

1 Introduction

Subjectivity and correlation, though formally related, are conceptually distinct and independent issues. We start by discussing subjectivity.

A mixed strategy in a game involves the selection of a pure strategy by means of a random device. It has usually been assumed that the random device is a coin flip, the spin of a roulette wheel, or something similar; in brief, an “objective” device, one for which everybody agrees on the numerical values of the probabilities involved. Rather oddly, in spite of the long history of the theory of subjective probability, nobody seems to have examined the consequences of basing mixed strategies on “subjective” random devices, i.e. devices on the probabilities of whose outcomes people may disagree (such as horse races, elections, etc.). Even a fairly superficial such examination yields some startling results, as follows:

- a. Two-person zero-sum games lose their “strictly competitive” character. It becomes worthwhile to cooperate in such games, i.e. to enter into binding agreements.¹ The concept of the “value” of a zero-sum game loses some of its force, since *both* players can get more than the value (in the utility sense).
- b. In certain n -person games with $n \geq 3$ new equilibrium points appear, whose payoffs strictly dominate the payoffs of all other equilibrium points.²

Result (a) holds not just for certain selected 2-person 0-sum games, but for practically³ all such games. Moreover, it holds if there is any area whatsoever of subjective disagreement between the players, i.e., any event in the world (possibly entirely unconnected with the game under consideration) for which players 1 and 2 have different subjective probabilities.

The phenomenon enunciated in Result (b) shows that not only the 2-person 0-sum theory, but also the non-cooperative n -person theory is modified in an essential fashion by the introduction of this new kind of strategy. However, this phenomenon cannot occur⁴ for 2-person games

This chapter originally appeared in *Journal of Mathematical Economics* 1 (1974): 67–96. Reprinted with permission.

1. Example 2.1 and sect. 6.

2. Example 2.3.

3. Specifically, wherever there is one payoff greater than the value and another one less than the value (for player 1, say).

4. Except in a very degenerate sense; see Example 2.2.

(zero-sum or not); for such games we will show⁵ that the set of equilibrium payoff vectors is not changed by the introduction of subjectively mixed strategies.

We now turn to correlation. Correlated strategies are familiar from cooperative game theory, but their applications in non-cooperative games are less understood. It has been known for some time that by the use of correlated strategies in a non-cooperative game, one can achieve as an equilibrium any payoff vector in the convex hull of the mixed strategy (Nash) equilibrium payoff vectors. Here we will show that by appropriate methods of correlation, even points outside of this convex hull can be achieved.⁶

In describing these phenomena, it is best to view a randomized strategy as a random variable with values in the pure strategy space, rather than as a distribution over pure strategies. In sect. 3 we develop such a framework; it allows for subjectivity, correlation, and all possible combinations thereof. Thus, a side product of our study is a descriptive theory (or taxonomy) of randomized strategies.

We are very grateful to Professors M. Maschler, B. Peleg, and R. W. Rosenthal for extremely stimulating correspondence and comments on this study.

2 Examples

(2.1) *Example* Consider the familiar two-person zero-sum game “matching pennies,” which we write in matrix form as follows:

1, -1	-1, 1
-1, 1	1, -1

Let D be an event to which players 1 and 2 ascribe subjective probability $2/3$ and $1/3$ respectively. Suppose now that the players bindingly agree to the following pair of strategies: Player 2 will play left in any event; player 1 will play top if D occurs, and bottom otherwise. The expectation of both players is then $1/3$, whereas the value of the game is of course 0 to both players.

To make this example work, one needs an event D whose subjective probability is $>1/2$ for one player and $<1/2$ for the other. Such an event

5. Sect. 5.

6. Examples 2.5 through 2.9.

can always be constructed as long as the players disagree about something, i.e., there is some event B with $p_1(B) \neq p_2(B)$ (where p_i is the subjective probability of i). For example, if $p_1(B) = 7/8$ and $p_2(B) = 5/8$, then we can construct the desired event by tossing a fair coin twice (provided these particular coin tosses were not involved in the description of B); if C is the event “at least one head,” then we have $p_1(B \cap C) = 21/32 > 1/2 > 15/32 = p_2(B \cap C)$. This kind of construction can always be carried out (see the proof of Proposition 5.1).

(2.2) *Example* Suppose that there is an event D to which players 1 and 2 ascribe subjective probabilities 1 and 0 respectively and such that only player 1 is informed before the play as to whether or not D occurs. Consider the following pair of strategies: Player 2 plays left; player 1 plays top if D occurs, bottom otherwise. This strategy yields a payoff of $(1, 1)$; moreover, it is actually in equilibrium, in the sense that neither player has any incentive to carry out a unilateral change in strategy. However, it has a somewhat degenerate flavor, since it requires that at least one of the two players be certain of a falsehood (see sect. 8). We now show that a similar phenomenon can also occur in a non-degenerate set-up; but it requires 3 players.

(2.3) *Example* Consider the three-person game given as follows:

0, 8, 0	3, 3, 3	0, 0, 0	3, 3, 3
1, 1, 1	0, 0, 0	1, 1, 1	8, 0, 0

here player 1 picks the row, player 2 the column, and player 3 the matrix. A triple of objectively mixed strategies is an equilibrium point in this game if and only if player 1 plays bottom and player 2 plays left; player 3 may play any mixed strategy. All these equilibrium points have the same payoff vector, namely $(1, 1, 1)$. Suppose now that D is an event to which players 1 and 2 ascribe subjective probabilities $3/4$ and $1/4$ respectively, and such that only player 3 is informed as to whether or not D occurs. Consider the following strategy triple: player 1 plays top; player 2 plays right; player 3 chooses the left matrix if D occurs, the right matrix if not. If player 1 moves down he will get $1/4 \cdot 8 = 2$ rather than $1 \cdot 3 = 3$; similarly for player 2 if he moves left; as for player 3, he certainly cannot profit from moving, since he is getting his maximum payoff in the whole game. Therefore this is an equilibrium point; its payoff is $(3, 3, 3)$, instead of the $(1, 1, 1)$ obtained at all objective equilibrium points.

That this kind of phenomenon cannot occur with fewer than 3 players will be shown in sect. 5.

An interesting feature of this example is that the higher payoff at the new, subjective, equilibrium point is not only “subjectively higher,” it is “objectively higher.” That is, unlike the case in Example 2.1, the payoff is higher not only because of the differing probability estimates of the players, but it is higher in any case, whether or not the event D takes place. The contribution of subjective probabilities in this case is to make the new point an equilibrium; once chosen, it is *sure* to make all the players better off than at any of the old equilibrium points.

We remark that by using different numbers we can find similar examples based on arbitrarily small probability differences.

It is essential in this example that only player 3 be informed as to whether or not D occurs. If, say, player 1 were also informed (before the time comes for choosing his pure strategy), he could do better by playing top or bottom according as to whether D occurs or not.⁷ This secrecy regarding D is quite natural. If we were to insist that all players be informed about D , it would be like insisting that in an objectively mixed strategy based on a coin toss, all players be informed as to the outcome of the toss. But much of the effectiveness of mixed strategies is based precisely on the secrecy, which would then be destroyed. In practical situations, it is of course quite common for some players to have differing subjective probabilities for events about which they are not informed, and on which other players peg their strategy choices.

Our last 6 examples deal with correlation.

(2.4) *Example* Consider the following familiar 2-person non-zero-sum game:

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

There are exactly 3 Nash equilibrium points: 2 in pure strategies, yielding (2, 1) and (1, 2) respectively, and one in mixed strategies, yielding $(2/3, 2/3)$. The payoff vector $(3/2, 3/2)$ is not achievable at all in (objectively) mixed strategies. It is, however, achievable in “correlated” strategies, as follows: One fair coin is tossed; if it falls heads, players 1 and 2 play top and left respectively; otherwise, they play bottom and right.

The interesting aspect of this procedure is that it not only achieves $(3/2, 3/2)$, it is also in equilibrium; neither player can gain by a unilateral change. Any point in the convex hull of the Nash equilibrium payoffs of

7. I am grateful to R. W. Rosenthal for this remark.

any game can be achieved in a similar fashion, and will also be in equilibrium. This is not new; it has been in the folklore of game theory for years. I believe the first to notice this phenomenon (at least in print) were Harsanyi and Selten (1972).

In the following 4 examples we wish to point out a phenomenon that we believe *is* new, namely that by the use of correlated strategies one can achieve a payoff vector that is in equilibrium in the same sense as above, but that is *outside* the convex hull of the Nash equilibrium payoffs. In fact, except in Example 2.7, it is better for all players than any Nash equilibrium payoff.

(2.5) *Example* Consider the 3-person game given as follows:

0, 0, 3	0, 0, 0	2, 2, 2	0, 0, 0	0, 0, 0	0, 0, 0
1, 0, 0	0, 0, 0	0, 0, 0	2, 2, 2	0, 1, 0	0, 0, 3

Here player 1 picks the row, player 2 the column, and player 3 the matrix. If we restrict ourselves to pure strategies, there are only 3 equilibrium payoffs, namely $(0, 0, 0)$, $(1, 0, 0)$, and $(0, 1, 0)$. If we allow (objectively) mixed strategies, some more equilibrium payoffs are added, but none of their coordinates exceed 1. Consider now the following strategy triple: Player 3 plays the middle matrix. Players 1 and 2 get together and toss a fair coin, but do not inform player 3 of the outcome of the toss. If the coin falls on heads, players 1 and 2 play top and left respectively; otherwise, they play bottom and right respectively. The payoff is $(2, 2, 2)$. If player 3 would know the outcome of the toss, he would be tempted to move away; for example, if it was heads, he would move left. Since he does not know, he would lose by moving. Thus the introduction of correlation among *subsets* of the players can significantly improve the payoff to everybody.

(2.6) *Example* Another version of Example 2.5 is the following:

0, 1, 3	0, 0, 0	2, 2, 2	0, 0, 0	0, 1, 0	0, 0, 0
1, 1, 1	1, 0, 0	2, 2, 0	2, 2, 2	1, 1, 1	1, 0, 3

This version has the advantage that there is only one Nash equilibrium payoff, namely $(1, 1, 1)$; and it can be read off from the matrices by the use of simple domination arguments, as compared to the slightly laborious computations needed in the previous example. The advantage of

Example 2.5 is that the new correlated strategy equilibrium point has the property that any deviation will actually lead to a loss (not only a failure to gain).

In both examples, player 3 will not even *want* to know the outcome of the toss; he will *want* players 1 and 2 to perform it in secret. It is important for him that players 1 and 2 know that he does not know the outcome of the toss; otherwise they cannot depend on him to choose the middle matrix, and will in consequence themselves play for an equilibrium point that is less advantageous for all. Thus in Example 2.6, player 3 knows that if he can “peep,” then players 1 and 2 will necessarily play bottom and left respectively, to the mutual disadvantage of all.

(2.7) *Example* In Examples 2.5 and 2.6, equilibrium payoffs outside of the convex hull of the Nash equilibrium payoffs were achieved by “partial” correlation—correlation of strategy choices by 2 out of the 3 players. We now show that a similar phenomenon can occur even in 2-person games; here again a kind of partial correlation is used, but it is subtler than that appearing previously.

Consider the 2-person game given as follows:

6, 6	2, 7
7, 2	0, 0

This game has two pure strategies equilibrium points, with payoff (2, 7) and (7, 2) respectively, and one mixed strategy equilibrium point, with payoff $(4\frac{2}{3}, 4\frac{2}{3})$. Consider now an objective chance mechanism that chooses one of three points A , B , C with probability $1/3$ each. After the point has been chosen, player 1 is told whether or not A was chosen, and player 2 is told whether or not C was chosen; nothing more is told to the players. Consider the following pair of strategies: If he is informed that A was chosen, player 1 plays his bottom strategy; otherwise, his top strategy. Similarly, if he is informed that C was chosen, player 2 plays his right strategy; otherwise, he goes left. It is easy to verify that this pair of strategies is indeed in equilibrium, and that it yields the players an expectation of (5, 5)—a payoff that is outside the convex hull of the Nash equilibrium payoffs (listed above).

The strategies in this equilibrium point are not stochastically independent, like the mixed or pure strategies appearing in Nash equilibrium points, nor are they totally correlated, like the strategies appearing in Example 2.4. It is precisely the partial correlation that enables the phenomenon we observe here.

(2.8) *Example* Consider the 2-person game given as follows:

6, 6	0, 0	2, 7
0, 0	4, 4	3, 0
7, 2	0, 3	0, 0

This is obtained from the previous example by adding a middle row and a middle column, with appropriate payoffs. The strategy pairs (top, right) and (bottom, left), which previously yielded the equilibrium payoffs (2, 7) and (7, 2), are no longer in equilibrium; the equilibrium property has been “killed” by the addition of the new strategies. The mixed-strategy pair which previously yielded $(4\frac{2}{3}, 4\frac{2}{3})$ remains in equilibrium here; moreover we get a new pure strategy equilibrium point yielding (4, 4), and a new mixed strategy equilibrium point yielding $(2\frac{10}{23}, 2\frac{10}{23})$. These are all the Nash equilibrium payoffs. Thus we see that $4\frac{2}{3}$ is the maximum that either player can get at any Nash equilibrium point.

The equilibrium point in “partially correlated” strategies described in Example 2.7 remains in equilibrium here, and yields (5, 5)—more, for both players, than that yielded by any of the Nash equilibria.

Though the random device needed in Examples 2.7 and 2.8 is conceptually somewhat more complex than the coin tosses (secret or joint) used in classical game theory and in Examples 2.4, 2.5, and 2.6, it is not at all difficult to construct. Given a roulette wheel, it is easy to attach electrical connections that will do the job. If the reader wishes, he can think of the players as jointly supervising the construction of the device, and then retiring to separate rooms to get the information and choose their strategies. It is advantageous for both players to build the device, to satisfy each other that it is working properly, and then to follow the above procedure. Once chosen, the procedure is of course self-enforcing; neither player will wish to renege at any stage.

(2.9) *Example* Consider again the 2-person 0-sum game “matching pennies” already treated in Example 2.1. Suppose that nature chooses one of four points $A, B, C,$ or D , and that players 1 and 2 ascribe to these choices subjective probabilities $(1/3, 1/6, 1/6, 1/3)$ and $(1/6, 1/3, 1/3, 1/6)$, respectively. Under no condition is either player told which point was chosen, but player 1 is told whether the point chosen is in $\{A, B\}$ or in $\{C, D\}$, and player 2 is told whether the point chosen is in $\{A, C\}$ or in $\{B, D\}$. Consider the following strategy pair: Player 1 chooses top or bottom according as to whether he is told “ $\{A, B\}$ ” or “ $\{C, D\}$ ”; player 2 chooses left or right according as to whether he is told “ $\{A, C\}$ ” or

“ $\{B, D\}$.” Like the strategy pair in Example 2.1, this yields each player an expectation of $1/3$, whereas the value of the game is of course 0 to both players; unlike the strategy pair of Example 2.1, *this pair is in equilibrium*.

If, in a 2-person zero-sum game with value v , we permit correlation but rule out subjectivity, then the payoffs to all strategy pairs continue to sum to 0, and it is easy to see that the only equilibrium points have payoffs $(v, -v)$. If we rule out correlation (i.e. permit mixed or pure strategies only) but permit subjectivity, we still get only $(v, -v)$ as an equilibrium payoff; this follows from Proposition 5.1. But if we permit both subjectivity and correlation, then this example shows that there may exist mutually advantageous equilibrium points, even in 2-person 0-sum games.

3 The Formal Model

In this section, definitions are indicated by italics.

A *game* consists of:

1. a finite set N (the *players*); write $N = \{1, \dots, n\}$;
2. for each $i \in N$, a finite set S_i (the *pure strategies* of i);
3. a finite set X (the *outcomes*);
4. a function g from the cartesian product $S = \mathbf{x}_{i \in n} S_i$ onto X (the *outcome function*).

This completes the description of the game as such; formally, however, we still need some equipment for randomizing strategies, and for defining utilities and subjective probabilities for the players. Thus to the description of the game we append the following:

5. A set Ω (the *states of the world*), together with a σ -field \mathcal{B} of subsets of Ω (the *events*);
6. For each player i , a sub- σ -field \mathcal{I}_i of \mathcal{B} (the events in \mathcal{I}_i are those *regarding which i is informed*).
7. For each player i , a relation \succsim_i (the *preference order of i*) on the space of lotteries on the outcome space X , where a *lottery* on X is a \mathcal{B} -measurable⁸ function from Ω to X .

The intuitive scenario associated with this model involves the following steps, to be thought of as occurring one after the other as follows:

8. This means that for each $x \in X$, the set $\{x = x\}$ [i.e. the set $\{\omega: x(\omega) = x\}$] is in \mathcal{B} .

- i. Nature chooses a point ω in Ω .
- ii. Each player i is informed as to which events in \mathcal{I}_i contain ω .
- iii. Each player i chooses⁹ a pure strategy s_i in S_i .
- iv. The outcome is determined¹⁰ to be $g(s_1, \dots, s_n)$.

Returning to the formal model, let us define a *randomized strategy* (or simply *strategy*) for player i to be a measurable function s_i from (Ω, \mathcal{I}_i) to S_i . Note that s_i must be measurable w.r.t. (with respect to) \mathcal{I}_i , not merely w.r.t. \mathcal{B} ; this means that i can peg his strategies only on events regarding which he is informed. [In the above scenario, strategies are chosen before step (ii).]

If s is an n -tuple of strategies, then $g(s)$ is a lottery on X . An *equilibrium point* is an n -tuple $s = (s_1, \dots, s_n)$ of strategies such that for all i and all strategies t_i of i , we have

$$g(s) \succeq_i g(s_1, \dots, s_{i-1}, t_i, s_{i+1}, \dots, s_n).$$

We will assume:

ASSUMPTION I For each i , there is a real function u_i on X (the *utility function* of i) and a probability measure¹¹ p_i on Ω (the *subjective probability* of i) such that for all lotteries x and y on X , we have

$x \succeq_i y$ if and only if

$$\int_{\Omega} u_i(x(\omega)) dp_i(\omega) \geq \int_{\Omega} u_i(y(\omega)) dp_i(\omega);$$

moreover, p_i is unique.¹²

This assumption says that player i 's preferences between lotteries are governed by the expected utility of the outcomes of the lotteries. There are well known systems of axioms on the preferences \succeq_i that lead to utilities and subjective probabilities; see for example Savage (1954) or Anscombe and Aumann (1963).

Note that p_i is defined on all of \mathcal{B} , not only on \mathcal{I}_i ; that is, i assigns a probability to all events, not only those regarding which i is informed.

9. Without informing the other players.

10. Note that g does not depend on ω ; see subsect. (d) of sect. 9.

11. σ -additive non-negative measure with $p_i(\Omega) = 1$.

12. I.e. if (u'_i, p'_i) also satisfy the condition of the previous sentence, then $p'_i = p_i$. Axiom systems leading to subjective probabilities usually imply the uniqueness as well. In our case, if we had wanted to minimize our assumptions we could have avoided assuming uniqueness, at this stage. Indeed, the uniqueness of p_i follows from the *existence* of a non-atomic p_i ; and this is assumed in Assumption II. It is, however, more convenient to assume uniqueness already at this stage, since it enables us to refer immediately to "the" subjective probability of player i , and this simplifies the entire discussion.

Two events A and B are *i-independent* if

$$p_i(A \cap B) = p_i(A)p_i(B);$$

intuitively, this means that i can get no hint about A from information regarding B . More generally, the events A_1, \dots, A_k are *i-independent* if

$$p_i(B_1 \cap \dots \cap B_k) = p_i(B_1) \cdots p_i(B_k)$$

whenever each B_j is either A_j or Ω . Events are *independent* if they are *i-independent* for all i . An event is *i-secret* if it is in \mathcal{S}_i , and for each j other than i , it is j -independent of all events in the σ -field generated by all the \mathcal{S}_k with $k \neq i$. Thus i is informed regarding each *i-secret* event, but players other than i can get no hint of it, even by pooling their knowledge. The family of *i-secret* events is denoted \mathcal{S}_i . We assume

ASSUMPTION II For each i , there is a σ -field \mathcal{R}_i of *i-secret* events such that each p_j is non-atomic on \mathcal{R}_i .

Intuitively, \mathcal{R}_i can be constructed from a roulette spin conducted by i in secret. To obtain the non-atomicity, it is sufficient to assume that each player j assigns subjective probability 0 to each particular outcome;¹³ or equivalently, that there exist finite partitions of Ω into \mathcal{R}_i -events whose p_j -probabilities are arbitrarily small. Such non-atomicity assumptions are familiar in treatments of subjective probability.¹⁴ Note that the roulette wheel need not be “objective,” i.e. the players may disagree about the probabilities involved.

A strategy s_i of i is *mixed* if it is \mathcal{S}_i -measurable, i.e. pegged on *i-secret* events. Strategies s_1, \dots, s_n are *uncorrelated* if for each $s \in S$, the n events $\{s_j = s_j\}$ are independent. Mixed strategies are uncorrelated, but the converse is false; strategies may be uncorrelated even though they are pegged on events regarding which all players are informed.

An event A is called *objective* if all the subjective probabilities $p_i(A)$ coincide; in that case their common value is called the *probability* of A , and is denoted $p(A)$. A strategy s_i of i is called *objective* if it is pegged on objective events, i.e. if $\{s_i = s_i\}$ is an objective event for all $s_i \in S_i$. The strategies occurring in the classical non-cooperative theory are precisely those that are both objective and mixed.

An event or strategy is called *subjective* if it is not objective.

The preference order of i determines his utility function on X only up to a monotonic linear transformation. Given a specific choice u_1, \dots, u_n of utility functions of the players, define the *payoff function* of player i in

13. The set of outcomes is taken to be the unit interval (not $\{0, 1, \dots, 36\}$).

14. Cf. Savage (1954, pp. 38–40, especially postulates P6' and P6).

the usual way; that is, if s is an n -tuple of pure strategies (i.e. $s \in S$), define

$$h_i(s) = u_i(g(s));$$

and if s is an n -tuple of arbitrary strategies, define

$$H_i(s) = E_i(h_i(s)) = \int_{\Omega} h_i(s(\omega)) dp_i(\omega), \quad (3.1)$$

E_i being the expectation operator w.r.t. the probability measure p_i on Ω . Then $g(s) \succ_i g(t)$ if and only if $H_i(s) > H_i(t)$. Write

$$h = (h_1, \dots, h_n),$$

$$H = (H_1, \dots, H_n).$$

Any vector of the form $H(s)$ is called a *feasible payoff*; if s is an equilibrium point, then $H(s)$ is called an *equilibrium payoff*; and if s is an equilibrium point in objective mixed strategies, then $H(s)$ is an *objective mixed equilibrium payoff*.

To orient the reader, we mention that Example 2.1 involves a pair of subjective strategies that is not in equilibrium; Examples 2.2 and 2.3 involve equilibrium points in mixed subjective strategies; and Examples 2.4 through 2.8 involve equilibrium points that are objective, but are not in mixed strategies (the strategies are in fact correlated, i.e. not uncorrelated).

4 Preliminaries and Generalities

This section is devoted to the *statement* of several lemmas needed in the sequel, of an existence theorem for equilibrium points (Proposition 4.3), and of a proposition concerning the nature of the sets of feasible and equilibrium payoffs. We also discuss the notion of “correlation.”

Though most of the results stated in this section are intuitively unsurprising, the proofs are a little involved (because of the relatively weak assumptions we have made). *We therefore postpone these proofs to sect. 7*, in order to get as quickly as possible to the conceptually more interesting results of the paper (Propositions 5.1 and 6.1). Readers who are interested only in the statements of these propositions (as opposed to the proofs) may skip this section.

The first lemma asserts that objective mixed strategies can be constructed with arbitrary probabilities, i.e. that all the “classical” mixed strategies appear in this model as well. Define a *distribution* on S_i to be a real-valued function on S_i whose values are non-negative and sum to 1.

(4.1) LEMMA *Let $i \in N$, and let σ_i be a distribution on S_i . Then i has an objective mixed strategy s_i such that for all $s_i \in S_i$,*

$$p\{s_i = s_i\} = \sigma_i(s_i).$$

Next, we have

(4.2) LEMMA *For any i in N , any event B , and any α between 0 and 1, there is an objective i -secret event with probability α that is independent of B .*

Intuitively, Lemmas 4.1 and 4.2 depend on the existence of an objective roulette wheel that i can spin in secret. Thus they appear to go somewhat further¹⁵ than Assumption II, and it is of some interest that in fact, they follow from it.

The next proposition asserts that the classical equilibrium points of Nash (1951) appear in this model as well. If $\sigma = (\sigma_1, \dots, \sigma_n)$ is an n -tuple of distributions on S_1, \dots, S_n respectively, and if $i \in N$, define

$$F_i(\sigma) = \sum_s h_i(s) \prod_{j \in N} \alpha_j(s_j),$$

where the sum runs over all pure strategy n -tuples $s = (s_1, \dots, s_n)$. The payoff to σ is the n -tuple $F(\sigma) = (F_1(\sigma), \dots, F_n(\sigma))$. Recall that σ is a Nash equilibrium point if for any i and any distribution τ_i on S_i , we have

$$F_i(\sigma) \geq F_i(\sigma_1, \dots, \sigma_{i-1}, \tau_i, \sigma_{i+1}, \dots, \sigma_n);$$

in that case $F(\sigma)$ is called a Nash equilibrium payoff.

(4.3) PROPOSITION *The set of Nash equilibrium payoffs coincides with the set of objective mixed equilibrium payoffs.*

From this proposition and Nash's theorem it follows immediately that there is an equilibrium point in every game.

We now come to the concept of correlation. Set $S_N = \bigcap_{i \in N} \mathcal{F}_i$; the members of S_N are called public events. A continuous chance device, or roulette for short, is a sub- σ -field \mathcal{R} of Ω on which each p_j is non-atomic; if \mathcal{R} consists of public events, it is called a public roulette.

Parallel to Lemma 4.2, we have

(4.4) LEMMA *Assume that there is a public roulette. Then for any event B , and any α between 0 and 1, there is an objective public event with probability α that is independent of B .*

15. See subsection (e) of sect. 9 for a discussion of this point.

A public roulette can be used as a correlating device on which all players can peg their choices; this leads (cf. Example 2.4) to

(4.5) PROPOSITION *Assume that there is a public roulette. Then the set of equilibrium payoffs and the set of feasible payoffs are both convex.*

It is not known whether the set of equilibrium—or feasible—payoffs is closed, whether or not one assumes the existence of a public roulette.

Public roulettes enable *all* players to correlate their strategies, as in Example 2.4. In Examples 2.5 through 2.9, the correlation is of a subtler kind. To describe the situation, let us call the triple consisting of the pair (Ω, \mathcal{B}) , the n -tuple $(\mathcal{I}_1, \dots, \mathcal{I}_n)$, and the n -tuple (p_1, \dots, p_n) a *randomizing structure*. In Examples 2.4, 2.5, and 2.6, the randomizing structure is of a particularly simple kind, which we call *standard*, and which is described as follows:

For each $T \subset N$, let J_T be a copy of the unit interval $[0, 1]$ with the Borel sets. Let $(\Omega, \mathcal{B}) = \mathbf{x}_{T \subset N} J_T$, and let π_T be the projection of Ω on J_T . For i in N , call two members ω_1 and ω_2 of Ω *i -equivalent* if $\pi_T(\omega_1) = \pi_T(\omega_2)$ for all T containing i . Let \mathcal{I}_i be the σ -field of all events in Ω that are unions of i -equivalence classes. Let all the p_i be Lebesgue measure on Ω .

Intuitively, the points of J_T are outcomes of an objective roulette spin conducted in the presence of the members of T only. Note that all the probabilities are the same, so that subjectivity does not enter the picture. Even in this relatively simple case, though, it is not clear that the set of equilibrium payoffs is closed (it is convex by Proposition 4.5). However, when $n = 2$ and the randomizing structure is standard, then it is easily verified that the set of equilibrium payoffs is precisely the convex hull of the Nash equilibrium payoffs.

In Example 2.7, the randomizing structure is not standard; but the probabilities are objective, i.e. all the p_i coincide. It is easily verified that whenever the probabilities are objective and there is a public roulette, the set of feasible payoffs is simply the convex hull of the pure strategy payoffs. In particular, this will be the case when the randomizing structure is standard.

5 Equilibrium Points in Two-Person Games

In Example 2.3 we showed that by pegging strategies on subjective events, it is possible to find a 3-person game with an equilibrium point in mixed strategies whose payoff is higher for *all* players than the payoff to

any equilibrium in objective mixed strategies. In this section we show that this cannot happen in the case of 2-person games.

(5.1) PROPOSITION *Let G be a 2-person game,¹⁶ and assume that for any event B , $p_1(B) = 0$ if and only if $p_2(B) = 0$.* (5.2)

Then for each equilibrium point s in mixed strategies, there is an equilibrium point t in objective mixed strategies such that $H(s) = H(t)$.

Proof Let (s_1, s_2) be an equilibrium point in mixed strategies. Let t_1 be an objective mixed strategy for player 1 that mimics the way player 2 sees s_1 , i.e., such that for all pure strategies s_1 of player 1,

$$p\{t_1 = s_1\} = p_2\{s_1 = s_1\}; \quad (5.3)$$

such a strategy exists because of Lemma 4.1. Similarly, let t_2 be an objective mixed strategy for player 2 such that for all s_2 ,

$$p\{t_2 = s_2\} = p_1\{s_2 = s_2\}. \quad (5.4)$$

Let s_{11}, s_{12}, \dots be the pure strategies that enter actively into s_1 , i.e., with positive p_1 -probability. Because s_1 is in equilibrium,

$$E_1 h_1(s_{1j}, s_2) = H_1(s_1, s_2) \quad (5.5)$$

for all j . But from (5.2) it follows that the s_{1j} are also precisely the pure strategies that enter actively into t_1 ; hence (5.5) yields

$$H_1(t_1, s_2) = H_1(s_1, s_2). \quad (5.6)$$

Since s_1 maximizes H_1 (if 2 plays s_2), it follows from (5.6) that t_1 also does. But by (5.4), player 1 ascribes the same distribution to t_2 as to s_2 ; therefore

$$H_1(t_1, t_2) = H_1(t_1, s_2) = H_1(s_1, s_2),$$

and furthermore, t_1 maximizes H_1 if 2 plays t_2 . Similarly,

$$H_2(t_1, t_2) = H_2(s_1, t_2) = H_2(s_1, s_2),$$

and t_2 maximizes H_2 if 1 plays t_1 . Therefore (t_1, t_2) is an objective mixed equilibrium point with the same payoff as (s_1, s_2) , and the proof of this proposition is complete.

It would have been reasonable to conjecture that this proposition remains true if both occurrences of the word “mixed” are deleted. Example 2.9 shows that this is false.

16. Not necessarily 0-sum.

6 Two-Person Zero-Sum Games

A game is called *2-person 0-sum* if $n = 2$ and if there are utility functions u_1 and u_2 for the players such that

$$u_1(x) + u_2(x) = 0$$

for all $x \in X$. After adopting a specific such pair of utility functions, one defines the value of the game as usual. Specifically, Proposition 4.3 provides an equilibrium point in objective mixed strategies; by the minimax theorem, all such equilibrium points must have the same payoff, which is of the form $(v, -v)$. Then v is called the *value* of the game.

(6.1) PROPOSITION *Let G be a 2-person 0-sum game with value v . Assume that*

there are outcomes x and y with $u_1(x) > v > u_1(y)$, and (6.2)

for each player i , there is a subjective event B_i regarding which i is informed. (6.3)

Then there is a pair s of strategies such that

$$H_1(s) > v, H_2(s) > -v. \quad (6.4)$$

Remark The proposition is considerably easier to prove if one assumes that

there is a public subjective event B and there is a public roulette. (6.5)

To see this, assume without loss of generality that $p_1(B) > p_2(B)$ (otherwise, use the complement of B). If C is a public event and $0 \leq \theta \leq 1$, denote by θC a public event such that

$$p_1(\theta C) = \theta p_1(C) \quad \text{and} \quad p_2(\theta C) = \theta p_2(C);$$

the existence of such a θC follows¹⁷ from Lemma 4.4. From (6.2) it follows that there is a β with $0 < \beta < 1$ such that

$$v = \beta u_1(x) + (1 - \beta)u_1(y).$$

For $0 < \delta < 1$, set

$$B^\delta = \Omega \setminus \delta(\Omega \setminus B).$$

17. It also follows from the theorem of Lyapunov (1940) on the range of a vector measure. We shall see in sect. 8 that Proposition 4.4 is itself proved via Lyapunov's theorem; intuitively, though, Proposition 4.4 is more transparent than Lyapunov's theorem.

Then for δ sufficiently small, we have $p_1(B^\delta) > p_2(B^\delta) > \beta$; hence if we let $A = \theta B^\delta$ for an appropriate θ , then

$$p_1(A) > \beta > p_2(A). \tag{6.6}$$

Now define s as follows: Both players choose pure strategies leading to x , or both players use pure strategies leading to y , according as to whether A does or does not occur. Then (6.4) follows, and the proposition is proved under the assumption of (6.5).

Proof of Proposition 6.1 In (6.2), let $x = g(r_1, r_2)$ and $y = g(t_1, t_2)$; then $h_1(r_1, r_2) > v > h_1(t_1, t_2)$.

Suppose first that $h_1(r_1, t_2) \geq v$. Setting

$$H^*(\varepsilon) = \varepsilon h_1(r_1, r_2) + (1 - \varepsilon)h_1(r_1, t_2)$$

$$H_*(\varepsilon) = \varepsilon h_1(t_1, r_2) + (1 - \varepsilon)h_1(t_1, t_2),$$

we find that for $\varepsilon > 0$ sufficiently small,

$$H^*(\varepsilon) > v > H_*(\varepsilon).$$

Hence there is a β in $[0, 1]$ with

$$v = \beta H^*(\varepsilon) + (1 - \beta)H_*(\varepsilon).$$

Now let A_1 be a subjective event in \mathcal{I}_1 with $p_1(A_1) > \beta > p_2(A_1)$. The existence of such an event may be established as in the remark [cf. (6.6)], except that now the mapping $C \rightarrow \theta C$ takes \mathcal{I}_1 into itself (rather than $\mathcal{I}_1 \cap \mathcal{I}_2$ into itself), Lemma 4.2 (rather than 4.4) is used to prove the existence of θC , and B_1 is substituted for B . Again using Lemma 4.2, we obtain an objective 2-secret event A_2 with probability ε that is independent of A_1 . Define a strategy pair s by stipulating that s_i takes the value r_i or t_i according as to whether A_i does or does not occur. Then s is uncorrelated and satisfies (6.4).

If $h_1(r_1, t_2) \leq v$, the proof is similar, the roles of the two players being then reversed. This completes the proof of Proposition 6.1.

Proposition 6.1 fails if it is only assumed that there is a subjective event of which at least one player is informed. For example, if in the 2-person 0-sum game with matrix

1	1
2	0

,

there is a subjective event in \mathcal{S}_1 but not in \mathcal{S}_2 , then there is no pair of strategies satisfying (6.4). In this game it is sufficient for (6.4) that there be a subjective event in \mathcal{S}_2 . In general, the proof of Proposition 6.1 shows that in any specific game, only one player need use a subjective strategy, while the other can use an objective one; but which player it is that uses the subjective strategy may depend on the game.

It is perhaps worth noting that in Proposition 6.1, the strategies constituting s may be taken to be independent; indeed, the strategies constructed in the proof are independent.

7 Proof of the Propositions of Sect. 4

We start with a lemma that is basic to the proofs in this section.

(7.1) LEMMA *Let \mathcal{R} be a roulette, and let B^1, \dots, B^l be events. Then for any α between 0 and 1, there is an objective event in \mathcal{R} with probability α that is independent of each of the B^k .*

Proof Consider the $(1+l)n$ -dimensional vector measure u on (Ω, \mathcal{R}) defined by

$$\mu_i(A) = p_i(A), \quad i = 1, \dots, n,$$

$$\mu_{kn+i}(A) = p_i(A \cap B^k), \quad i = 1, \dots, n, \quad k = 1, \dots, l.$$

That μ is non-atomic follows from the non-atomicity of the p_i on \mathcal{R} . By the theorem of Lyapunov (1940) on the range of a vector measure, the range of μ is convex. Now

$$\mu(\Omega) = (1, \dots, 1, p_1(B^1), \dots, p_n(B^1), \dots, p_1(B^l), \dots, p_n(B^l))$$

and

$$\mu(\phi) = (0, \dots, 0).$$

Hence $\alpha\mu(\Omega)$ is in the range of μ , i.e. there is a set A in \mathcal{R} such that $\mu(A) = \alpha\mu(\Omega)$. This means that $p^i(A) = \alpha$ for all i and

$$p_i(A \cap B^k) = \alpha p_i(B^k) = p_i(A)p_i(B^k)$$

for all i . The proof of the lemma is complete.

Proof of Lemma 4.1 By induction on the number m of members s_i of S_i for which $\sigma_i(s_i) > 0$. Let \mathcal{R}_i be the i -secret roulette provided by Assumption II. As often in the case of inductive proofs, it is convenient to prove a somewhat stronger statement than is needed; here we show that

a strategy s_1 can be chosen obeying Proposition 4.1, so that the events $\{s_i = s_i\}$ are in \mathcal{R}_1 . (7.2)

If $m = 1$ there is nothing to prove. Let $m > 1$, and suppose the proposition true for $m - 1$. Let $S_i = \{s_i^1, \dots, s_i^m, \dots\}$, where $\sigma_i(s_i^j) > 0$ if and only if $j \leq m$. By the induction hypothesis, there are disjoint objective events B^1, \dots, B^{m-1} such that $p(B^j) = \sigma_i(s_i^j)/(1 - \sigma_i(s_i^m))$. By Lemma 7.1 with $\mathcal{R} = \mathcal{R}_i$, there is an objective i -secret event B^m independent of each of B^1, \dots, B^{m-1} , such that $p(B^m) = 1 - \sigma_i(s_i^m)$. Set $A^j = B^j \cap B^m$ for $j < m$, $A^m = \Omega \setminus B^m$, and $A^j = \phi$ for $j > m$; then $A^j \in \mathcal{R}_i$ for all j . Define s_i by $s_i(\omega) = s_i^j$ if and only if $\omega \in A^j$. Then s_i satisfies (7.2). This completes the proof of Lemma 4.1.

Proof of Lemma 4.2 Use Assumption II, and apply Lemma 7.1 with $\mathcal{R} = \mathcal{R}_i$ and $l = 1$.

Before proving Proposition 4.3 we need a lemma.

(7.3) LEMMA *Let (s_1, \dots, s_n) be an n -tuple of strategies, and let $s \in S$. For some $i \in N$, suppose that all the s_j except possibly s_i are mixed. Then*

$$\begin{aligned} p_i\{s = s\} &= p_i\{s_i = s_i\}p_i\{s_j = s_j \text{ for all } j \neq i\} \\ &= p_i\{s = s_1\} \dots p_i\{s_n = s_n\}. \end{aligned}$$

Proof W.l.o.g. let $i = 1$, and let $A_j = \{s_j = s_j\}$ for all j ; then

$$\{s = s\} = A_1 \cap \dots \cap A_n.$$

Moreover, $A_1 \in \mathcal{S}_1$ and $A_j \in \mathcal{S}_i$ for $j > 1$. By definition of S_j , each A_j with $j > 1$ is 1-independent of $A_{j+1} \cap \dots \cap A_n$. Hence

$$\begin{aligned} p_1(A_2 \cap \dots \cap A_n) &= p_1(A_2)p_1(A_3 \cap \dots \cap A_n) \\ &= \dots = p_1(A_2) \dots p_1(A_n). \end{aligned}$$

Again by definition of \mathcal{S}_j , each A_j with $j > 1$ is 1-independent of $A_1 \cap A_2 \cap \dots \cap A_{j-1}$. Hence by the previous equation,

$$\begin{aligned} p_1(A_1 \cap A_2 \cap \dots \cap A_n) &= p_1(A_1 \cap A_2 \cap \dots \cap A_{n-1})p_1(A_n) \\ &= \dots = p_1(A_1)p_1(A_2) \dots p_1(A_n) = p_1(A_1)p_1(A_2 \cap \dots \cap A_n). \end{aligned}$$

This completes the proof of the lemma.

(7.4) COROLLARY *Let (s_1, \dots, s_n) be an n -tuple of mixed strategies. Then the s_i are independent.*

Proof of Proposition 4.3 If $s = (s_1, \dots, s_n)$ is an objective mixed equilibrium point, then for each strategy t_i of i ,

$$H_i(\mathbf{s}) \geq H_i(s_1, \dots, s_{i-1}, \mathbf{t}_i, s_{i+1}, \dots, s_n). \quad (7.5)$$

In particular, this holds when \mathbf{t}_i is itself objective and mixed. From this it follows that $H(\mathbf{s})$ is a Nash equilibrium payoff.

Conversely, if $F(\sigma)$ is a Nash equilibrium payoff, then from Lemma 4.1 and Corollary 8.4 it follows that there is an n -tuple $\mathbf{s} = (s_1, \dots, s_n)$ of objective mixed strategies with payoff $F(\sigma)$ that is in equilibrium against deviations that are restricted to objective mixed strategies; i.e. that (7.5) holds when \mathbf{t}_i is an objective mixed strategy. In particular (7.5) holds when \mathbf{t}_i is “pure,” i.e. takes on a particular value in S_i with probability 1. To show that \mathbf{s} is an equilibrium in our sense, we must show that (7.5) holds when \mathbf{t}_i is *any* strategy of i , i.e. that i cannot “correlate into” the strategies of the other players. Intuitively this follows from the fact that the s_j are mixed, i.e. pegged on j -secret events; formally, though, we must show that our definition of secrecy does indeed yield this result.

W.l.o.g. let $i = 1$. Write $S' = S_2 \times \dots \times S_n$, and $\mathbf{s}' = (s_2, \dots, s_n)$.

For any s_1 in S_1 , write

$$H_1(s_1, \mathbf{s}') = E_1(h_1(s_1, \mathbf{s}')).$$

Since (7.5) holds when \mathbf{t}_1 is “pure,” it follows that for all $s_1 \in S_1$,

$$H_1(s_1, \mathbf{s}') \leq H_1(\mathbf{s}). \quad (7.6)$$

Then for an arbitrary strategy \mathbf{t}_1 of 1, it follows from Lemma 7.3 and (7.6) that

$$\begin{aligned} H_1(\mathbf{t}_1, \mathbf{s}') &= \sum_{s \in S} p_1\{\mathbf{t}_1, \mathbf{s}' = s\} h(s) \\ &= \sum_{s_1 \in S_1} \sum_{s' \in S'} p_1\{\mathbf{t}_1 = s_1\} p_1\{\mathbf{s}' = s'\} h_1(s) \\ &= \sum_{s_1 \in S_1} p_1\{\mathbf{t}_1 = s_1\} \sum_{s' \in S'} p_1\{\mathbf{s}' = s'\} h_1(s) \\ &= \sum_{s_1 \in S_1} p_1\{\mathbf{t}_1 = s_1\} H_1(s_1, s_2, \dots, s_n) \\ &\leq \sum_{s_1 \in S_1} p_1\{\mathbf{t}_1 = s_1\} H_1(\mathbf{s}) = H_1(\mathbf{s}). \end{aligned}$$

This completes the proof of Proposition 4.3.

Proof of Lemma 4.4 Apply Lemma 7.1 with $l = 1$, taking \mathcal{R} to be a public roulette.

Proof of Proposition 4.5 We first consider the equilibrium payoffs. Let \mathbf{s} and \mathbf{t} be equilibrium points, and let $0 \leq \alpha \leq 1$. Let \mathcal{R} be the public

roulette. By Lemma 7.1, there is an objective event A in \mathcal{R} with probability α that is independent of all the events $\{s = s\}$ and $\{t = s\}$ for all s in S . Now let each player i play the strategy r_i defined as follows: if A occurs, play s_i ; if not, play t_i . In other words, i plays a given pure s_i in S_i if and only if the event

$$[A \cap \{s_i = s_i\}] \cup [(\Omega \setminus A) \cap \{t_i = s_i\}] \quad (7.7)$$

occurs. This is indeed a strategy, since the event (7.7) is in the σ -field \mathcal{I}_i . It may then be verified that r is an equilibrium point and that

$$H(r) = \alpha H(s) + (1 - \alpha)H(t), \quad (4.12)$$

so the set of equilibrium payoffs is convex. The proof that the set of feasible payoffs is convex is similar, so the proof of Proposition 4.5 is complete.

8 A Posteriori Equilibria

We would like to view an equilibrium point as a *self-enforcing agreement*. In the scenario described in sect. 3, if the players agree on an equilibrium point s before stage (ii) (the stage at which information about ω is received), no player will want to renege [i.e. choose a pure strategy other than $s_i(\omega)$] after stage (ii). More precisely, when making the agreement each player i assigns subjective probability 0 to the possibility that he will want to renege after receiving his information.

There is, however, a difficulty here. Though the subjective probability of wanting to renege is 0, this possibility is not entirely excluded; more important, it is quite possible that a player assigns positive probability to a *different* player's wanting to renege. In that case the equilibrium point can no longer be considered self-enforcing; the possibility of *somebody* wanting to renege is not negligible. This is exactly the phenomenon that is responsible for the equilibrium payoff of (1, 1) in Example 2.2.

To legitimize the view of an equilibrium point as a self-enforcing agreement, one can either

- a. make assumptions under which the possibility of some player wanting to renege is assigned probability 0 by all players; or
- b. construct a model in which it is possible to define equilibrium points at which no player *ever* wants to renege.

Specifically, we may assume that though the players may have different subjective probabilities, the concept of 'impossible' or 'negligible' is the

same for all. That is, if $p_i(B) = 0$ for one i then $p_i(B) = 0$ for all i ; in other words, the p_i are absolutely continuous with respect to each other. [For the case of two players, this is the same as (5.2).] In that case, the possibility that any player will want to renege is negligible for all players, and so can be safely ignored.

Alternatively, one could replace the concept of equilibrium point by that of ‘a posteriori equilibrium point’—a strategy n -tuple from which no player i ever wishes to move (unilaterally), even after receiving his information about nature’s choice of ω . Formally, this can be done by adding the following to the 7 items that define a game (see sect. 3).

For each player i and each ω in Ω , a relation \succsim_i^ω on the space of lotteries on X (the preference order of i given his information about ω). (8)

The relations \succsim_i^ω will also be called a *posteriori preferences*. An *a posteriori equilibrium point* is then defined to be an n -tuple s of strategies such that for all ω , all i , and all t_i in S_i , we have

$$g(s) \succsim_i^\omega g(s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_n).$$

Under appropriate assumptions one can then prove the following results:

(8.1) Theorem 5.1 holds without the hypothesis of mutual absolute continuity (5.2), if “equilibrium point” is replaced by “a posteriori equilibrium point.”

(8.2) If all the p_i are mutually absolutely continuous, then the set of payoffs to a posteriori equilibrium points coincides with the set of payoffs to equilibrium points.

The proof of (8.1) is like proof of Proposition 5.1, except that a pure strategy s_{ij} is now said to “enter actively” into the strategy s_i if there is *some* ω for which $s_i(\omega) = s_{ij}$, even if i assigns probability 0 to the set of all such ω . Note that from (8.1) it follows that $(1, 1)$ is not a posteriori equilibrium payoff in Example 2.2.

One can also prove a posteriori results that are analogous to Propositions 4.3 and 4.5.

A logically complete development of the a posteriori theory involves rewriting Assumptions I and II, redefining many of the concepts defined above (such as i -secrecy and mixed strategy), and making assumptions that relate preferences to a posteriori preferences.¹⁸ The development

18. The a posteriori preferences of player i can be derived from his a priori preferences, but (like conditional probabilities) only for p_i -almost all ω . Of course, as we saw above, it is precisely the sets of p_i -measure 0 that cause the difficulty.

therefore becomes somewhat lengthy. In view of result (8.2), it does not seem worthwhile, at least at this stage, to impose a formal description of the a posteriori theory on the reader.

9 Discussion

a. The fact that differing subjective probabilities can yield social benefit is perhaps obvious. Wagers on sporting events, stock market transactions, etc., though they can be explained by convexities in utility functions, can also be explained by differing subjective probabilities; and in reality, the latter explanation is probably at least as significant as the former. Correlation of strategies is basically also a fairly obvious idea. It is all the more surprising that these ideas have heretofore not been more carefully studied in game theory in general, and in the context of randomization in particular.

b. Our more substantive results can be broadly divided into two classes: Those having to do with equilibrium points (notably Examples 2.2 through 2.9, and Proposition 5.1), and those having to do with feasible points (notably Example 2.1 and Proposition 6.1). The former belong to the “non-cooperative” theory, the latter to the “cooperative” theory. An understanding of this distinction rests on an understanding of the concepts of communication, correlation, commitment, and contract. Let us examine these concepts in the context of this paper, noting in particular how they are related to “cooperative” and “non-cooperative” games.

Recall that in step (ii) of the scenario of sect. 3, each player i is given his information about nature’s choice of ω . By *communication* we mean communication between the players before step (ii). Strategies of the players are *correlated* if they are statistically dependent. A *commitment* is an irrevocable undertaking on the part of a player, entered into before step (ii), to play in accordance with a certain strategy.¹⁹ A *contract* (or “binding agreement” or “enforceable agreement”) is a set of commitments simultaneously undertaken by several players, each player’s undertaking being in consideration of those of the others.

Of the four terms just introduced, only “correlation” has a formal meaning in the framework of our formal model (see sect. 3, where

19. One can imagine broader meanings for the word ‘commitment,’ in which the undertaking, though irrevocable, need not be to play in accordance with a certain strategy. For example, it could be to play in accordance with one of a certain set of strategies, or it could be contingent on the choice of ω or on other commitments. We have not found it necessary to complicate the discussion by considering these broader meanings.

“uncorrelated” is defined). The other three refer not to the model itself, but to how the model is to be interpreted.

If the players can enter into commitments and contracts, we have a *cooperative game*. If not,²⁰ we have a *non-cooperative game*. Formally, cooperative and non-cooperative games are described in the same way, namely by the model of sect. 3. The difference between the two types of games lies not in the formal description itself, but in what we have to say about it—what kind of theorems we prove about it.

In a non-cooperative game, at each point of time, each player acts so as to maximize his utility at that point of time, without taking previous commitments into account. Therefore the only strategy n -tuples of interest are those that are in equilibrium (cf. sect. 8). On the other hand, in the cooperative theory the players can enter into binding agreements before step (ii); therefore a strategy n -tuple does not have to be in equilibrium to be of interest, and one is led to study all feasible n -tuples.

In sect. 8 we saw that equilibrium points can be viewed as self-enforcing agreements. This is to be contrasted with feasible points that are not in equilibrium; if the players agree before step (ii) of the scenario to play such a point, some of them will wish to renege after step (ii). Therefore a contract—and an external enforcement mechanism—are required to make such agreements stick.

Thus we see that both the non-cooperative and the cooperative theory involve agreement among the players, the difference being only in that in one case the agreement is self-enforcing, whereas in the other case it must be externally enforced. Agreement usually involves communication, so that we conclude that communication normally takes place in non-cooperative as well as cooperative games.

As for correlation, this is sometimes taken as a hallmark of cooperative games. In our view, this is a fallacy; correlation may or may not be possible in a given game, whether or not it is cooperative. Examples 2.4 through 2.9 show that correlation may be involved in agreements that are not enforceable but are self-enforcing, so that it can be significant even in a non-cooperative game. On the other hand, even in a cooperative game it may be impossible to correlate, for example because the players are in different places and cannot observe the same random events; but nevertheless they can negotiate and even execute a contract, for example by mail.

Equilibrium points can be viewed in ways other than as self-enforcing agreements. For example, they can be viewed as strategy n -tuples with

20. These are the two extremes. In many situations, some commitments and contracts are possible, others not.

the property that if, for some extraneous reason, they “come to the fore” or are suggested to the players, there will be no tendency to move away from them. Alternatively, they can be viewed as providing only a necessary condition for a satisfactory theory of non-cooperative games: if a theory is to “recommend” a specific strategy to each player, then the n -tuple of the recommended strategies must be in equilibrium. To make sense, these points of view do not require pre-play communication. But even without communication, correlation is by no means ruled out. Thus in Example 2.4, the payoff $(3/2, 3/2)$ is no more the result of communication than the payoff $(2, 1)$. All that is needed to achieve $(3/2, 3/2)$ is that a fair coin be tossed, or a similar random device be actuated, and that the outcome be communicated to both players before the beginning of play. If that is done, the payoff $(3, 2, 3/2)$ becomes a full-fledged equilibrium payoff, conceptually indistinguishable from $(2, 1)$; and that is so even in the absence of communication.

Of course, communication may be important to choose a particular equilibrium point from among all those whose payoff is $(3/2, 3/2)$. But the problem of choosing a particular equilibrium point has nothing to do with correlation—it exists in most non-cooperative games. Thus in Example 2.4, even without the opportunity for correlation, it is difficult to see why one equilibrium point rather than another should be chosen. This, incidentally, is one reason for preferring to view an equilibrium point as a self-enforcing agreement.

Finally, a word about the distinction between “commitment” and “contract.” Unlike a contract, a commitment is an *individual* undertaking; it may, for example, consist of giving irreversible instructions to an agent or a machine to act in accordance with a given randomized strategy, or it may involve an obligation to pay a large indemnity if one does not act in accordance with this strategy. If there is no such agent or machine, or nobody to enforce collection of the indemnity, it may be impossible to undertake commitments, and we will have a non-cooperative game. Commitments are important in game theory even when contracts are impossible [see for example Aumann and Maschler (1972)]. However, when it is possible to make commitments as well as to communicate, it should usually be possible to make contracts as well. Moreover, by permitting commitments one opens the door to pre-emptive tactics (such as threats) in the pre-play stage, which may easily lead the outcome away from equilibrium. Thus we feel that when one has admitted commitments, one has already gone much of the way from the non-cooperative to the cooperative theory.

To sum up: The cooperative theory requires communication as well as commitment and contracting power; and it is a priori concerned with

all feasible outcomes. The non-cooperative theory requires that there be neither commitment nor contracting power, but it permits communication; and it is concerned with equilibrium outcomes. Both theories can accommodate correlation, but do not require its presence.

c. It is interesting that those of our results that belong to the “cooperative” theory are concerned with 2-person 0-sum games, which have traditionally been considered the epitome of non-cooperative games; it has been asserted that such a game can only be played non-cooperatively, since it can never be worthwhile to reach an agreement concerning it. Our results show that this is incorrect. It is not true that a 2-person 0-sum game is strictly competitive, i.e. that the preferences of the players are always in direct opposition;²¹ both players can gain by the use of a binding agreement. Thus 2-person 0-sum games can profitably be played cooperatively.

On the other hand, when correlation is ruled out, the introduction of subjective strategies creates no new equilibrium payoffs [see (5.1)]. This indicates that subjectivity alone (without correlation) at least does not affect the non-cooperative theory of 2-person 0-sum games. However, we must be careful before reaching even this modest conclusion. The arguments for the use of minimax strategies, as presented by von Neumann and Morgenstern (1953), are of two kinds: the equilibrium arguments, which they (§17.3) call “indirect,” and what they (§14, §17.4–§17.8) call “direct” arguments.²² The indirect or equilibrium arguments are of course not affected by the introduction of subjectivity. But it is the direct arguments that have traditionally been used to support the contention that the theory of 2-person 0-sum games is conceptually more satisfactory than that of more general games. To a certain extent, these direct arguments *are* affected; implicitly, they depend on the strictly competitive character of the game, which as we know is removed by subjectivity.

d. It must be stressed that the chance elements of our model—the space (Ω, \mathcal{B}) and the σ -fields \mathcal{I}_i —are used for purposes of randomization only. Nature’s choice of a point ω in Ω does not directly affect the outcome of the game; the outcome function g does not depend on ω . This model is therefore quite different from extensive game models in which chance makes “moves” that affect the outcome; here chance is only “made available” to the players for purposes of randomization. [See also subsection (g) below.]

21. Nevertheless, even after the introduction of subjective strategies, the game remains “almost strictly competitive” in the sense of Aumann (1961).

22. These are called “guaranteed value arguments” by Aumann and Maschler (1972).

e. Our use of the terms “subjective” and “objective” is a little different from the usual. The term “subjective” often signifies a personalistic definition of probability, based on preferences, such as that of Savage (1954), whereas “objective” signifies a physical (i.e. frequency) definition.²³ This is not the distinction being made here. Here the distinction is based solely on whether the players in the game do or do not agree about the numerical value of the probability in question.²⁴

f. In this paper we have extended the theory of games in strategic (i.e. normal) form by the consideration of subjective events. A parallel extension is possible for the theory of games in extensive form. Basically, such an extension would consist of replacing the probabilities now appearing in the extensive form [Kuhn (1953)] on chance moves by *vectors* of probabilities, with one component for each player. The definition of “strategy” remains unchanged; the definition of “payoff” is modified only in that one uses the different probabilities on the chance moves in calculating the payoffs for the different players. It may be verified that the basic theorems [Kuhn (1953)] of extensive game theory—namely the theorem on pure strategy equilibrium points in games of perfect information, and on behavior strategies in games of perfect recall—go through in this case as well.

One immediate application of this remark is to the theory of Harsanyi (1967, 1968) of games of incomplete information. Harsanyi has shown that in what he calls the “consistent case,” a game of incomplete information in strategic form corresponds to a certain game of complete information in extensive form. In the “inconsistent case,” there is no such correspondence. However, if one extends the definition of extensive games in the way indicated above, then the inconsistent case can be taken care of in essentially the same way as the consistent case.

g. An alternative approach to that of this paper would be to introduce the possibilities for correlation and subjective randomization explicitly into the extensive form of the game. This would mean starting the game

23. Cf. Anscombe and Aumann (1963), where personalistic and frequency probabilities were called “probability” and “chance” respectively.

24. The fact that they may disagree forces us to adopt a personalistic definition of probability for at least some events; but there is nothing to prevent us from adopting a physical definition for other events (e.g. roulette spins) on whose probabilities the players agree. Hybrid systems in which the two kinds of probability exist side by side, have been investigated; cf. Anscombe and Aumann (1963). On the other hand, a purely personalistic view, such as Savage’s, is also perfectly satisfactory for the purposes of this paper. Incidentally, if one does adopt such a purely personalistic view, then an a priori assumption on the existence of objective roulette wheels loses its intuitive attractiveness; thus the fact that Lemmas 4.1 and 4.2 follow from Assumption II gains in significance (see the discussion after the statement of Lemma 4.2).

with chance's choice of a point in Ω , and then giving each player i that information which is "his" in accordance²⁵ with the σ -field \mathcal{I}_i . The strategic form of this enlarged game could be calculated as indicated under (f) above; the Nash equilibrium points of the enlarged game correspond to the equilibrium points, as defined in this paper, of the original game. Thus in the appropriate context, our equilibrium points are special cases of Nash equilibrium points. One difficulty with this approach is that it would lead to an infinite extensive game, since Ω is infinite. But this is by no means an insuperable difficulty; models have been studied that are entirely adequate to cover such a situation [Aumann (1964)]. We did not construct our model in this way for a number of reasons. First, none of the examples or propositions would have been simplified by such a procedure; on the contrary, the extensive form is clumsy to work with, and the mathematical treatment would presumably have been more complex. Second, one usually thinks of the extensive form as representing the originally given rules of the game. In that case extraneous random events, which do not directly affect the outcome of the game, do not belong in the extensive form; one does not introduce all the possibilities for objective mixing explicitly into the game either [see subsect. (e) above].

h. It is nevertheless useful to point out how some of our examples involving correlation look in an extensive framework. Example 2.4 may be restated as follows: Suppose we add a move at the beginning of the game in which chance simply announces the results of a coin toss; no other change is made in the game (see Figure 1). One would have thought that this could not possibly effect an essential change in the game. But in the classical theory of Nash (1951), it does; $(3/2, 3/2)$ becomes the payoff to an equilibrium point, whereas it was not one before. Similarly if chance announces the result of a roulette spin, every point in the interval connecting $(2, 1)$ to $(1, 2)$ becomes an equilibrium payoff, whereas of these points, only the end points were equilibrium payoffs before. As we said in sect. 2, all this appears already in Harsanyi and Selten (1972). But examples 2.5 and 2.6 can be treated in a similar manner. Here one must add a move in which chance tosses a fair coin and informs players 1 and 2—but not player 3—of the outcome. Again, one would have thought that the addition of such a move could not possibly affect the game in any essential manner. But it adds a Nash equilibrium point with payoff $(2, 2, 2)$, whereas in the original game no player could get more than 1 at any Nash equilibrium point. Examples 2.7 and 2.8 may be treated similarly.

25. The information sets of i immediately after chance's choice would be those subsets B of Ω such that for all A in \mathcal{I}_i , either $B \cap A = \emptyset$ or $B \subset A$.




Figure 1

i. The phenomena in Examples 2.1 and 2.3, which are the basic examples regarding subjective probabilities, do not depend on each player knowing the other players' subjective probabilities *precisely*. Thus in Example 2.1, it is sufficient that both players know that players 1 and 2 ascribe to D probabilities that are $>1/2$ and $<1/2$ respectively; in that case it will already be worthwhile to enter into a binding agreement. A similar remark holds for Example 2.3, in which it is sufficient that players 1 and 2 ascribe to D subjective probabilities that are *approximately* $3/4$ and $1/4$ respectively (the precise limits of the approximation are easily calculated).

Of course there is no particular reason to treat subjective probabilities differently from utilities—if we assume that the players' utilities are known to each other, we may as well assume the same for the subjective probabilities. The above remark only points out that precise knowledge is

not crucial. Situations in which the players' utilities and/or subjective probabilities are not known to each other can be treated by the methods of games of incomplete information [Harsanyi (1967, 1968)].

j. The view is sometimes held that when people have different subjective probabilities for the same event, this can only be due to differences in the information available to these people. Such a view has been eloquently set forth by John Harsanyi (1968, §16), and we will call it the *Harsanyi doctrine*. Suppose, for example, that players 1 and 2 have subjective probabilities $2/3$ and $1/3$ respectively that a given horse A will run faster than another horse B in a given race. One could simply say that the subjective probabilities are different and leave it at that. But one could also imagine that both players previously had a uniform prior on the probability p that A beats B; that 1 had seen A beat B in one race, and 2 had seen B beat A in another race. Harsanyi's view is that differences in subjective probabilities can always be accounted for in such a fashion.²⁶

The holders of such a view would probably consider the approach of this paper invalid. Suppose that in Example 2.1, D is the event "horse A beats horse B." Suppose further that each player knows that the other has observed exactly one previous race. Now in our context we assume that each player knows the other's probability for D . But in that case, the very knowledge of the other player's probability will cause revision of each player's own probability. The result of this revision will necessarily be $p_1(D) = p_2(D) = 1/2$, so that the previously "subjective" probabilities become "objective."

But there is also another possibility, namely that each player has made several observations, and that these observations lead the players to assign probabilities $2/3$ and $1/3$ respectively to D . For example, $2/3$ would be the result of 3 wins for A and 1 win for B, as well as 1 win for A only. Though the players may know each other's probability, they may not know on precisely what observations it is based. Thus each player would not know if his own information is or is not more reliable than the other player's, and might be inclined to stick with his own information. Of course some revision of probabilities would certainly be called for even in such a case; but whether such a revision must always ultimately lead to equal probabilities for D is not clear.

In any event, under the Harsanyi doctrine cases of differing subjective probabilities known to all players either do not occur, or if they do occur, they should be treated by the methods of games of incomplete informa-

26. Though Harsanyi (1968) does not state this position in such absolute terms, there is little doubt that that is his belief.

tion. Basically, therefore, that part of this paper dealing with subjective randomization assumes that different people can have irreconcilable priors—precisely what Harsanyi calls the “inconsistent case.” Most workers in the field would probably agree that the inconsistent case can occur, and we are perfectly willing to let our contribution stand or fall on this basis.

But there is also another argument for our theory, an argument based on Savage’s “Small Worlds” theme (1954, §2.5 and §5.5). Suppose we are faced with a game like that in Example 2.1 or 2.3, and wish to apply the Harsanyi doctrine; this leads to an analysis by means of the theory of games of incomplete information. Now such an analysis will in general involve an enormous expansion of the formal description of the game. It may be necessary to use a population of many millions of types for each of the 2 or 3 players.²⁷ The resulting game will in practice be completely unanalyzable. We suggest that an equally valid practice would be to accept apparent differences in subjective probabilities at their face value, even if one adheres to the Harsanyi doctrine.

This question is not too different from the question “What is an outcome?” (cf. Savage, *op. cit.*). Suppose we are playing a given matrix game for money. The normal procedure would be to assign utilities to the amounts of money involved, and then solve the game using any of the standard theories. But a sum of money is not in itself valuable; it depends on how one wishes to use it. So to analyze the game properly we would have to decide on how to invest the money or what consumer article to buy with it. But even the investment or the consumer article itself is often only a means to an end; and this “end” in turn, is often again largely a means. Following through all this in a formal fashion, even if it were theoretically possible, would make the simplest of games totally unmanageable. So one isolates the given game as a “small world,” using the utilities to sum up all that follows. We suggest that even if in principle one accepts the Harsanyi doctrine, one can live with differing subjective probabilities as a summation of the complex informational situation in which the players find themselves.

k. We end this paper with a discussion of two possible objections that could be raised against the idea of subjective strategies. To fix ideas, consider the situation of two politicians, Adams and Brown, running for the office of mayor of their town. Each one has a number of (pure) campaign strategies open to him. Each pair of such strategies yields a proba-

27. . . . *every* possible combination of attributes . . . will be represented in this population . . . [Harsanyi (1968, p. 176) italics in the original].

bility for the election of Adams, the complementary probability being that of the election of Brown; for simplicity, assume that these probabilities are objective. This is a classical example of a constant-sum game, since there are only two possible final outcomes; the utilities of Adams and Brown may thus be taken equal to the respective success probabilities, and the sum is therefore always constant ($= 1$).

What is now being suggested is that Adams and Brown use randomized campaign strategies that are pegged on events for which they have different subjective probabilities. These events may be entirely disconnected with the political campaign in question; to take an extreme example, if the candidates are Americans, they may peg their choices on the outcome of a cricket match in England.

One objection that may be raised to this procedure is that whereas a prudent person might be willing to use mixed strategies based on a coin toss with known objective probabilities, he would hesitate to risk his career on the outcome of an event of which he knows little or nothing. Adams and/or Brown may know little or nothing about cricket in England; how can we suggest that they peg their decisions on it?

Though this objection has great intuitive force, it is not consistent with the theory of subjective probability. According to this theory, people have subjective probabilities for any event, no matter how well or ill they are informed about it; of course the probability will depend on the available information, but there will always *be* a probability. Once this probability is determined, it enters the decision-making process in all respects like an objective probability. A certain amount of introspection will, moreover, support the conclusions of the theory despite their apparent strangeness. Suppose Adams is asked to choose between winning the election with objective probability $1/3$ and winning the election if the Manchester cricket team beats the Liverpool team. If he knows nothing at all about the cricket situation, and subscribes to the “principle of insufficient reason,” he might well choose the second possibility, despite—or rather because of—his lack of information. The same point may similarly be made by considering the situation in which it is not a pre-determined coin that is tossed, but rather a biased coin that is picked out of a hat with a wide distribution of biases; in such a situation, the utility theory of von Neumann and Morgenstern (1953) still applies, although in a certain sense the user knows much less about the coin.

1. Another possible objection runs as follows: Both players realize that the game is basically zero-sum, since in the last analysis, only one of them can win the election. Both know it is an illusion to think that both can gain more than the value of the game; this illusion is based on what

might be called a mistaken appraisal of probabilities by at least one of them (and possibly both). Would they not be more prudent, then, to guarantee to themselves the objective value of this game by the use of objectively randomized strategies, thus avoiding all possible “mistakes”?

On one level, this question can be answered simply by saying that each side is here taking advantage of what it sees as the other side’s mistake, and that this is perfectly rational. On a deeper level, though, it is incorrect to speak of “mistakes” at all. Each side has well-defined systems of preferences, each perfectly consistent in itself, and these preferences are *not* directly opposed. Thus, it is incorrect to say that the situation is “basically zero sum.” This situation is analogous to a situation in which the players’ preferences as between *pure* outcomes are directly opposed, though between mixtures of outcomes they may not be; such a game is not usually called zero-sum or even “strictly competitive.”

Consider, for example, the game

	<i>L</i>	<i>R</i>
<i>T</i>	5, 0	2, 2
<i>B</i>	1, 3	0, 5

Preference-wise, this matrix has a saddle point at (T, R) . It is, however, not constant sum, since a contract to play (T, L) and (B, R) with probability $1/2$ each is preferred by both players to the saddle point. Rather than signing the contract before tossing a coin, one could imagine that the coin has already been tossed, but that the players are not informed of the outcome. Paralleling the above argument, one could then say that both players realize that the game is “basically strictly competitive,” since the preferences as between pure outcomes are directly opposed. Thus both know that it is an illusion to think that both can gain more than the 2 units assigned at the saddle point; this illusion is based on a mistaken appraisal of how the coin actually fell. Would they not be more prudent to guarantee a sure payoff of at least 2?

Certainly if they knew how the coin fell, there would be no point in using it to mix outcomes; but in the current situation, they both prefer the $1/2$ - $1/2$ combination of (T, L) and (B, R) to the saddle point (T, R) . This agreement is therefore a perfectly natural one in spite of the fact that one of the players must necessarily lose out. Completely similar reasoning holds in the case of a game that is zero-sum in utilities, but that can be made non-strictly competitive by the appropriate use of subjectively randomized strategies.

Acknowledgement

Research supported in part by National Science Foundation grant GS-3269 at the Institute for Mathematical Studies in the Social Sciences, Stanford University, and in part by CORE, the Center for Operations Research and Econometrics at the Catholic University of Louvain.

References

- Anscombe, F. J. and R. J. Aumann, 1963, A definition of subjective probability, *Ann. Math. Stat.* 34, 210–212 [Chapter 17].
- Aumann, R. J., 1961, Almost strictly competitive games, *J. Soc. Ind. Appl. Math.* 9, 544–550 [Chapter 27].
- Aumann, R. J., 1964, Mixed and behavior strategies in infinite extensive games, in: M. Dresher, L. S. Shapley and A. W. Tucker, eds., *Advances in game theory* (Princeton University Press, Princeton) 627–650 [Chapter 28].
- Aumann, R. J. and M. Maschler, 1972, Some thoughts on the minimax principle, *Management Science* 18, P-54–P-63 [Chapter 29].
- Harsanyi, J., 1967–68, Games of incomplete information played by Bayesian players, Parts I–III, *Management Science* 14, 159–182, 320–334, 486–502.
- Harsanyi, J. and R. Selten, 1972, A generalized Nash solution for two-person bargaining games with incomplete information, *Management Science* 18, P-80–P-106.
- Kuhn, H. W., 1953, Extensive games and the problem of information, in: H. W. Kuhn and A. W. Tucker, eds., *Contributions to the theory of games II* (Princeton University Press, Princeton) 193–216.
- Lyapunov, A., 1940, Sur les fonctions-vecteurs completement additives, *Bull. Acad. Sci. USSR Ser. Math.* 4, 465–478.
- Nash, J. F., 1951, Non-cooperative games, *Ann. of Math.* 54, 286–295.
- Savage, L. J., 1954, *The foundations of statistics* (John Wiley, New York).
- Von Neumann, J. and O. Morgenstern, 1944, *Theory of games and economic behavior* (Princeton University Press, Princeton); 3rd ed., 1953.