that $v_2(G) \geq 0$, and Pareto efficiency would imply that $v(G) = (0, 0)$ to complete the proof of this special case. Now, when $\kappa(G) = (0, 0)$, the minmax value of G^{ad} must also be zero. Let σ_1^* be any minmax strategy of player 1 in G^{ad} . By the coco decomposition, in the game G this strategy guarantees player 1 an expected payoff at least as large as that of player 2. Now consider the game H in which player 1 is forced to play σ_1^* (a new pure action is created corresponding to σ_1^* and all other actions are deleted). By axioms 3 and 5, $v_1(G) \geq v_1(H)$. By payoff dominance (see the proof of Theorem 1 for how to address the weak vs. strong inequality), $v_1(H) \geq v_2(H)$. If the team optimum of H were the same as G , i.e., $v_1(H) + v_2(H) = 0$, then these three facts would imply that $v_1(G) \geq 0$, and we would be done.

However, forcing player 1 to play σ_1 may decrease the team optimum. To overcome this difficulty in this complete-information case, before we force player 1 to play σ_1^* , we augment the game G as follows: we add to player 2 a new simple action that yields the constant payoffs $(0, 0)$, no matter what player 1 plays. This does not change the coco value, and axioms 1 and 3 imply that this new strategy cannot increase player 1's value (in particular, player 2 is no worse off while the team-opt remains 0). Furthermore, when player 1 is now forced to play σ_1 , player 2's new action guarantees that the team optimum remains 0, and hence the argument in the previous paragraph goes through.

In the case of incomplete information, the approach of the proof above fails for two reasons. First, the addition of a constant $(0, 0)$ action for player 2 could very well change the team optimum and the value of the advantage game, because this action may be taken based upon the player's information. Second, in order to apply the payoff-dominance axiom, we remove all of player 1's information, which might decrease the team optimum. We now show how to address these subtleties.

Lemma 1. *Let G be a finite two-person Bayesian game such that $\kappa(G) = (0, 0)$. Then axioms 1-6 above imply that $v_1(G) \geq 0$.*

Proof. We will construct a sequence of games and argue that $v_1(G) \geq v_1(G') \geq v_1(G'') \geq v_1(H) \geq 0$.

To construct G' , we add a new action b_2 ($\notin A_2$) to player 2's set of actions, so that the sets of actions of G' are $A'_1 = A_1$ and $A'_2 = A_2 \cup \{b_2\}$. Next we define the possible payoff functions U' of G' . We fix any action $a_2^* \in A_2$ for player 2, we fix some team-optimal coordinated strategy $c : T \rightarrow A$ in G (see Definition 2), and we define the gain from cooperation at (a_1, a_2^*) and t to be $g = u_1(c(t)) + u_2(c(t)) - u_1(a_1, a_2^*) - u_2(a_1, a_2^*)$. Every payoff function of G is extended in up to $|T|$ payoff functions in G' so that when player 2 selects the new action b_2 , their payoffs are those of G at (a_1, a_2^*) plus the gain g divided equally between the two players.

Formally, for any $t \in T$ and $u : A \rightarrow \mathbf{R}^2$, define $f^{tu} : A' \rightarrow \mathbf{R}^2$ by

$$f^{tu}(a) = \begin{cases} u(a) & \text{if } a_2 \neq b_2, \\ u(a_1, a_2^*) + \left(\frac{g}{2}, \frac{g}{2}\right) & \text{if } a_2 = b_2, \end{cases}$$

with $g = u_1(c(t)) + u_2(c(t)) - u_1(a_1, a_2^*) - u_2(a_1, a_2^*)$.

Now the prior probability distribution of G' is the one induced by μ , i.e., $\mu'(t, u') = \mu(\{(t, u) : f^{tu} = u'\})$.

It is easy to see that the team optimum of G' is still zero, because the total achieved by any coordinated strategy in G' can also be achieved in G , so $v_1(G') + v_2(G') = 0$. By monotonicity in actions, $v_2(G') \geq v_2(G)$. Hence, $v_1(G') \leq v_1(G)$.

Next, because $\kappa(G) = (0, 0)$, the value of the advantage game G^{ad} must be $(0, 0)$. Hence, there must exist a mixed strategy for player 1, σ_1^* , which guarantees player 1 at least as much as player 2, in expectation, i.e., $\mathbf{E}_\mu[u_1(\sigma_1^*(t), \sigma_2(t)) - u_2(\sigma_1^*(t), \sigma_2(t))] \geq 0$ for any σ_2 . In particular, fix any such σ_1^* which is a minmax optimal strategy for player 1 in G^{ad} . Note that σ_1^* also guarantees player 1 at least as much as player 2 in G' , because b_2 is equivalent to a_2^* in terms of the difference in the players' payoffs.

Now, using σ_1^* defined above, we define a new action $b_1 \notin A_1$; we then consider the game G'' , obtained from G by restricting the actions of player 1 to be $A_1'' = \{b_1\}$, with payoffs $u''(b_1, a_2) = u'(\sigma_1^*, a_2)$. Hence, in G'' player 1 must play like σ_1^* in G' (in expectation). By monotonicity in actions, this means that $v_1(G'') \leq v_1(G')$. (To see this formally, one must first consider the game with actions $(A_1 \cup \{b_1\}) \times A_2'$, which has the same value as G' because b_1 is redundant by axiom 5; then remove all remaining actions for player 1.) Finally, $\text{team-opt}(G'') = 0$, since when player 2 plays b_2 , they achieve the same expected total as when they play c in G .

Now, by design, b_1 guarantees player 1 at least as much as player 2, in expectation. However, to apply payoff dominance, we must argue that, *even if the players coordinate*, player 1 gets *strictly more* than player 2, in expectation. Even though player 1 has only one action in G'' , he may have information that can help player 2 achieve an advantage.

To address this coordination problem, we remove all information from player 1. In particular, fix any $t_1^* \in T_1$ and define the game H by changing only the set of types of G'' so as to obtain $T_1^H = \{t_1^*\}$ with $\mu^H((t_1^*, t_2), u'') = \mu''(\{(t_1, t_2), u''\} : t_1 \in T_1)$.

By axiom 6, $v_1(H) \leq v_1(G'')$. Also, the team optimal of H remains zero, because player 2 still has the option of playing the fixed action b_2 . Finally, notice that player 1 is guaranteed an expected amount at least as large as that of player 2, due to our choice of b_1 . Coordination is impossible since player 1 has only one action and one possible type.

We are almost ready to apply Axiom 4. The remaining issue is that we have a weak inequality rather than a strong one. To complete the proof, imagine translating the payoff of player 1 up by any constant $\epsilon > 0$. By Axiom 2, this would only shift his value up by ϵ . However, once his payoff has been shifted, Axiom 4 does apply, in which case player 1's value is at least as large as that of player 2. Hence, $v_1(H) + \epsilon \geq v_2(H)$. Since this holds for every $\epsilon > 0$, it follows that $v_1(H) \geq v_2(H)$. Combining this with $v_1(H) + v_2(H) = 0$ implies that $v_1(H) \geq 0$, and we have already argued that $v_1(G) \geq v_1(G') \geq v_1(G'') \geq v_1(H)$. ■

We now prove Theorem 1.

Proof of Theorem 1. 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 value maximizes the expected sum of payoffs. 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 to a shift of $w_1 + w_2$ 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, and hence also for the coco value. Similarly, removing a redundant

action for i in G corresponds to removing the redundant action in the zero-sum and team games, which does not change their value. Removing a redundant state also does not change the value of a team game or a zero-sum game.

The proof of the converse, namely that the only value v that satisfies the axioms is $v(G) = \kappa(G)$, follows easily from Lemma 1. Specifically, translate the payoffs of G by $-\kappa(G)$ to get game G' where $\kappa(G') = (0, 0)$. Lemma 1 states that $v_1(G') \geq 0$. Since the axioms are symmetric, the same reasoning implies $v_2(G') \geq 0$. Pareto efficiency then implies that $v(G') = (0, 0)$. Finally, axiom 2 implies that $v(G) = v(G') + \kappa(G) = \kappa(G)$. ■

6. NONCOOPERATIVE IMPLEMENTATION OF THE COCO VALUE

Solutions to strictly-cooperative and to strictly-competitive games deal with private information easily, but in opposite ways. In Bayesian team games (where the players' payoffs are identical), the obvious incentive is to fully disclose all private information, enabling the players to choose the mutually best pair of actions. Conversely, in every Bayesian zero-sum game (where one player's gain is the other's loss), the obvious incentive is not to disclose any private information, keeping the opponent from gaining any advantage. It follows that through the decomposition of a Bayesian game into the sum of a team game and a zero-sum game, the coco value can also deal with private information easily, provided that the incentives in the play of each component are independent of the play of the other component. This idea is exploited in this section, where we study the implementability of the coco value.

Would unobligated strategic players choose to use the coco value? In the case of complete-information games, it is simple to implement the coco value through the use of individual commitments and binding agreements. But in the case of incomplete information, implementation is more difficult due to the need to share and to make a coordinated use of private information.

We use the term *protocol* to describe a two-person procedure that involves communication and simple commitments, without the use of joint randomization devices (other than the original choice of types and payoff functions in the game being implemented). We say that a protocol *implements* the coco value if it admits a Bayesian Nash equilibrium with expected payoffs that match the coco value.

As readers familiar with the implementation literature know, achieving first-best efficiency in general Bayesian games may be impossible, even if we do not insist on simultaneously achieving the other properties (such as fairness) of the coco value. We overcome such difficulties by restricting ourselves to applications that satisfy the revealed-payoff assumptions discussed in Section 3.1. Under the weaker assumption – that the realized payoff vector $u(a)$ is revealed after the play of the game – *ex-ante* Bayesian implementation of the coco value is possible. And under the stronger assumption – that the entire realized payoff function u is revealed after the play of the game – an *interim* Bayesian implementation of the coco value is possible. The latter implementation is more realistic because it is interim; thus, the decision of whether to adopt the coco value takes place after the players know their types, and its protocol does not require knowledge of the prior probability of the game being implemented. Further discussion about the merits of and the need for of these assumptions is deferred to later in this section.

Our protocol mirrors some of the methods used in the formation of real-life 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.¹² 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, if there are differences in information, resources, and opportunity costs – the partners may agree up-front to make a compensating payoff transfer. If the size of the transfer is independent of their performance in the joint venture (for example, they commit to the size of the transfer before the actual play), then the partnership should still be able to achieve first-best efficiency in an incentive compatible manner.

The observation above can be used to achieve efficiency in different ways. In subsections 6.1 and 6.2 we use it to construct specific incentive-compatible protocols that implement the coco value in some restricted but important classes of games.

Throughout the remainder of this section, $G = (A, T, U, \mu)$ is assumed to be an arbitrary, fixed two-player finite Bayesian game as discussed above.

6.1. Ex-ante implementation. We make two assumptions in this subsection: (1) both realized payoffs are revealed to both players after the game is played, and (2) 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 maximize the total payoff. However, to make up for the imposed equal division of the payoffs when the game is not symmetric, a second side payment is made from the weaker player to the stronger one. When the two side payments are combined, the coco value is achieved at equilibrium. A direct consequence of this protocol is that the coco value is ex-ante individually rational.

Ex-ante partnership protocols for an arbitrary, finite, two-player Bayesian game $G = (A, T, U, \mu)$.

Fix any *optimal coordinated strategy* $c : T \rightarrow A$.¹³

- (1) Players simultaneously choose whether to commit to participate or not.
 - If either one 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, (t_1, t_2, u) , is drawn by the prior distribution μ , and each player i is informed of her *realized type* t_i .
 - (3) Players $i = 1, 2$ simultaneously declare their supposed types $\tilde{t}_i \in T_i$.
 - (4) The players are committed to play the pair of actions $a = c(\tilde{t})$, after which the pair of payoffs $u(a)$ is revealed.
 - (5) A side payment is made so that the net payoff to player i is $u^{\text{eq}}(a) + \text{val}_i(G^{\text{ad}})$. In other words, she is paid one-half of the total payoffs obtained through the actual play in stage 4, plus her minmax value (positive or negative) of the advantage component-game of G , computed without knowledge of the types.
-

¹²The use of such monotonicity conditions is common in cooperative game theory; see, for example, Kalai (1977) and Myerson and Thomson (1980).

¹³Recall that for every other coordinated strategy, $c' : T \rightarrow A$, $\mathbf{E}_{t,u} [u^{\text{eq}}(c(t))] \geq \mathbf{E}_{t,u} [u^{\text{eq}}(c'(t))]$.

Theorem 2. *The coco value $\kappa(G)$ of any finite two-player Bayesian game $G = (A, T, U, \mu)$ is the expected payoff vector of a Nash equilibrium in any ex-ante partnership protocol of the game.*

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} .
- 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 $\frac{u_1(a)+u_2(a)}{2}$ 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} . We can nonetheless compute her expected payoffs via the coco decomposition. In particular, player i 's expected payoff is the sum of the expected payoffs in G^{ad} and G^{eq} . The expected payoff in G^{ad} is at most the (minmax) value of G^{ad} , and the expected payoff of σ in G^{eq} is at most the (team) value of G^{eq} ; hence, her total is at most the coco value for i . ■

In addition to its direct implementation message, the theorem above serves as an easy way to establish the following.

Corollary 1. *The coco value is individually rational (ex-ante).*

Proof. Notice that a player may decline to participate, and choose to use her G minmax strategy in the ex-ante protocol above, guaranteeing herself her minmax value of G as the payoff in the protocol. Thus, the minmax values of the protocol are at least as high as the minmax values of G . Moreover, being equilibrium payoffs of the protocol game, the coco payoffs must be at least as high as the minmax values of the protocol. Thus, the coco payoffs are at least as high as the minmax payoffs of G . ■

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 own types before committing to play. If either player knows some information about their own 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) advocates the use of mechanisms with rules and payoff functions that do not depend on the prior probability distribution of the game being implemented. The mechanism above violates the Wilson doctrine in two respects. First, in order to compute the optimal coordinated strategy c used in the definition of the protocol, one must know the prior distribution. Second, in order to compute the value of G^{ad} used in allocating the protocol's payoffs, one must know the prior as well.

In the next section, we show how these deficiencies may be overcome by imposing a further restriction on the environment.

6.2. Interim implementation. In this section we make the strong revealed-payoff assumption: The realized payoff function $u : A \rightarrow \mathbf{R}^2$ (equivalently, the *state of nature* if it is incorporated into the model) becomes known after the play of the game, and the players can compute what the realized payoffs $u(a)$ would have been for every chosen pair of actions a . In Section 8, we give examples where this assumption does not hold but the coco value can be implemented nonetheless.

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 weather forecasts (their types) before deciding on a location. The payoffs in this example depend on the weather and not on the forecasts, 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 state of nature becomes known, even if the types (the weather forecasts) remain unknown.

Under the above assumption, one can design effective interim protocols to implement the coco value.

Interim partnership protocol for an arbitrary, finite, two-player Bayesian game $G = (A, T, U, \mu)$.

- (1) A triple (t_1, t_2, u) is drawn by the prior distribution μ , 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 *noncooperative 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 *noncooperative action*.

The **YES/NO** declarations are revealed to both players and then:

- If either player declares **NO**, then the noncooperative pair of actions $(\tilde{a}_1, \tilde{a}_2)$ selected above is played, the game stops, and the players collect their respective G -payoffs, $u(\tilde{a})$.
 - 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 “*cooperative*” actions $a_i \in A_i$ and play G using a . Both $u(a) \in \mathbf{R}^2$ and the realized payoff function $u : A \rightarrow \mathbf{R}^2$ are then revealed.
 - (4) Based upon the cooperative actions a_i from stage 3 and the noncooperative actions \tilde{a}_i from stage 2 (the envelopes are now opened), side payments are made so that the net payoff to each player i is $u_i^{\text{eq}}(a) + u_i^{\text{ad}}(\tilde{a})$.
-

Note that in the protocol above the participation decision is *voluntary*, and that by choosing not to participate, each player can force the play of the unmodified game. However, this also means that any Nash equilibrium of G can be converted to a nonparticipatory equilibrium of the interim partnership protocol. (In some games, such as Prisoner’s Dilemma, participation is a dominant strategy.) While it is possible to employ refinements and “implement away” these equilibria,¹⁴ we feel that it is also reasonable to model the possibility that players may choose not to

¹⁴One difficulty is evident even in a pure coordination game, such as

1,1	0,0
0,0	2,2

, where there is a (1,1) equilibrium. However, a *team game refinement*, which is natural among cooperative players,

participate and to play noncooperatively. Before stating our theorem, we point out two practical considerations regarding the protocol.

- (1) In stage 3 above, one might allow extra communication, in the form of cheap talk, to aid the players in selecting the same coordinated optimal strategy c . However, since Nash equilibrium allows for coordinated selection when multiple equilibria are available, this is not necessary for the formal theorem below. Similarly, for many games one might consider protocols with lower *communication complexity* (see, e.g., Kushilevitz and Nisan, 1996), which is defined as the number of bits transmitted in a binary communication. In many games, an optimal c may be computed using significantly less communication than when players reveal all of their private information.
- (2) The definition of the protocol above is in no way dependent on the prior. Moreover, the implementing strategies rely on very 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. Hence, the resulting solution inherits some appealing stability and polynomial-time computability properties from these more robust solution concepts.¹⁵ These issues are discussed in the concluding section.

In the equilibrium discussed in this theorem, players choose to truthfully share information and to optimally coordinate, with threats defined through the relative-advantage game.

Definition 5.

- (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 strategy (see Definition 2) $c : T \rightarrow A$ 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 a 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 ($\tilde{t}_i \neq t_i$, which is a probability zero event), then she selects a_i which maximizes $\mathbf{E}_u [u^{eq}(a_i, c_{-i}(\tilde{t})) | (t_i, \tilde{t}_{-i})]$.¹⁶

A c -coordinated strategy is clearly participatory, honest, and ex-post efficient; and also enjoys the additional properties specified in the following theorem:

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

- (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 , $\kappa(G)$.

could rule out such equilibria. A suitable implementation may then have *all* equilibria yielding the coco value in expectation.

¹⁵We thank Robert Wilson for pointing out the solution's stability.

¹⁶It is necessary to specify how to act under such zero-probability events in order to argue, as we do below, that we have a *sequential* equilibrium.

- (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.
- (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 a Nash equilibrium of the partnership protocol for both players to declare NO and select \tilde{a}_i according to σ_i .

Proof of Theorem 3. For part 1 assume, for example, that player 2 uses her c -coordinated strategy. Notice first that by definition, in every one of player 1's stage-3 information sets, player 1 acts optimally, since his opponent tells the truth and follows the c -optimal selection.

So, for part (1), it remains to be shown that player 1 acts optimally at his first information set, namely, in stage 2. Suppose that instead of the above, he chooses not to participate, and to play \tilde{b}_1 . By the decomposition, his expected payoff, conditioned on t , is:

$$\mathbf{E}_u [u_1^{\text{ad}}(\tilde{b}_1, \tilde{a}_2) + u_1^{\text{eq}}(\tilde{b}_1, \tilde{a}_2) \mid t].$$

But by switching his strategy to a c -coordinated one, player 1's payoff may be written as:

$$\mathbf{E}_u [u_1^{\text{ad}}(\tilde{a}_1, \tilde{a}_2) + u_1^{\text{eq}}(c(t)) \mid t],$$

where \tilde{a}_1 is chosen by his advantage-game minmax strategy against player 2's \tilde{a}_2 , chosen according to her minmax strategy. Hence, it is easy to see that the switch to the c -coordinated strategy can only increase both terms in the above two expectations. We can also easily see, using the same decomposition argument above, that under participatory strategies, player 1 cannot obtain a higher payoff than by following any other c -coordinated strategy. Thus, part (1) of the theorem holds.

For part (2) the decomposition implies 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 . ■

6.3. *Interim individual rationality and conditional values.* Part (1) of Theorem 3 offers a positive equilibrium answer to the question of whether the unobligated players would choose to participate after they know their types. For any pair of privately known types, if one player participates, it is a best response for the opponent to participate. For complete-information games, this means in particular that the coco value is individually rational; as we already discussed, the coco value is also individually rational ex-ante, before the players know their types.

One may ask whether the coco value is also individually rational *interim*, that is, after the players have observed their types. A related question is: what is the conditional coco value, i.e., how much should a player expect, conditioned on her type? In fact, this latter question is not well-defined, but it is not a deficiency particular to the coco value. As we demonstrate below, even among minmax strategies of a zero-sum game, there may be no unique payoff that a player may expect, conditioned on her type. Hence, a good definition of interim individual rationality is subtle. To illustrate, consider the following zero-sum Bayesian game:

wp .5

0, 0	1,-1
0, 0,	0, 0

wp .5

0, 0	-1, 1
0, 0	-1, 1

Now, suppose player 1 is completely informed, knowing which payoff table is being used, while player 2 is completely uninformed. Clearly the (minmax and coco) value of the game is $(0, 0)$. However, when player 1 knows the payoff table is the one on the left, player 1 must play up but it is not clear what he should expect. Player 2 may be playing left, in which case player 1 should expect zero. Or player 2 may be playing right, in which case player 1 should expect one. And all strategies are minmax strategies for player 2.

Furthermore, player 1's strategy of always playing down does meet the following tempting definition of conditional individual rationality: it guarantees him the most he can guarantee, given each type. This is because player 1 can only guarantee a payoff of 0 in the left payoff table. Similarly, in the right payoff table, player 1 can only guarantee a payoff of -1. The strategy of playing down does guarantee player 1 these minimal values. However, down is clearly an unsatisfactory strategy and does not even meet the definition of ex-ante individual rationality. Hence, the natural criterion of guaranteeing the most one can guarantee, conditioned on one's type, is a poor definition of individual rationality.

The following *threat-based* definition of individual rationality is preferable.

Definition 6. Let $G = (A, T, U, \mu)$ be a finite two-person Bayesian game. A payoff function, $p : T \rightarrow \mathbf{R}^2$, is interim individually rational for player i if there exists an opponent's threat strategy σ_{-i} such that, for any type $t_i \in T_i$ and any $a_i \in A_i$, $\mathbf{E}_\mu[u_i(a_i, \sigma_{-i}(t_{-i})) \mid t_i] \leq p_i(t_i)$.

That is, player i would rather receive $p_i(t_i)$ than face the threat of playing against σ_{-i} , for any type t_i . Theorem 3, in particular the existence of a participatory equilibrium that achieves the coco value, implies that there are individually rational payoff functions that achieve the coco value in expectation. In particular, fix any c -coordinated strategy profile. Letting $p : T \rightarrow \mathbf{R}^2$ be the expected payoff pair for any type profile, since each player has the option of not participating and the players are at Nash equilibrium, p must be individually rational for both players.

Remark 1. In certain cases, the conditional value of a zero-sum game is well-defined. For example, this is clearly the case when there are unique minmax strategies. The same is true for the coco value. In particular, when the advantage game admits unique minmax strategies, the conditional coco value is well defined.

7. 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, and 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 strong revealed-payoff assumption in this paper, the coco value offers an efficient and fair solution that can be implemented in the interim sense discussed above.

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

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

As for the advantage component, consider the constant strategies $P^{\text{dem}} = 100$ and $P^{\text{ofr}} = 0$ played by M and D, respectively, 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.

Under the definition of the coco value, the two paragraphs above imply the following:

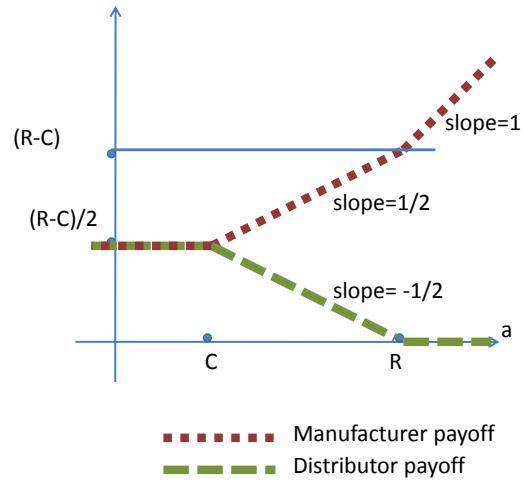
Proposition 1. *The coco value of the joint venture game above is:*

$$\left(\frac{1}{2} \mathbf{E}[\max\{R - C, 0\}] , \frac{1}{2} \mathbf{E}[\max\{R - C, 0\}] \right).$$

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 noncooperative strategies $P^{\text{dem}} = 100$ and $P^{\text{ofr}} = 0$, 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 that 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.

In general, it is difficult to compute a Bayesian-Nash equilibrium of a bargaining game, like the one above, especially when it involves an asymmetric prior probability distribution over dependent types. In contrast, the computation and implementation of the coco solution above is simple. The next example further illustrates the solution and ease of computation, by breaking the structural symmetry of the joint venture game.

7.1. 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? To keep the illustration simple, we also assume complete information, i.e., C , R , and a are common knowledge; and that trade is possible, i.e., $C < R$. Figure 2 illustrates how the coco payoffs of M and D vary as we increase a .

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

When the outside option is useless, $a < C$, it has no effect on the coco value. And when the outside option is sufficiently high to make D useless, M collects all the benefits and D is out. But in between, the coco value changes in a continuous (and piecewise linear) manner. Every extra dollar above the cost adds 50 cents to M and takes away 50 cents from D.

8. WEAKENING THE STRONG REVEALED-PAYOFF ASSUMPTION

For the purpose of achieving a general result, the interim implementation theorem uses sufficient conditions that are stronger than needed for many games. Moreover, the assessment of the player's relative advantage can sometimes be achieved in manners different from the interim protocol above. The following examples illustrate such situations.

- **Joint venture example: weak revealed payoffs suffice.** While the interim implementation theorem requires knowledge of the entire payoff function, implementation is possible under the (weak) revealed-payoff assumption. In general, revealing the entire payoff function serves only to assess the minmax value of the advantage game. So in games like the one above, where the minmax values can be assessed by easier means, the strong revealed payoff assumption is no longer needed. As was illustrated above, it is clear that the value of the advantage game is 0, hence the two players can simply form a partnership and share their net profits equally, after verifying the cost C of M and the revenue R collected by D. Furthermore, the same holds in the case of imperfect private information, where the players only have forecasts of their cost and revenue: when trade occurs, it is sufficient that the realized C and R (but not the forecasts) are revealed.
- **Hot-dog example with weak revealed payoffs.** Consider the hot-dog seller example from the beginning of the paper, in which it is known that the per-hot-dog profit of one player is twice the other's. However, for any

distribution over the number of buyers at the beach and airport, and any forecasts that these players have, the (weak) revealed-payoff assumption suffices. The strong revealed-payoff assumption requires that if both players go to the beach or airport, they would still know how many buyers were at the other location, which may be unreasonable. Note, however, that in our implementation one player will be at either location, thus the revealed pair of profits will reveal the number of buyers at each location.

- **Professional wrestling game.** Two professional wrestlers are about to participate in a match for which \$1000 will be awarded to the winner and nothing to the loser. Moreover, 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* 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 agrees 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 have a full scrimmage match. Similar in spirit, such proxy’s for determining the winner of a war are exhibited in animals in nature and in the story of *David and Goliath*.

An interesting feature of some of the examples above is that the protocols may make sense even if the players *do not have a common prior*. For example, when two wrestlers fight, the private information is in fact quite involved, including knowledge of what moves they are themselves particularly good at, beliefs about the opponent,

and higher-order beliefs. The assumption that all of these probabilities are derived from a common prior is certainly questionable in such situations. Nonetheless, in the real world, it is perfectly plausible to tell two wrestlers to “go wrestle.”

9. FURTHER REMARKS

9.1. On security levels, threats, and externalities. A common indirect approach to determine a cooperative value of a strategic game is through a “bridge” that connects the strategic theory with the cooperative theory. To every strategic game G , one associates 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 *characteristic 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 solution offers direct cooperative values for the players of strategic games, without the need for a bridge, it can still be interpreted as a special case of the bridge approach. We now proceed to explain how the coco bridge compares to another common bridge, used early on by Aumann (1961). For more recent examples and additional references, see Forge, et al. (2002), and Carpentre et al. (2005,2006).

For a two-person bimatrix game $G = (A, B)$, an associated characteristic function is determined by three numbers: V_{12} , V_1 and V_2 . To have a unique value φ_i associated with each player of G , we consider here the Shapley (1953) value: $\varphi_i \equiv V_i + \frac{1}{2}[V_{12} - (V_1 + V_2)] = \frac{1}{2}V_{12} + \frac{1}{2}(V_i - V_j)$. Thus, to determine cooperative values for G , 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 , i.e., the team-opt of G in the language of this paper. But how should we define the worth of singleton coalition V_i ? On this issue, the coco value differs from the *alternative* method used by the Aumann and the authors mentioned above.

The alternative bridge computes 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 for herself, assuming that her opponent’s goal is to minimize her payoff. Under the coco value, one computes $V_1^\kappa = \min\max\left(\frac{A-B}{2}, \frac{B-A}{2}\right)$ and $V_2^\kappa = \min\max\left(\frac{(B-A)^T}{2}, \frac{(A-B)^T}{2}\right)$, 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 example for player 1, we can see clearly the contrasts between the two methods. 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 left-hand-side game (lhs) only in the boldface entry.

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

In both games $V_{12} = 2$ is obtained by player 1 playing up. Notice, however, that the two games show substantial differences in player 1's ability to threaten and extract side payments from player 2. In the lhs game, player 1 can bring player 2's payoff down from 1 to 0 at no cost to himself, unlike the rhs game. In other words, the punishment has different externality on player 1's own outcomes.

Should such a difference be reflected in the solution of the game? The coco value reflects this difference by giving player 1 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 $\text{minmax}(\frac{A+B}{2})$, used by the coco value, reflects the difference in such externalities, whereas the individual minmax values, $\text{minmax}(A)$ and $\text{minmax}(B^T)$, cannot do so because they look at the two payoff tables separately.

The example above also sheds light on the role of the payoff-dominance axiom, used to characterize the coco value. Consider the rhs game above, with -100 being replaced by any negative number $-M$. No matter how large M is, the alternative value rewards player 1 an extra \$0.5, due to his ability to threat player 2. This is true even if the cost of carrying out the threat to player 1 is bigger than the damage to player 2. In effect, the axiom of payoff dominance puts a bound on the level of such extortion. In this particular example, player 2 gets to keep her \$1 payoff, if $-M < 0$.

The two games above motivate 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. Nevertheless, the strategic threat possibilities seem significant to us. Thus, going straight to a folk-theorem analysis, and then using any (bargaining) method of selecting a feasibly point from the individually rational feasible set, is bound to miss the effect of such externalities.

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

$$\begin{array}{|c|} \hline 0,0 \\ \hline 1,5 \\ \hline \end{array}$$

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 she can get by cutting out communication with P1 and letting him play his dominant strategy. While this intuition is clear in purely strategic environments, where communication, threats, side payments, and binding agreements are limited, in cooperative environments the outcome (3,3) may be more reasonable. To see why, consider the following example.

Example 1. *The sprinkler game.*

Two neighbors, each having to decide whether or not to water a shared lawn, play the following game:

	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).

But what if Player 2's sprinkler breaks, so that she cannot water? Then we are back in the one-player up/down game above. And if the solution of the up/down game were (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 break 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 herself inaccessible for the purpose of making side payments) in order 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 instructions to his gardener to water if and only if Player 2 gives the gardener \$4.

	pay \$4	don't pay
5,1		0,0

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

Finally, it is worth noting that a cooperative solution is inappropriate in certain contexts. For example, if the lawn is owned by P1, and P2 simply enjoys looking at it, then it may be inappropriate for P1 to demand a payment for watering his own lawn, and the noncooperative solution may be preferred. However, this is mainly a criticism of the bimatrix representation of a game for failing to capture such information. Any of the above three games could be considered a model of a situation in which player 1 owns the lawn (player 2 may volunteer to water his neighbor's lawn, especially if his neighbor does not have a sprinkler). Hence, external factors should be used to determine whether a cooperative or noncooperative solution should be applied to a particular game.

9.3. The value of information. One feature of the coco value, is that it makes optimal use of information, and compensates players for providing it, as was illustrated in the hot-dog seller example from Section 2.1.

The valuation of information has long been studied in game theory. (See, for example, Kamien, Tauman, and Zamir (1990), and the more recent references in De

Meyer, Lehrer, and Rosenberg (2009)). A natural measure of the value of information may be developed through the coco value, by considering how the coco value of the players change, as you change the information of one or both players. Such a measure is fairly sophisticated, since it takes 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.

In the hot-dog sellers game of Section 2.1, consider the possibilities that each player is either completely informed of the weather, or has no information about the weather. This gives rise to four different games. The coco values of these games, rounded to the nearest integers, are as follows.

	P2 uninformed	P2 informed
P1 uninformed	95, 145	85,175
P1 informed	115,145	100,160

Starting from the case of no information at all, the respective values of perfect weather information (PWI) acquired by player 1 are (20,0), whereas the respective values of PWI acquired by player 2 are (-10,30). Notice that player 2's PWI lowers player 1's value, while player 1's PWI does not lower player 2's value. This is a result of the tradeoff between how much the information increases the team total and how much it increases one player's advantage over the other.

9.4. Computational complexity. In the case of complete information, the coco value can be computed in polynomial time, that is, time which is 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 $(\text{size})^{\mathcal{O}(m)}$. 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 $T_i = \{1, 2, \dots, |T_i|\}$, respectively. The prior distribution μ is represented by a list of triples, $t, u, \mu(t, u)$, where t is a type profile, u is a matrix, and $\mu(t, u)$ is in $(0, 1]$. As is standard, we assume that all these numbers are rational numbers encoded as the ratios of binary integers. The size of the game, $|G|$, is simply the total number of bits used to describe the game.

Observation 2. *There is a constant $c > 0$ and an algorithm such that, given any two-player Bayesian game $G = (A, T, U, \mu)$, 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 $\mathbf{E}_\mu[u(a)|t]$ is straightforward to evaluate, and the team optimal is:

$$\sum_{t \in T} \Pr_\mu[t] \max_{a \in A} \mathbf{E}_\mu[u_1(a) + u_2(a)|t].$$

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. Specifically, one constructs the $|A_1|^{|T_1|} \times |A_2|^{|T_2|}$ bimatrix game in which each strategy (a 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 computable 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. ■

9.5. Composability. Composability of protocols has become increasingly recognized as an important topic in computer science and specifically within cryptography. While cryptographic protocols have typically been 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 does not 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 exhibit *universal composability*. First, optimal play in a repeated zero-sum game is simply optimal play in each stage. Moreover, suppose two players are to play m fixed zero-sum games G_1, G_2, \dots, G_m , either in parallel or serially, or by some combination thereof. 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 sums of the payoffs achieved in the constituent games. The minmax value of G , regardless of the particular order in which moves in G_i ’s are played, is equal to the sum of the minmax 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, all 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 the composition of team games is simple. The coco value inherits the appealing composability properties of both team and zero-sum games. Suppose a Bayesian game is played repeatedly, with types drawn freshly each round. Then the coco value of the infinitely repeated game is equal to the value of G .

10. CONCLUSION

While strictly competitive (zero-sum) and strictly cooperative (team) games represent opposite extremes in strategic interactions, both have beautiful game-theoretic solutions with a great number of desirable properties, even in the presence of incomplete information. Every Bayesian zero-sum game has optimal strategies in which players do not reveal any of their information. Every Bayesian team game has optimal cooperative play in which players share all relevant information. The coco solution to Bayesian games, via a simple decomposition of any game into its zero-sum

and team components, inherits many of these desirable properties. Moreover, under a strong revealed payoff assumption, it may be implemented by simple partnership protocols in which the two games are played separately. First, the zero-sum game is played fictitiously; this determines a “fair” compensation which can be viewed as a measure of the strategic and informational advantage that one player has over the other. In this phase, they have no incentive to share information. Then, the two form a partnership and play cooperatively, sharing all information and dividing payoffs equally.

It is interesting that cooperation is most often associated with other game representations: coalitional studies usually employ the characteristic form, and bargaining and implementation studies generally consider feasible sets of allocations or outcomes. These representations miss important strategic considerations and externalities that are captured by the normal-form representation and by Bayesian games. The setting we study 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 (or perhaps because of) the fact that we look at both types of issues, the formula and computations in our approach are relatively simple.

In complex negotiations with incomplete information, it is useful to have a programmatic way to determine a “fair” agreement in order to make sure that an agreement is reached. While the solution here may not be immediately suitable for every such negotiation, it may shed light on how people should and do reach agreements in the presence of asymmetric private information.

Our justification of the coco solution is both axiomatic, based on principles of fairness and efficiency; and strategic, showing that, through the use of a partnership type of protocol, it can be implemented as required by the Nash program. It can even be implemented in the interim sense, after players acquire private information. While axiomatization and strategic implementation are standard methods of justifying solutions in cooperative and strategic game theory, respectively, there are synergies in having both justifications for the same solution concept.

It is worth pointing out that the coco value is intended for games *where cooperation makes sense*. Of course, there are many situations in which a cooperative approach is inappropriate. For example, suppose your neighbor threatened to paint your house pink. The coco value of such a game would have you paying your neighbor not to paint your house. His threat is probably not credible, and possibly illegal as well. For such a game, a noncooperative solution makes more sense.

Finally, this work suggests several directions for future directions. First, is there a natural coco value for games with more than two players? Even the case of three players is interesting, and it is not clear whether there will be a single solution that possesses the great number of appealing properties shared by two-player zero-sum and team games. Second, can the coco value of two (or more) players be extended to the general case of (NTU) strategic games? This direction may require a substantially more refined discussion of the various axioms, as suggested by the conflicts among the many bargaining solutions in NTU cooperative games. But this direction is important if we wish to study applications where cooperation involves the optimal allocation of risk in Bayesian environments. Third, it would be interesting to consider applications of the proposed solution, which is very general, to cooperation in specific types of two-player games where players may benefit from cooperation

but suffer from not having a principled method for finding a mutually satisfactory agreement. Finally, perhaps the most interesting research direction would be to experimentally test the solution across a number of two-player games, as well as to try to identify experimentally which axioms are most violated in real-world or experimental play.

11. REFERENCES

- Aumann, R. J., “The Core of a Cooperative Game Without Side Payments,” *Transactions of the American Mathematical Society*, **98**(3):539–552, 1961.
- Binmore, K., “Game Theory and the Social Contract, Vol 1: 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.
- Canetti, R., *Universally composable security: A new paradigm for cryptographic protocols Proceedings of the 42nd IEEE symposium on the Foundations of Computer Science (FOCS)*, 136–145, 2001.
- 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.
- De Clippel, G. and E. Minelli, “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, 2009.
- Fehr, E. and K. M. Schmidt, “A Theory Of Fairness, Competition, and Cooperation,” *Quarterly Journal of Economics*, **114**(3):817–868, 1999.
- Fershtman, C., K. Judd, and E. Kalai, “Observable contracts: strategic delegation and cooperation,” *International Economic Review*, **32**(3):551–559, 1991.
- Forges, F., Mertens, J.F. and Vohra, R., “The Ex Ante Incentive Compatible Core in the Absence of Wealth Effects,” *Econometrica*, **70**: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*, **14**(3):159–182, 1967.
- Ichiishi, T. and A. Yamazaki, *Cooperative Extensions of the Bayesian Game*. Hackensack, NJ: World Scientific Publishing Co. Pte Ltd, 2006.
- Jackson, M. O., “A Crash Course in Implementation Theory,” *Social Choice and Welfare*, **18**(4):655–708, 2001.

- Jackson, M. O. and S. Wilkie, “Endogenous Games and Mechanisms: Side Payments among Players,” *Review of Economic Studies* **72**(2):543–566, 2005.
- Kalai, A. T., E. Kalai, E. Lehrer, and D. Samet, “A Commitment Folk Theorem,” *Games and Economic Behavior*, **69**(1):127–137, 2010.
- 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**:129–153, 1990.
- Kushilevitz, E., and N. Nisan, *Communication Complexity*, Cambridge University Press, 1996.
- Mezzetti, C., “Mechanism Design with Interdependent Valuations: Efficiency,” *Econometrica* **72**(5):1617–1626, 2004.
- Myerson, R.B., “Two-Person Bargaining Problems with Incomplete Information,” *em Econometrica*, **52**(2):461–488, 1984.
- Myerson, R.B. and W.L. Thomson, “Monotonicity and independence axioms,” *International Journal of Game Theory*, **9**:37–49, 1980.
- Myerson, R. B., and Satterthwaite, M.A., “Efficient Mechanisms for Bilateral Trading,” *Journal of Economic Theory* **29**, 1983, 265–281.
- 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*, 1928, **100**, 295–300.
- 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,” *American Economic Review*, **83**(5):1281–1302, 1993.
- Raiffa, H., “Arbitration schemes for generalized two-person games,” in *Contributions to the Theory of Games II*, H. Kuhn, 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**(2):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 University Press, **52**:577–626, 1964.
- Shapley, L., “A value for n-person games,” *Contributions to the Theory of Games*, Vol. II, Princeton University Press, 307–317, 1953.
- Tennenholtz, M., “Program equilibrium,” *Games and Economic Behavior*, **49**:363–373, 2004.

- 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.