

Risk Neutral Equilibria of Noncooperative Games

Robert Nau

Fuqua School of Business
Duke University
Durham, NC 27708 USA
robert.nau@duke.edu
October 8, 2013

Abstract: Game-theoretic solution concepts such as Nash and Bayesian equilibrium start from an assumption that the players' sets of possible payoffs, measured in units of von Neumann-Morgenstern utility, are common knowledge, and they go on to define rational behavior in terms of equilibrium strategy profiles that are either pure or independently randomized and which, in applications, are often taken to be uniquely determined or at least tightly constrained. A mechanism through which to obtain common knowledge of payoff functions measured in units of utility (or common priors over pre-determined sets of such functions) is not part of the model. This paper describes an operational method of constructing a state of common knowledge of the key parameters of the players' utility functions in terms of conditional small bets on the game's outcome. When the rationality criterion of joint coherence (no arbitrage) is applied in this setting, the solution of a game is typically characterized by a convex set of correlated equilibria. In the most general case, where players are risk averse, the parameters of the equilibria are risk neutral probabilities, interpretable as products of subjective probabilities and relative marginal utilities for money, as in financial markets. Risk aversion generally enlarges the set of equilibria and may present opportunities for Pareto-improving modifications of the rules of the game.

Keywords: Nash equilibrium, correlated equilibrium, coherent previsions, subjective probabilities, risk neutral probabilities, lower and upper probabilities, asset pricing, arbitrage, matching pennies

Forthcoming in *Theory and Decision*, © Springer Science+Business Media New York, 2013

1. Introduction

Game theory occupies the large middle ground of rational choice theory: the problem of “2, 3, 4... bodies” in which agents must reason about the strategic behavior of other rational agents as well as reflect on their own preferences and compete in markets. The modeling of interactive decisions of this kind requires strong assumptions. First, the rules of the game are (in the most general case) parameterized in units of von Neumann-Morgenstern utility rather than money or material goods in order to allow for differences in tastes and attitudes toward risk. Second, the utility payoffs of different players are assumed to be common knowledge, enabling them to model each other's decisions as well as their own, and to all know that they can all do this, and so on. Third, common knowledge of rationality and common knowledge of the rules of the game are assumed to lead to an equilibrium, usually a Nash equilibrium or one of its refinements or extensions, in which the decision of each player is individually rational given the decisions simultaneously made by the other players, and randomization (if any) is performed independently. And fourth, when there is uncertainty about any of the game parameters, the beliefs of the players are assumed to be consistent with a common prior distribution, which generates an infinite hierarchy of mutually consistent reciprocal beliefs. These assumptions are often applied at maximum strength in order to tightly constrain the solution, yet each of them is problematic in its own way.

Common knowledge of real-valued von Neumann-Morgenstern utilities among potentially risk averse players is an idea that is rather tricky to make precise, let alone credible, except in simple settings where it is true by construction. In textbooks on game theory (e.g., Fudenberg and Tirole 1991, p. 4) an informal analogy may be drawn in which players are seated at separate computer terminals where they all see each other's utility payoffs displayed. Formal

treatments of common knowledge generally refer to knowledge of the occurrence of events, following Aumann (1976), and where this concept is applied to common knowledge of real-valued payoffs or reciprocal uncertainty about payoffs, it is done in terms of events that select elements from a set of possible payoff functions whose parameters are already given. In experimental games, common knowledge of utilities may be induced by having the players read from a common rulebook in which their payoffs are specified in units of the probability of winning a single large prize (e.g., Cooper et al. 1989, Ochs 1995). By definition this yields commonly known utilities, but the players' risk attitudes are rendered irrelevant by the fact that there are only two terminal outcomes. A modeling approach that gives a behavioral role to risk aversion in simple games is the concept of quantal response equilibrium (McKelvey and Palfrey 1995) in which the players' behavior has a noise component that is distinct from randomization in the usual sense of deliberately mixed strategies. Goeree et al. (2003) show that when this solution concept is applied to generalized matching-pennies games in the presence of risk aversion, it leads players to prefer moves that are safe in the sense of having payoffs that exhibit little variation in response to the moves of others, which is incompatible with Nash equilibrium behavior in that setting.

This paper asks a more fundamental question, namely: how does risk aversion affect the amount of information about the rules of a game that can be made commonly known via material communication, and what are its implications for the determinacy of the game's solution, even in a simple case such as matching pennies? The modeling approach follows that of Nau and McCardle (1990) and Nau (1992), which is a multi-player extension of de Finetti's (1937, 1974)

operational approach to defining subjective probabilities,¹ which in turn is a microcosm of a financial market. Its behavioral primitives are offers to accept small conditional bets whose payoffs depend on the outcome of the game. This leads to solutions of games that are characterized by exactly the same rationality conditions as those of individual decisions and competitive markets, and they are not necessarily uniquely determined nor probabilistically independent between players: they may consist of convex sets of correlated equilibria rather than unique Nash equilibria. It will be shown here that in the case of risk averse players the common priors and equilibria are expressed in terms of sets of “risk neutral probabilities” that need not represent the players’ true subjective beliefs. In strictly competitive games such as matching pennies, players may be able to hedge some of their risks and alter the game’s rules to their mutual benefit by additionally accepting finite bets that partially reveal the solution of the game, i.e., bets that are rationalized by their present beliefs in addition to the bets that are rationalized by the game’s rules.

2. Coherent lower and upper previsions

Virtually all of the fundamental theorems of rational choice theory—subjective probability, expected utility, subjective expected utility, asset pricing, welfare economics, cardinal utilitarianism, *and* noncooperative games—are duality theorems that can be proved by a separating hyperplane argument. In the versions of these theorems that involve finite sets of states and/or consequences, it is a variant of Farkas’ lemma, the basis of the duality theorem of linear programming:

¹ Lad (1996) provides a comprehensive treatment of de Finetti’s theory of probability and subjective statistical methods that are based on it.

LEMMA 1: For any matrix G , there exists a non-negative vector α such that $\alpha \square G < \mathbf{0}$ or else there exists a non-negative vector π such that $G \pi \geq \mathbf{0}$, $\pi \neq \mathbf{0}$.

LEMMA 2: For any matrix G , there exists a non-negative α such that $\alpha \square G \leq \mathbf{0}$ and $[\alpha \square G]_k < 0$ or else there exists a non-negative vector π , with $\pi_k > 0$, such that $G \pi \geq \mathbf{0}$.

De Finetti's (1974) fundamental theorem of probability provides a rationale for the Bayesian viewpoint as it applies to personal beliefs. Its general form, which is stated in terms of lower and upper probabilities and expectations, can be proved as follows using the language of financial markets. Consider an agent ("she") who is uncertain about which element of a finite set S of states of the world will occur. Let N denote the number of states and let \mathbf{x} denote an *asset*, which is an N -vector of payoffs assigned to states. The agent's *lower prevision for* \mathbf{x} is the price $\underline{P}(\mathbf{x})$ that she is publicly willing to pay per unit of \mathbf{x} in arbitrary (but small) quantities chosen by someone else. This means that for any small positive number α chosen by an observer ("he") the agent will accept a bet whose payoff vector for her is $\alpha(\mathbf{x} - \underline{P}(\mathbf{x}))$, with the opposite payoffs to the observer.² For example, if $N = 3$, $\mathbf{x} = (30, 10, -20)$, and $\underline{P}(\mathbf{x}) = 14$, the agent will accept a bet whose payoff vector for her is $(16\alpha, -4\alpha, -34\alpha)$ for any small positive α chosen by the observer.

² Notational conventions: Lower-case boldface letters such as \mathbf{x} and \mathbf{e} are used interchangeably for payoff vectors of assets and indicator vectors of events as well as for their proper names (e.g., "event \mathbf{e} " is the event whose indicator vector is \mathbf{e}). In the expression $\alpha(\mathbf{x} - \underline{P}(\mathbf{x}))$, \mathbf{x} is a vector and α and $\underline{P}(\mathbf{x})$ are scalars, and the multiplication and subtraction are performed pointwise, yielding a vector whose n^{th} element is $\alpha(x_n - \underline{P}(\mathbf{x}))$. If \mathbf{x} and \mathbf{y} are vectors of the same length, then $\mathbf{x}\mathbf{y}$ denotes their pointwise product (another vector of the same length), and $\mathbf{x} \square \mathbf{y}$ denotes their inner product (a scalar). If G is a matrix and \mathbf{x} and \mathbf{y} are vectors of appropriate length, then $\mathbf{x} \square G$ and $G \mathbf{y}$ denote matrix multiplication of G by \mathbf{x} on the left or by \mathbf{y} on the right, yielding vectors. If π is a probability distribution on states and \mathbf{x} is a payoff vector and \mathbf{e} is an indicator vector for an event, then $P_\pi(\mathbf{x})$ is the corresponding expected value of \mathbf{x} and $P_\pi(\mathbf{e})$ is the probability of \mathbf{e} , i.e. $P_\pi(\mathbf{x}) = \pi \square \mathbf{x}$ and $P_\pi(\mathbf{e}) = \pi \square \mathbf{e}$. $P_\pi(\mathbf{x}|\mathbf{e})$ denotes the conditional expectation of \mathbf{x} given the occurrence of \mathbf{e} that is determined by π , i.e. $P_\pi(\mathbf{x}|\mathbf{e}) = P_\pi(\mathbf{x}\mathbf{e})/P_\pi(\mathbf{e})$ provided that $P_\pi(\mathbf{e}) > 0$.

A lower prevision for an asset may be considered as a *lower expectation*, i.e., a lower bound on its subjective expected value for the agent. In the special case where \mathbf{x} is a binary vector that is the indicator of an event, its lower prevision is a *lower probability* for the event.

Lower previsions can also be assessed conditionally. If \mathbf{x} is the payoff vector of an asset and \mathbf{e} is the indicator vector of an event, the agent's *conditional lower prevision for \mathbf{x} given \mathbf{e}* is the price $\underline{P}(\mathbf{x}|\mathbf{e})$ that she is publicly willing to pay per unit of \mathbf{x} in arbitrary small multiples chosen by an observer, subject to the condition that the deal will be called off if \mathbf{e} does not occur. This means that the agent will agree to accept a bet whose payoff vector for her is $\alpha(\mathbf{x} - \underline{P}(\mathbf{x}|\mathbf{e}))\mathbf{e}$, for any small positive α . To continue the previous example, if $\mathbf{e} = (1, 1, 0)$, which is the indicator vector for the event in which either state 1 or state 2 occurs, and $\underline{P}(\mathbf{x}|\mathbf{e}) = 21$, then the agent will accept a bet whose payoff vector for her is $(9\alpha, -11\alpha, 0)$. In the special case where $\underline{P}(\mathbf{x}|\mathbf{e}) = 0$, the agent is willing to pay zero for \mathbf{x} conditional on \mathbf{e} , which means she will accept a small bet whose payoff vector is proportional to \mathbf{x} conditional on the occurrence of \mathbf{e} . This is equivalent to an unconditional bet with payoffs proportional to $\mathbf{x}\mathbf{e}$.

It remains to show that rational lower previsions satisfy the laws that ought to be satisfied by lower bounds on probabilities and expectations. Suppose that the agent assigns a conditional lower prevision $\underline{P}(\mathbf{x}_m|\mathbf{e}_m)$ to asset \mathbf{x}_m given the occurrence of event \mathbf{e}_m , and the observer chooses a small quantity α_m to sell her at that price, for $m = 1, \dots, M$, subject to the condition that the transactions are additive, which is the way a bookmaker or financial market normally operates. The agent is rational *ex ante* if her previsions do not expose her to arbitrage, i.e., if the observer is not able to earn a profit in every state through a clever combination of sales to her at her posted prices. The agent is rational *ex post* in state k if her previsions do not allow the observer

to earn a profit if state k occurs, without any exposure to loss. These rationality conditions are called *coherence* and *ex post coherence*, respectively. More precisely:

DEFINITION: The conditional lower previsions $\{\underline{P}(x_1|e_1), \dots, \underline{P}(x_M|e_M)\}$ are *coherent* if there do not exist non-negative numbers $\{\alpha_1, \dots, \alpha_M\}$ such that $\sum_{m=1}^M \alpha_m (x_{mn} - \underline{P}(x_m | e_m)) e_{mn} < 0 \forall n$. They are *ex post coherent in state k* if there do not exist non-negative numbers $\{\alpha_1, \dots, \alpha_M\}$ such that $\sum_{m=1}^M \alpha_m (x_{mn} - \underline{P}(x_m | e_m)) e_{mn} \leq 0 \forall n$ with strict inequality when $n = k$.

Coherence entails ex post coherence in at least one state. In these terms, de Finetti's fundamental theorem is as follows:

THEOREM 1: The conditional lower previsions $\{\underline{P}(x_1|e_1), \dots, \underline{P}(x_M|e_M)\}$ are coherent [ex post coherent in state k] if and only if there exists a non-empty convex set Π of probability distributions on states of the world [satisfying $\pi_k > 0$] such that, for all m and all $\pi \in \Pi$, $P_\pi(x_m|e_m) \geq \underline{P}(x_m|e_m)$ or else $P_\pi(e_m) = 0$.³

Proof: The coherence [ex post coherence] result follows from Lemma 1 [2] by letting G be the matrix whose mn^{th} element is $(x_{mn} - \underline{P}(x_m|e_m))e_{mn}$.

Coherent lower previsions therefore have the properties of lower probabilities and expectations determined by a convex set of probability distributions, which can be interpreted to represent the possibly-imprecise beliefs of the agent, assuming that she has linear utility for money—an assumption that will be relaxed later on.

³ The necessary and sufficient conditions for coherence [ex post coherence] follow directly from Lemma 1 [2]. De Finetti (1937) separately derived the additive and multiplicative laws of probability from the requirement of coherence by evaluating the determinants of matrices whose rows are the payoff vectors of acceptable bets. A proof of the more general statement of the theorem via a separating hyperplane method was given by Smith (1961). The ex post variant of the theorem is discussed by Nau (1995).

An under-appreciated property of de Finetti's operational definition of subjective probabilities and expectations is that it does not merely define them: it makes them common knowledge in the everyday specular sense of the term. The prices are visible to both actors in the scene, and the actors both know it, and both know that they both know it, and so on, and the meaning of the numbers is commonly understood by virtue of the opportunities that they create for reciprocal financial transactions.⁴ This is a property of posted prices in general. They do not merely facilitate trade and market efficiency: they are also credible and commonly known numerical measurements of the comparative beliefs and values of those who post them.

It might be argued that game-theoretic techniques should be used to address the question of why and how the agent should offer distinct bid and ask prices in her interaction with the observer, or whether she should offer to bet at all. There might be asymmetric information or incentives for secrecy or deception or speculation that would motivate the agent to set her bid prices for assets at levels other than her true lower bounds on their expected payoffs, whatever "true" might mean. This would merely beg the question of how the rules of the higher-order game would come to be commonly known in numerical terms. If an infinite regress is to be avoided, then at some level of description the amount of private information about her beliefs and values that an agent is willing to publicly reveal is a behavioral primitive. In what follows, the game-theoretic argument will be turned on its head: the fundamental theorem of noncooperative games is merely an extension of the fundamental theorem of probability to multiple actors in the same scene.

⁴ Proper scoring rules (Savage 1971, de Finetti 1974) provide an alternative operational method of defining and eliciting subjective probabilities, in which the forecaster is rewarded for predictive accuracy according to a fixed formula (e.g. a quadratic or logarithmic function) and the coherence criterion enforces the laws of probability via marginal analysis. Probabilities elicited by this method are arguably common knowledge between the probability assessor and whomever is responsible for paying her an amount of money equal to her score.

3. Correlated equilibrium as the game-theoretic expression of coherent previsions

In the assessment of previsions via offers to bet, there is no requirement that states of the world should be events that are beyond anyone's control. However, an observer might be reluctant to take the other side of any bet whose payoff depends on an event that they both know the agent *does* control, and conversely the agent might be reluctant to offer to bet on events that she knows to be controlled by others. An important special case is one in which the state space can be partitioned as $\mathcal{S} = \mathcal{S}_1 \times \mathcal{S}_2$, where \mathcal{S}_1 is a set of events that the agent controls (her own moves) while \mathcal{S}_2 is a set of events outside her control (moves of nature or other agents). If e is an event that is measurable with respect to \mathcal{S}_1 and \mathbf{x} is the payoff vector of an asset that is measurable with respect to \mathcal{S}_2 , it might be reasonable for the agent to assert a lower prevision for \mathbf{x} conditional on e . If she asserts that $\underline{P}(\mathbf{x}|e) = 0$, it means that she will accept a small bet whose payoff vector is proportional to \mathbf{x} under the same conditions in which she would choose the move e , or equivalently, she will accept a small bet whose payoff vector is proportional to $\mathbf{x}e$. Such a bet reveals some information about the agent's payoff function in the game she is playing against nature or her adversaries without necessarily revealing the move she intends to make. Namely, her payoffs in the game are such that her best move is e only under conditions where her prevision for \mathbf{x} is non-negative. This method for revealing limited information about one's payoff function yields enough detail about the rules of a noncooperative game to determine its equilibria, as will be shown next.

Let \mathcal{G} denote a noncooperative game among I players, each having a finite set of strategies. Let $\mathcal{S} = \mathcal{S}_1 \times \dots \times \mathcal{S}_I$ denote the set of outcomes, where \mathcal{S}_i is the set of index numbers for strategies of player i . Let $s = (s_1, \dots, s_I)$ denote a generic outcome, in which s_i is the strategy chosen by player i . Let \mathbf{x}_i denote the payoff function (an $|\mathcal{S}|$ -dimensional vector) for player i ,

whose value in outcome s is $x_i(s)$. Assume that payoffs are measured in units of money and that the players are risk neutral. (The risk neutrality assumption will be relaxed later.) The true game \mathcal{G} is thus defined by the sets of strategies $\{S_i\}$ and payoff vectors $\{x_i\}$.

Let e_{ij} denote the event in which player i plays her j^{th} strategy, and for every $j \in S_i$, let x_{ij} denote a vector of payoffs that is obtained from x_i as follows: $x_{ij}(s) = x_i(s_1, \dots, j, \dots, s_N)$, where the j occurs in the i^{th} position. In other words, $x_{ij}(s)$ is the payoff that player i receives by playing her j^{th} strategy while the rest play according to s . (There is some duplication of information in the structure of x_{ij} : it contains multiple copies of the payoff profile that player i obtains by playing j , because $x_{ij}(s_1, \dots, s_i, \dots, s_N)$ is the same for all values of s_i .) Suppose that the payoff functions $\{x_i\}$ are not commonly known *a priori* and must therefore be revealed through some credible language of communication. The language that will be used here is the same one that was sketched in the previous section. To see how it works in the game, observe that in the event that player i chooses her j^{th} strategy, she must weakly prefer the profile of payoffs she gets by playing strategy j to the profile of payoffs she would have gotten by playing any other strategy k . In the terms introduced above, she evidently prefers x_{ij} over x_{ik} in the event that e_{ij} occurs, which means that she would trade x_{ik} for x_{ij} conditional on e_{ij} . Such a trade is equivalent to an unconditional bet with a payoff vector of $(x_{ij} - x_{ik})e_{ij}$. If the agent wants to let this information about her payoff function be common knowledge, she can publicly offer to accept a small bet whose payoff vector is proportional to $(x_{ij} - x_{ik})e_{ij}$ at the discretion of an observer. Or, to turn the story around, if by some means her payoff function x_i is already common knowledge, then it is also common knowledge that she will accept such a bet.⁵ Note that she is not betting directly on

⁵ Strictly speaking, the *choice* of strategy j in the presence of k can only be interpreted to mean a *preference* for j over k if the agent has complete preferences, requiring precise beliefs. Here, offers to bet are assumed to occur at a

her own strategy. Rather, her own strategy is used as a conditioning event for bets on what other players will do. Bets that are conditioned on the player's own strategy, which may be uncertain to the observer and the other players, do not necessarily reveal her actual state of information or her intended move.

Suppose that all the players offer to accept small conditional bets that are determined by their true payoff functions in the manner described above. Let \mathbf{G} denote the matrix whose columns are indexed by outcomes of the game, whose rows are indexed by ijk , and whose ijk^{th} row is $(x_{ij} - x_{ik})e_{ij}$, the payoff vector of the bet that is acceptable to player i in the event that she chooses strategy j in preference to strategy k . Then, under the assumption that such bets may be non-negatively linearly combined, an observer of the game may choose a non-negative vector of multipliers α to construct an acceptable bet that yields a total payoff vector of $\alpha \square \mathbf{G}$ to the players and the opposite payoffs to himself.

The matrix \mathbf{G} will be henceforth called the *revealed rules of the game* because, as will be shown, it contains all the commonly-knowable information about the rules that is actually used in determining noncooperative equilibria. However, \mathbf{G} does not contain all the information about the true game \mathcal{G} that is materially important to the players. In particular, it does not reveal the benefits that a given player might obtain from changes in the strategies of the other players, holding her own strategy fixed. The latter information is subtracted out when the calculation $(x_{ij} - x_{ik})e_{ij}$ is performed. All that remains is information about how a given player would benefit by changing her own strategy, holding the strategies of the *other* players fixed. This is the

point in time when the agents may not yet have formed precise beliefs about what their opponents will do, but they expect that they will have done so by the time they are called upon to move. In the meantime they are making assertions about constraints that precise beliefs would have to satisfy in order for them to prefer one strategy over another, thereby partially revealing their payoff functions.

essence of noncooperative game-playing. The players do not consider the implications of their own play for the payoffs of other players, nor do they expect the other players to show that consideration to them.

Under the assumptions given above, we can define what it means for the game to be played rationally by applying the concept of ex post coherence jointly to all the players. Consider an observer who knows nothing about the game except the bets that the players have offered, which is the minimal information about the game's rules that is common knowledge. Suppose that he does not want to speculate on the game's outcome, but he would like to make a riskless profit if possible. From the observer's perspective, if several bets are placed on the same table at the same time, it doesn't matter whether they are offered by one individual or by many who are all looking each other in the eye. If the observer manages to pick their pockets, the players have behaved irrationally.

DEFINITION: The strategy s is *jointly coherent* if there does *not* exist a non-negative α such that $\alpha \cdot G \leq 0$ and $[\alpha \cdot G](s) < 0$, i.e., if, under the revealed rules of the game, there is no system of system of bets under which the observer cannot lose and will win a positive amount from the players if they play s .

Fortunately for the players, there is always at least one jointly coherent strategy: they are not doomed to exploitation if they honestly reveal some information about their payoff functions.⁶ The interesting question is whether there are strategies that are *not* jointly coherent, and if so, how they are characterized.

⁶ A proof of this result is given by Nau and McCardle (1990). A proof of the dual condition, which (by Theorem 2) is the existence of a correlated equilibrium, is given by Hart and Schmeidler (1989). These proofs are more elementary than the proof of existence of a Nash equilibrium insofar as they do not invoke a fixed-point theorem. In Nau and McCardle's proof, the result follows from the existence of a stationary distribution of a Markov chain.

In general, the players might choose either pure or randomized strategies, and randomized strategies might be either independent or correlated. Correlated randomization of strategies could be carried out with the help of a mediator but does not necessarily require it: flipping a coin or drawing straws or playing paper-scissors-rock are familiar correlation devices that do not require a mediator, and a taking-turns convention in repeated play could be viewed as a correlation device from the perspective of an observer who doesn't know whose turn it is. Let π denote a (possibly-degenerate) probability distribution over the outcomes of the game, and suppose that the players *do* employ a mediator who is instructed to randomly draw a joint strategy s according to the distribution π and then privately recommend to each player that she should play her own part of it. Thus, player i hears only her own recommended strategy, s_i , not those of the other players. Under these conditions, π is a common prior distribution over recommended joint strategies in the game, and each player can use Bayesian updating to compute a posterior distribution for the recommendations that were received by the other players, given her own recommendation. If each player's recommended strategy is optimal for her *a posteriori* when the others play their own recommended strategies, then π is a correlated equilibrium of the game (Aumann 1974, 1987). More precisely:

DEFINITION: π is a *correlated equilibrium* of \mathcal{G} if $G\pi \geq \mathbf{0}$, which means that for every player i and every recommended strategy j and alternative strategy k of that player, either $P_\pi(e_{ij}) = 0$ (the probability of strategy j being recommended to player i is zero) or else $P_\pi(x_{ij} - x_{ik} | e_{ij}) \geq 0$ (the conditional expected payoff of strategy j is greater than or equal to the conditional expected payoff of strategy k when j is recommended, assuming that other players adhere to their own recommendations).

Because the set of all correlated equilibria of \mathcal{G} is determined by a system of linear inequalities, it is a convex polytope—a tractable geometrical object—which will henceforth be denoted by $\square_{\mathcal{G}}$. A *Nash equilibrium* is a special case of a correlated equilibrium in which π is independent between players, allowing each player to perform her own randomization (if necessary) without a mediator. The set of Nash equilibria is not necessarily convex or connected or bounded by points with rational coordinates, and it can be rather difficult to compute, particularly in games with more than 2 players. Nash equilibria cannot lie just anywhere in the polytope of correlated equilibria: they are always on its boundary, i.e., on supporting hyperplanes, although not necessarily at extreme points (Nau et al. 2004, Viossat 2006, 2010).

Aumann (1987) has asserted that correlated equilibrium, rather than Nash equilibrium, is the natural expression of Bayesian rationality in noncooperative game theory. And indeed, in the terms given above, we can prove a fundamental theorem of games that is a strategic variant of the fundamental theorem of probability and leads straight to correlated equilibrium. The theorem and its proof are merely a *restatement* of the fundamental theorem of probability and its proof (the ex post version) for the special case in which conditional previsions are jointly announced by two or more individuals and the assets and conditioning events to which they refer have a special structure that is determined by a noncooperative game they are playing.

THEOREM 2 (Nau and McCardle 1990): In a game among risk neutral players, a strategy is jointly coherent if and only if there exists a correlated equilibrium in which it has positive probability.

Proof: By Lemma 2, either there exists $\alpha \geq \mathbf{0}$ such that $\alpha \square G \leq \mathbf{0}$ and $[\alpha \square G](s) < 0$ or else there exists $\pi \geq \mathbf{0}$, with $\pi(s) > 0$, such that $G \pi \geq \mathbf{0}$. \square

Hence, the players are rational ex post if and only if they behave as if they had implemented a correlated equilibrium, i.e., if they play a strategy that could have occurred with positive probability in such an equilibrium.⁷ But even more can be said: lower and upper bounds can be placed on the players' jointly-held previsions for outcomes of the game and any side bets that might be attached to it, namely the bounds that are determined by the convex polytope $\square_{\mathcal{G}}$ of correlated equilibria. On this basis it is appropriate to consider $\square_{\mathcal{G}}$ to be the rational solution of the game when it is played noncooperatively in the absence of any constraints other than coherence, and in general it is a solution in terms of imprecise probabilities.⁸

4. Risk aversion and risk neutral probabilities

The results of the preceding section require the players to be risk neutral, i.e., to have state-independent linear utility for money. The more general case of risk averse players will now be considered, and it will be shown that risk aversion leads them to hedge their bets, making the revealed set of equilibria larger than it would have been otherwise. Furthermore, when players are risk averse, side bets may provide opportunities for Pareto-improving modifications of the rules of the game, which blurs the distinction between strategic and competitive equilibria. In extreme cases, players may be able to hedge their positions so as to decouple their payoff

⁷ In games of incomplete information, joint coherence leads similarly to a correlated generalization of Bayesian equilibrium (Nau 1992). Correlated equilibria of games of incomplete information have been extensively studied by Forges (1986, 1993, 2006).

⁸ This approach can be generalized to the situation in which players do not exactly know their own payoffs. If each payoff in the game matrix is known by its recipient only to lie within some interval, then the ijk^{th} row of \mathbf{G} becomes $(\mathbf{x}_{ij}^{\max} - \mathbf{x}_{ik}^{\min})\mathbf{e}_{ij}$, where \mathbf{x}_{ij}^{\max} and \mathbf{x}_{ik}^{\min} are pointwise maxima and minima of the possible payoffs of strategies j and k for player i . This means that in the event that player i chooses strategy j over strategy k , the minimal requirement that her conditional beliefs must satisfy is that her best possible lower prevision for the payoff of j should be at least as great as her worst possible lower prevision for the payoff of k .

functions and exit from the game altogether. To set the stage, some general remarks on the modeling of risk aversion are appropriate.

If the agent is risk averse and has substantial prior stakes in events, i.e., background risk, then Theorem 1 still holds, but its parameters have a different interpretation. Suppose that she has subjective expected utility preferences and her risk attitude is represented by a strictly concave von Neumann-Morgenstern utility function $U(x)$, with its derivative denoted by $U'(x)$, and suppose that her background risk is represented by a payoff vector \mathbf{z} whose elements differ across states by amounts that are large enough to cause substantial variations in the marginal utility of money. Then her acceptance of an additional small bet \mathbf{x} will not be based on its expected value but rather on its expected marginal utility in the context of \mathbf{z} . If the agent's true beliefs are represented by a subjective probability distribution \mathbf{p} , then her status quo expected utility is $E_p[U(\mathbf{z})]$. A bet \mathbf{x} will be acceptable to her if it maintains or increases her expected utility, i.e., if $E_p[U(\mathbf{z}+\mathbf{x})] - E_p[U(\mathbf{z})] \geq 0$.

If the elements of \mathbf{x} are small enough in magnitude so that only first-order effects are important, then \mathbf{x} is acceptable if $E_p[U'(\mathbf{z})\mathbf{x}] \geq 0$, or equivalently if $E_\pi[\mathbf{x}] \geq 0$, where π is a probability distribution obtained by multiplying the true probability distribution \mathbf{p} pointwise by the marginal utility vector $U'(\mathbf{z})$ and then re-normalizing, so that $\pi(\mathbf{s}) \propto p(\mathbf{s})U'(\mathbf{z}(\mathbf{s}))$. This is the *risk neutral probability distribution of the agent* at \mathbf{z} , because she evaluates small bets in a seemingly risk neutral way using π rather than her true subjective probability distribution \mathbf{p} . The agent's risk neutral distribution depends on her beliefs, her background risk, and her attitude

toward it, and the effects of these three conceptually distinct factors may be inseparable from the viewpoint of an observer who bets with her, but this is not necessarily a problem.⁹

In a financial market, the necessary and sufficient condition for asset prices to create no arbitrage opportunities is that there should exist a probability distribution under which every asset's expected payoff, discounted at the risk-free rate of interest if time is a factor, lies between its bid and ask prices. This result is known as the fundamental theorem of asset pricing, and it is merely de Finetti's fundamental theorem of probability applied to asset prices offered simultaneously by possibly-different individuals. The probability distribution that prices the assets is called the *risk neutral probability distribution of the market*, because it prices them in a seemingly risk neutral way, and it can be determined from prices of options or Arrow securities.¹⁰ Because of friction and incompleteness, the market's risk neutral distribution is not necessarily unique. Rather, there may be a convex set of risk neutral distributions determined by bid and ask prices for assets.

In equilibrium, the marginal prices that agents are willing to pay for financial assets must agree with market prices, which means that the risk neutral probability distributions of all the agents must agree with the risk neutral probability distribution of the market. More precisely,

⁹ For example, phenomena such as ambiguity aversion as well as risk aversion can modeled and measured in terms of properties of the agent's risk neutral probabilities, even when her beliefs and tastes cannot be uniquely separated (Nau 2001, 2003, 2006, 2011).

¹⁰ The literature on risk neutral probabilities and asset pricing by arbitrage traces back to the work of Black and Scholes, Merton, Cox, Ross, Rubinstein, and many others in the 1970's, although the connection with de Finetti's use of the no-arbitrage principle in subjective probability, dating to the 1930's, was not noticed until later. Even today, where de Finetti's work is mentioned at all in the literature of financial economics, it is usually in connection with having anticipated Markowitz's mean-variance portfolio theory and the Arrow-Pratt measure of risk aversion. Observations of the fact that de Finetti's coherence principle for subjective probabilities is the same as the no-arbitrage principle for contingent claim markets and that his fundamental theorem of probability is the same as the fundamental theorem of asset pricing are found instead in the applied mathematics literature, e.g., Ellerman (1984).

the set of risk neutral distributions that is determined by bid and ask prices in the market is the intersection of all the sets of risk neutral distributions that are determined by bid and ask prices of individuals, which is non-empty if and only if there are no arbitrage opportunities. Thus, rational behavior in markets requires the agents to agree on risk neutral probabilities to the extent that their sets of personal risk neutral probabilities must have at least one point in common. In the special case where the agents have complete preferences and the market is also complete and frictionless, the risk neutral probabilities of the agents and the market are uniquely determined and must be identical.

5. Risk neutral equilibria

When agents are risk averse with significant prior stakes in events, their lower and upper previsions determined by offers to accept small bets must be interpreted as lower and upper expectations with respect to convex sets of risk neutral probabilities, rather than true subjective probabilities, as discussed above. The same consideration applies to the analysis of games. A game's own payoffs are a source of background risk with respect to bets on its outcome, and if the players are sufficiently risk averse, this will give rise to distortions when the rules of the game are revealed through betting. The result will be that a rational solution of the game is characterized by a convex set of equilibria whose parameters are risk neutral probabilities.

Suppose that each player has strictly risk averse subjective-expected-utility preferences with respect to profiles of monetary payoffs in the game, and let U_i denote the strictly concave von Neumann-Morgenstern utility function of player i . Then the payoff profiles $\{x_i\}$ translate into utility profiles $\{U_i(x_i)\}$. Let \mathcal{G}^* denote the true game that is determined by the utility profiles. If U_i' denotes the first derivative of U_i , strict concavity requires that $U_i'(x) < U_i'(y)$ whenever $x > y$. Let \mathbf{u}_i denote the utility payoff vector for player i , whose value in outcome s is

$U_i(x_i(s))$, and let \mathbf{u}_i' denote the corresponding marginal utility vector whose value in outcome s is $U_i'(x_i(s))$. Also, let \mathbf{u}_{ij} denote the vector constructed from \mathbf{u}_i in the same way that \mathbf{x}_{ij} was constructed from \mathbf{x}_i , namely $u_{ij}(s) = U_i(x_{ij}(s))$. In other words, $u_{ij}(s)$ is the utility that player i would receive by playing her j^{th} strategy when all others play according to s . Let \mathbf{u}_{ij}' denote the corresponding profile of marginal utilities for money, i.e., $u_{ij}'(s) = U_i'(x_{ij}(s))$. As in the case of \mathbf{x}_{ij} , there is some duplication of information insofar as $u_{ij}(s)$ and $u_{ij}'(s)$ do not depend on s_i .

By the same argument as in the risk neutral case, player i will choose strategy j in preference to strategy k only if her beliefs are such that she would be willing to exchange the utility profile \mathbf{u}_{ik} , for the utility profile \mathbf{u}_{ij} , hence a small monetary bet yielding a profile of changes in *marginal* utility that is proportional to $\mathbf{u}_{ij} - \mathbf{u}_{ik}$ should be acceptable if the event \mathbf{e}_{ij} is observed to occur. When strategy j is chosen, the agent's profile of marginal utilities for money is \mathbf{u}_{ij}' , and a monetary bet that yields a profile of marginal utilities proportional to $\mathbf{u}_{ij} - \mathbf{u}_{ik}$ can be obtained by dividing the utilities by the corresponding marginal utilities. Therefore agent i should be willing to accept a small bet whose monetary payoffs are proportional to $(\mathbf{u}_{ij} - \mathbf{u}_{ik})/\mathbf{u}_{ij}'$ conditional on the occurrence of \mathbf{e}_{ij} . Such a bet has an unconditional payoff vector of $((\mathbf{u}_{ij} - \mathbf{u}_{ik})/\mathbf{u}_{ij}')\mathbf{e}_{ij}$ in units of money.

Let \mathbf{G}^* now denote the matrix whose rows are indexed by ijk and whose columns are indexed by s and whose ijk^{th} row is the vector $((\mathbf{u}_{ij} - \mathbf{u}_{ik})/\mathbf{u}_{ij}')\mathbf{e}_{ij}$. This is the revealed rules matrix for the game \mathcal{G}^* , representing the information about the game that can be made common knowledge through unilateral offers to accept small bets when the players are risk averse. If an observer chooses a small non-negative vector $\boldsymbol{\alpha}$ of multipliers for these bets, the players as a group will receive the vector of payoffs $\boldsymbol{\alpha} \square \mathbf{G}^*$ and the observer will receive the opposite payoffs. The same rationality criterion that was applied in the risk neutral case also applies here:

an outcome s is jointly coherent if there is no $\alpha \geq \mathbf{0}$ such that $\alpha \sqcap G^* \leq \mathbf{0}$ and $[\alpha \sqcap G^*](s) < 0$.¹¹

The definition of correlated equilibrium and the fundamental theorem of games can now be generalized accordingly. The proof is the same.

DEFINITION: π is a *risk neutral equilibrium* of G^* if and only if $G^*\pi \geq \mathbf{0}$, which means that for every player i and every strategy j and alternative strategy k of that player, either $P_\pi(e_{ij}) = 0$ or else $P_\pi(((u_{ij} - u_{ik})/u_{ij}')e_{ij}) \geq 0$.

THEOREM 4: In a game among risk averse players, a strategy is jointly coherent if and only if there is a risk neutral equilibrium in which it has positive probability.

To provide a story to go with this solution concept, suppose that the players employ a mediator who will use a possibly-correlated randomization device to recommend strategies to them privately, but in this more general case they do not necessarily agree on the true prior probabilities of the outputs of the device. For example, the device may take some of its input data from financial markets or from political or sporting or weather events in addition to or instead of objective randomization. Suppose that through side bets with each other or through participation in a public betting market for the input events, they have arrived at a common prior *risk neutral* probability distribution π for the outputs of the device. Finally, suppose they will not receive any information about realizations of the input events prior to making their moves

¹¹ When the utility functions of the players are strictly concave rather than linear, the bet with payoff vector $((u_{ij} - u_{ik})/u_{ij}')e_{ij}$ is technically only “marginally” acceptable to player i , so a bet with an aggregate payoff vector of $\alpha \sqcap G^*$ may not be quite acceptable to the players for finite α . In such a case the observer may need to make a small side payment to the players to get them to agree to the deal, which makes the observer’s position not entirely riskless. However, if $\alpha \sqcap G^* \leq \mathbf{0}$ and $[\alpha \sqcap G^*](s) < 0$, then by choosing α sufficiently small, the magnitude of the required side payment can be made arbitrarily small in relative terms in comparison to the aggregate loss the players will suffer if they play s , which will be considered here as sufficient grounds for not playing s . This could be made precise by using the concept of ε -acceptable bets (Nau 1995), but it will not be pursued here in the interest of brevity.

except for the private recommendations they receive from the mediator, who will have observed the events by then. Under these conditions, for all i, j , and k , the constraint $P_{\pi}((\mathbf{u}_{ij} - \mathbf{u}_{ik})/\mathbf{u}_{ij}' | e_{ij}) \geq 0$ implies $p_{ij} \square (\mathbf{u}_{ij} - \mathbf{u}_{ik}) \geq 0$, which means that according to player i 's own private beliefs, strategy j yields an expected utility greater than or equal to that of the alternative strategy k when j is recommended to her, so it is optimal for each player to follow the mediator if all others do, and this is common knowledge. Hence a game among risk averse players is played coherently if and only if it is played as-if with the help of a mediator who uses an incentive-compatible device with respect to whose outputs the players have common prior risk neutral probabilities, although their unobserved true probabilities may differ.

A risk neutral equilibrium is a special case of a *subjective correlated equilibrium* (Aumann 1974, 1987), which is an equilibrium that can be implemented through a mediator who uses a randomizing device about whose properties the players may hold differing beliefs. Such a device would be welcome in playing a zero-sum game: all players might believe their expected payoffs to be positive! Aumann (1987) remarks that such a result depends on "a conceptual inconsistency between the players." By permitting such inconsistencies, subjective correlated equilibrium places only weak restrictions on solutions of many games. A risk neutral equilibrium adds the nontrivial restriction that the players' risk neutral prior probabilities should be mutually consistent, as in an equilibrium of a financial market. Whenever agents are risk averse with significant prior investments in events, their true probabilities are not revealed by their preferences among financial assets, and inconsistencies among them are neither surprising nor problematic. This is the norm in financial markets.

As in the risk neutral case, there is more to be said about the rational solution of the game than to identify the outcomes that are jointly coherent. It is also possible to place bounds on risk

neutral probabilities of events and on risk neutral expectations of financial assets that depend on the outcome of the game, namely whatever bounds are determined by the system of inequalities $G^*\pi \geq 0$ that defines the convex polytope of risk neutral equilibria. These bounds are bid-ask spreads for assets that the players are jointly offering to the observer through their bets that reveal information about the rules of the game.

A simple example of the concept of risk neutral equilibrium is provided by the zero-sum game of matching pennies, whose payoff matrix is:

	Left	Right
Top	1, -1	-1, 1
Bottom	-1, 1	1, -1

When played by risk neutral players, the revealed-rules matrix G , scaled to a maximum value of 1, is:

	TL	TR	BL	BR
1TB	1	-1	0	0
1BT	0	0	-1	1
2LR	1	0	-1	0
2RL	0	-1	0	1

This game has a unique correlated/Nash equilibrium in which the players use independent 50-50 randomization, so the graph of the set of equilibria consists of the single point $(\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4})$ in the center of the saddle-shaped set of distributions that are independent between $\{T,L\}$ and $\{B,R\}$.

Now suppose that both players are risk averse, and in particular assume that they have identical exponential utility functions, $U(x) = 1 - \exp(-\rho x)$, in which the risk aversion parameter is $\rho = \text{LN}(\sqrt{2})$. In units of utility, the payoff matrix of the matching-pennies game is then:

	Left	Right
Top	a, b	b, a
Bottom	b, a	a, b

where $a = 1 - \sqrt{1/2} \approx 0.293$ and $b = 1 - \sqrt{2} \approx -0.414$. The corresponding marginal utilities of money in the vicinity of the payoffs a and b are 0.245 and 0.490, respectively, which conveniently differ by a factor of exactly 2.

This game is constant-sum and strategically equivalent to the original one, having the same unique correlated/Nash equilibrium that uses independent 50-50 randomization. However, the rules matrix of the corresponding revealed game, G^* , is *not* equivalent because of the distortions of nonlinear utility for money. When each of its rows is scaled to a maximum value of 1, it is:

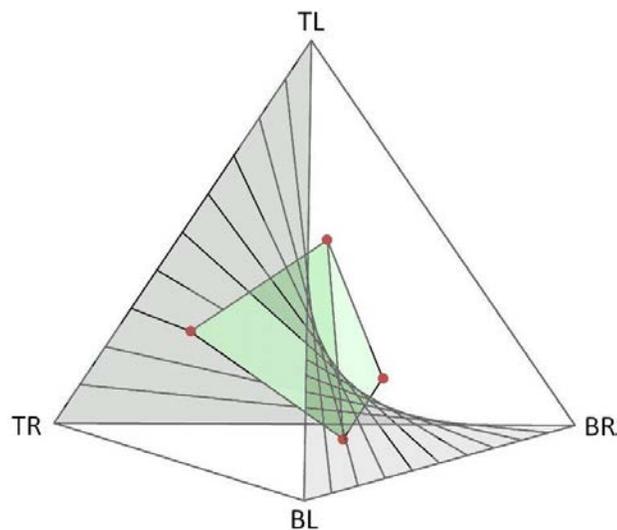
	TL	TR	BL	BR
1TB	1	-1/2	0	0
1BT	0	0	-1/2	1
2LR	-1/2	0	1	0
2RL	0	1	0	-1/2

The polytope of risk neutral equilibria determined by the inequalities $G^*\pi \geq \mathbf{0}$ is no longer a single point. Rather, it “blows up” to a tetrahedron with these vertices:

	TL	TR	BL	BR	EV>0?
Vertex 1	2/15	4/15	1/15	8/15	1BT
Vertex 2	8/15	1/15	4/15	2/15	1TB
Vertex 3	4/15	8/15	2/15	1/15	2RL
Vertex 4	1/15	2/15	8/15	4/15	2LR

None of them is independent between players, so none is a Nash equilibrium of a game with these strategy sets. Each of these probability distributions satisfies 3 out of the 4 incentive constraints with equality, i.e., it assigns an expected value of zero to 3 out of the 4 rows of G^* .

The label of the row whose expected value is positive is shown in the rightmost column of the table. The graph of this solution is shown below. The tetrahedron of risk neutral equilibria is suspended in the middle of the probability simplex, and the saddle of independent distributions cuts through its interior, a situation that would be impossible for a set of correlated equilibria of a game among risk neutral players.



The uniform distribution that is the unique equilibrium of the game when the true utility functions of the players are common knowledge lies in the strict interior of this polytope. When players are risk averse, the small side bets they are willing to accept do not fully reveal the between-strategy differences in utility profiles that they face in the game, so the set of risk neutral equilibria is larger than the set of correlated equilibria. This is true in general, as formalized by:

THEOREM 5: The set of correlated equilibria of a game with monetary payoffs played by risk neutral players is a subset of the set of risk neutral equilibria of the same game played by risk averse players.

Proof: If player i is risk neutral, she will accept a bet with payoff vector $(x_{ij} - x_{ik})\mathbf{e}_{ij}$, while if she is risk averse, she will accept a bet with payoff vector $((u_{ij} - u_{ik})/u_{ij}')\mathbf{e}_{ij}$, where $u_{ij}(s) = U_i(x_{ij}(s))$, and $u_{ij}'(s) = U_i'(x_{ij}(s))$. The term \mathbf{e}_{ij} can be ignored because it zeroes-out the same elements of both vectors. By the subgradient inequality, $U(z) < U(y) - U'(y)(y - z)$, because the value of a strictly concave function U at z must lie below the tangent to its graph at any other point y . Plugging in $y = x_{ij}(s)$ and $z = x_{ik}(s)$ yields $u_{ik}(s) \leq u_{ij}(s) - u_{ij}'(s)(x_{ij}(s) - x_{ik}(s))$, which rearranges to $(u_{ij}(s) - u_{ik}(s))/u_{ij}'(s) \geq x_{ij}(s) - x_{ik}(s)$, with strict inequality if $x_{ij}(s) \neq x_{ik}(s)$. This means that $\mathbf{G}^* \geq \mathbf{G}$ pointwise, i.e., the bet that player i is willing to accept when she chooses strategy j in preference to k if she is risk neutral is weakly dominated by the bet she will accept in the same game if she is risk averse. It follows that $\mathbf{G}\pi \geq \mathbf{0}$ implies $\mathbf{G}^*\pi \geq \mathbf{0}$ for any probability distribution π , so if π is a correlated equilibrium of the game played by risk neutral players, then it is a risk neutral equilibrium of the same game when it is played by risk averse players. \square

Hence, risk aversion introduces even more imprecision into the probabilistic solutions of noncooperative games when their rules must be revealed through credible bets.

6. Rewriting the rules of the game

In situations where the acceptable bets that reveal the rules of the game do not determine a unique probability distribution over its outcomes, it is conceivable that players could choose to accept additional bets in order to reveal more precise information about their joint beliefs. Or, to turn the argument around again, if their joint beliefs are somehow already commonly known with more than the minimal precision, then it is as if they have offered to accept some additional bets. If they are risk neutral and have in fact implemented a Nash or correlated equilibrium, which induces a common prior distribution over outcomes of the game, they cannot both be

made strictly better off through bets with each other. When players are risk averse, this is not necessarily true, and the matching-pennies game provides a good example. When played by risk averse players, it is a negative-sum game in units of utility, and for both players the unique Nash equilibrium has an expected utility that is below their status quo utility. Risk averse players would rather not play this game at all. Furthermore, player 1's marginal utility of money is greater in outcomes TR and BL (her losing outcomes) than in the other two, and vice versa for player 2. The Nash equilibrium is therefore not a *competitive* equilibrium of a financial market in which it is possible for the players to make additional bets that reveal their *solution* of the game in addition to the bets that reveal the *rules* of the game (the latter being the rows of G^*). In the context of the Nash equilibrium, it is desirable to both players to make a bet in which player 1 wins $\$x$ if TR or BL occurs and player 2 wins $\$x$ if TL or BR occurs, for any positive $x \leq 1$. Such a bet changes the rules of the game in a substantive way, but coin-flipping remains a Nash equilibrium. By choosing $x = 1$ they can even zero-out their payoffs, dissolving the game altogether. If they do not bet with each other in this fashion, but instead bet separately with an observer, there is an arbitrage opportunity for the observer that arises from the fact that, at the outset, the players' risk neutral probabilities do not agree if their true probability distributions are uniform.

7. Conclusions

The language of financial markets can be used to address, in a rigorous if somewhat stylized way, the question of how the rules of a noncooperative game and the reciprocal beliefs of the players might come to be common knowledge. It also provides a practical standard of strategic rationality that can be applied when complete or incomplete information about the game's rules and the players' beliefs has become common knowledge by whatever mechanism,

namely the standard of no-arbitrage. A rational solution of a game played under these conditions is typically a convex set of correlated equilibria rather than a Nash equilibrium. The presence of aversion to risk changes the units of analysis to risk neutral probabilities, as in asset pricing theory, and it typically renders the solutions even more imprecise. When risk averse players make bets with each other that reveal information about their beliefs concerning the solution of the game as well as the rules from which they started, they may also be able to rewrite those rules in a mutually beneficial way, merging the concepts of strategic and competitive equilibrium.

References

- [1] Aumann, R. (1974) Subjectivity and Correlation in Randomized Games. *Econometrica* **30**, 445-462
- [2] Aumann, R. (1976) Agreeing to Disagree. *Annals of Statistics* **4**, 1236-1239
- [3] Aumann, R. (1987) Correlated Equilibrium as an Expression of Bayesian Rationality. *Econometrica* **55**, 1-18
- [4] Cooper, R., D. DeJong, R. Forsythe, and T. Ross (1989) Communication in the Battle of the Sexes Game: Some Experimental Results. *Rand Journal of Economics* **20**, 568-587
- [5] de Finetti, B. (1937) La Prévision: Ses Lois Logiques, Ses Sources Subjectives. *Ann. Inst. Henri Poincaré* **7**, 1-68. Translation reprinted in H.E. Kyburg and H.E. Smokler, Eds. (1980), *Studies in Subjective Probability*, 2nd ed., New York: Robert Krieger, 53-118.
- [6] de Finetti, B. (1974) *Theory of Probability, Vol. I*. New York: Wiley
- [7] Ellerman, D. (1984) Arbitrage Theory: A Mathematical Introduction. *SIAM Review* **26**, 241-261

- [8] Forges, F. (1986) Correlated Equilibria in Games With Incomplete Information: A Model With Verifiable Types. *International Journal of Game Theory* **15**, 65-82
- [9] Forges, F. (1993) Five Legitimate Definitions of Correlated Equilibria in Games With Incomplete Information. *Theory and Decision* **35**, 277-310
- [10] Forges, F. (2006) Correlated Equilibrium in Games With Incomplete Information Revisited. *Theory and Decision* **61**, 329-344
- [11] Fudenberg, D., and J. Tirole (1991) *Game Theory*. Cambridge: MIT Press
- [12] Goeree, J., C. Holt and T. Palfrey (2003) Risk Averse Behavior in Generalized Matching Pennies Games. *Games and Economic Behavior* **45**, 97-113
- [13] Hart, S., and D. Schmeidler (1989) Existence of Correlated Equilibria. *Mathematics of Operations Research* **14**, 18-25
- [14] Lad, F. (1996) *Operational Subjective Statistical Methods: A Mathematical, Philosophical, and Historical Introduction*. New York: Wiley
- [15] McKelvey, R., and T. Palfrey (1995) Quantal Response Equilibria for Normal Form Games. *Games and Economic Behavior* **10**, 6-38
- [16] Nau, R., and K. McCardle (1990) Coherent Behavior in Noncooperative Games. *Journal of Economic Theory* **50**, 424-444
- [17] Nau, R. (1992) Joint Coherence in Games of Incomplete Information. *Management Science* **38**, 374-387
- [18] Nau, R. (1995) Coherent Decision Analysis with Inseparable Probabilities and Utilities. *Journal of Risk and Uncertainty* **10**, 71-91
- [19] Nau, R. (2001) De Finetti Was Right: Probability Does Not Exist. *Theory and Decision* **51**, 89-124

- [20] Nau, R. (2003) A Generalization Of Pratt-Arrow Measure To Non-Expected-Utility Preferences And Inseparable Probability And Utility. *Management Science* **49**, 1089-1104
- [21] Nau, R., S. Gomez Canovas and P. Hansen (2004) On the Geometry of Nash Equilibria and Correlated Equilibria. *International Journal of Game Theory* **32**, 443-453
- [22] Nau, R. (2006) Uncertainty Aversion With Second-Order Utilities and Probabilities. *Management Science* **52**, 136-145
- [23] Nau, R. (2011) Risk, Ambiguity, and State-Preference Theory. *Economic Theory* **49**, 437-467
- [24] Ochs, J. (1995) Games With Unique Mixed Strategy Equilibria. *Games and Economic Behavior* **10**, 202-217
- [25] Savage, L. (1971) Elicitation of Personal Probabilities and Expectations. *Journal of the American Statistical Association* **66**, 783-801
- [26] Smith, C. (1961) Consistency in Statistical Inference and Decision. *Journal of the Royal Statistical Society B* **23**, 1-25
- [27] Viossat, Y. (2006) The Geometry of Nash Equilibria and Correlated Equilibria and a Generalization of Zero-Sum Games. S-WoPEc working paper 641, Stockholm School of Economics, Stockholm
- [28] Viossat, Y. (2010) Properties and Applications of Dual Reduction. *Economic Theory* **44**, 53-68