

Imprecise Probabilities in Non-cooperative Games

Robert Nau
Fuqua School of Business
Duke University

GDRR '11 (Belgirate) and ISIPTA '11 (Innsbruck)

- Historically there has been tension between game theory and subjective probability theory
 - Subjective probability theory allows arbitrary personal beliefs, subject only to internal consistency (coherence).
 - Game theory assumes that beliefs should be dictated by equilibrium conditions (often uniquely) and players should have “common priors” for exogenous uncertainties.
 - Exchange between Kadane/Larkey and Harsanyi (*Management Science* 1982) highlighted this conflict

- Also a diversity of rationality standards and equilibrium models for games and markets:
 - Game theory: Nash & Bayesian equilibrium & refinements
 - Consumer theory: Pareto efficiency, Walrasian equilibrium
 - Finance theory: CAPM, no-arbitrage...
- Competitive and strategic equilibrium are distinct concepts (the latter usually regarded as deeper and more fundamental)
- Nash and Walrasian equilibria are fixed points; Pareto efficiency and no-arbitrage are dual conditions for existence of equilibrium prices
- Pareto efficiency is itself a form of no-arbitrage: a Pareto inefficient allocation presents a riskless profit opportunity if its existence is public knowledge and transaction costs are low.

- This paper presents a unification of these ideas, extending Nau-McCardle's (1990) work on "joint coherence" in games
 - *Coherence* (no-arbitrage) is the primal rationality axiom that unifies subjective probability theory, game theory, consumer theory, and finance theory.
 - In non-cooperative games, this leads to the solution concept of *correlated equilibrium* (a coarsening of Nash equilibrium).
 - The "solution" of a game based only on common knowledge of the rules and common knowledge of rationality is generically the *convex set* of all correlated equilibria.
 - When agents are risk averse, the observable parameters of equilibria are *risk neutral probabilities* (products of probabilities and relative marginal utilities for money, as in asset pricing), and the set of risk neutral equilibria is typically larger than the set of correlated equilibria.

Modeling rational beliefs

- The starting point: de Finetti's use of monetary bets to reveal personal probabilities and expectations
- Your subjective lower expectation ("lower prevision") for an asset with payoff vector \mathbf{x} conditional on an event \mathbf{e} (" $\mathbf{x}|\mathbf{e}$ ") is the price p you would be willing to pay for small multiples of \mathbf{x} at the discretion of an observer, subject to the deal being called off if \mathbf{e} does not occur.
- This means you are willing to accept a bet with payoff vector $\alpha(\mathbf{x} - p)\mathbf{e}$ for any small positive α chosen by the observer.
- Example: $\mathbf{x} = (6, 3, 5)$, $\mathbf{e} = (1, 1, 0)$, $p = 4$
 \Rightarrow you will accept a bet with net payoff vector $(2\alpha, -\alpha, 0)$

(Here \mathbf{e} is the indicator vector for states 1 and 2.)

- Underappreciated fact: this is not only a definition of your lower expectation of \mathbf{x} , it is a definition of *common knowledge* of your lower expectation of \mathbf{x} .
- Your announcement creates a real financial opportunity for the observer, and he knows it, and you know he knows it, etc., and you both understand the meaning of the numbers (prices and quantities) in exactly the same way.
 - *This issue is usually glossed over in game theory: what does it mean, in material terms, for the numerical values of subjective parameters such as utilities and prior probabilities to be common knowledge? Directly confronting this issue (even if in a stylized way) is the key to a more unified theory.*
- Important special case: your lower expectation for \mathbf{x} is zero if \mathbf{x} is itself an acceptable bet, i.e., you are willing to accept $\$ \alpha \mathbf{x}$ for any small positive α chosen by an observer.

Notation for probabilities and expectations

- N denotes the number of states, $\mathbf{x} = (x_1, \dots, x_N)$ denotes the payoff vector of an asset, $\pi = (\pi_1, \dots, \pi_N)$ denotes a probability distribution, and \mathbf{e} denotes a 0-1 vector that is the indicator of an event (i.e., $e_n \in \{0, 1\}$). “Dot” (\cdot) denotes inner product. Non-dot products of vectors (e.g., “ $\mathbf{x}\mathbf{e}$ ”) are taken pointwise.
- $P_\pi(\mathbf{x}) := \pi \cdot \mathbf{x}$ is the expected value of \mathbf{x} determined by π
- $P_\pi(\mathbf{e})$ is the probability of \mathbf{e} determined by π
- $P_\pi(\mathbf{x}|\mathbf{e}) := P_\pi(\mathbf{x}\mathbf{e})/P_\pi(\mathbf{e})$ is the conditional expectation of \mathbf{x} given \mathbf{e} if $P_\pi(\mathbf{e}) > 0$
- Assertion of p as your conditional lower prevision for \mathbf{x} given \mathbf{e} means willingness to accept a bet with payoff vector proportional to $(\mathbf{x} - p)\mathbf{e}$, whose value in state n is $(x_n - p)e_n$

Fundamental theorem of subjective probability

- Definition: Conditional lower previsions $\{p_1, \dots, p_M\}$ for assets $\{\mathbf{x}_1|\mathbf{e}_1, \dots, \mathbf{x}_M|\mathbf{e}_M\}$ are *coherent* [ex post coherent in state n] if there do not exist non-negative bet multipliers $\{\alpha_1, \dots, \alpha_M\}$ such that $\sum_m \alpha_m (\mathbf{x}_m - p_m) \mathbf{e}_m < \mathbf{0}$ [$\leq \mathbf{0}$ and $< \mathbf{0}$ in state n]
 - Coherence means you are not exposed to a sure loss, and ex post coherence in state n means you are not exposed to a loss in state n with no possibility of a gain in some other state.
- Theorem 1: Conditional lower previsions are coherent [ex post coherent in state n] iff there is a nonempty nonconvex set of probability distributions Π [satisfying $\pi_n > 0$] such that, for all m , either $P_\pi(\mathbf{x}_m|\mathbf{e}_m) \geq p_m$ or else $P_\pi(\mathbf{e}_m) = 0$ for all $\pi \in \Pi$.
 - This is de Finetti's theorem, which can be proved by a separating hyperplane argument (linear programming duality).

- Thus, conditional lower previsions that are coherent are those that are “rationalized” by some non-empty convex set of probability distributions in the sense that they are not-necessarily-tight lower bounds on the conditional expectations determined by the distributions from that set which assign positive probability to their conditioning events.
- They are ex post coherent in state n if this condition holds with respect to a set of distributions which all assign strictly positive probability to state n .

Extension of de Finetti's theorem to games

- Consider a 2x2 game with payoff matrix:

	Left	Right
Top	a_1, a_2	b_1, b_2
Bottom	c_1, c_2	d_1, d_2

- How can these payoff functions be made common knowledge, and in how much detail?
- What restrictions does this knowledge place on rational beliefs, and why?
- What outcomes of the game are rational?

- To extend the logic of De Finetti's theorem to games, we must consider situations in which some of the events on which bets are placed are under your control, while others are under your opponent's control.
 - In games of incomplete information, nature could also be a player
- In general you will not want to make unconditional bets on your opponent's moves: they might reveal too much private information about your beliefs and expose you to exploitation.
- Nor will you want to make unconditional bets on your own moves: they might reveal too much about your intentions, and others would be reluctant to take the other side of them.

- However... it might make sense for you to offer to accept bets whose payoffs depend on your opponent's moves *conditional on moves of your own*.
- Such bets reveal properties of the beliefs you *would need to have* in order to prefer a given move over other moves that are available.
- In this manner they reveal (limited) information about your payoff function *without* necessarily revealing your current beliefs and intentions... and they do so in terms that are common knowledge.
 - It is not *necessary* for players to accept such bets, but it is a way for them to credibly make some information about their payoff functions common knowledge if they wish to. Or, if by magic their payoff functions are already common knowledge, there is no loss of generality in assuming they will offer to accept such bets

Modeling common knowledge of payoffs

- In the event that player 1 chooses Top when she could have chosen Bottom, she evidently prefers the payoff profile (a_1, b_1) over (c_1, d_1) based on her then-current beliefs about whether player 2 will choose Left or Right.
- In that case a small *bet* proportional to the more-preferred payoffs *minus* the less-preferred payoffs, $(a_1 - c_1, b_1 - d_1)$, should be desirable to her (assuming risk neutrality).
- She can make this fact common knowledge by publicly offering to accept such a bet *conditional on her choice of Top*, at the discretion of an observer.
- Similarly, she can offer to accept the opposite bet, whose payoff profile is $(c_1 - a_1, d_1 - b_1)$, conditional on Bottom.

- For both players combined, there are 4 such bets, whose payoff vectors are the rows of the following matrix (“G”):

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

Blank cells
contain
zeros

- Here, “1TB” is the bet acceptable to player 1 in the event she chooses T in preference to B, etc.
- Each column corresponds to an outcome of the game.
- A given bet yields non-zero payoffs only in outcomes where its chosen strategy is played, e.g., the 1TB bet yields non-zero payoffs only in outcomes TL and TR, where T is played.

- Claim: the matrix G consists precisely of the information about the game's rules which can be made common knowledge under non-cooperative conditions, i.e., via unilateral offers to accept bets at anyone else's discretion.
- It suffices to determine all the equilibria of the game (correlated, Nash, and refinements thereof).
- What it includes: information about how a given player's payoff depends on *her own* strategy, holding the others' strategies fixed.
- What it does *not* include: information about how a given player's payoff depends on *others'* strategies, holding her own strategy fixed (i.e., possible benefits of cooperation)
- This construction generalizes to any number of players and strategies: for player i there are $|S_i| \times (|S_i| - 1)$ acceptable bets, corresponding to every possible combination of a chosen strategy and an alternative strategy

Modeling common knowledge of rationality

- Let M denote the total number of acceptable bets for all players (# of rows of G) and let $\alpha = (\alpha_1, \dots, \alpha_M)$ denote a vector of non-negative multipliers for them, which is chosen by the observer. Then $\alpha \cdot G$ is the net payoff vector for the players as a group.
- Is there ever an *ex ante* arbitrage opportunity, i.e., a choice of α that yields a *strictly negative* payoff vector for the players as a group? Fortunately not!
 - This can be proved by showing that for any α there is a profile of independently randomized strategies that yields an expected payoff of exactly zero for each player (hence also for the group), which can be computed as a stationary distribution of a Markov chain whose transition matrix is determined from α .
- However, there may be *ex post* arbitrage opportunities...

- Suppose that for some game outcome \mathbf{s} , there is an $\alpha \geq 0$ such that $\alpha \cdot G \leq 0$ (players can't win money from the observer) and $[\alpha \cdot G](\mathbf{s}) < 0$ (they lose money to him if they play \mathbf{s}).
- Such an outcome is “jointly incoherent”: it yields an *ex post* arbitrage profit for the observer if it is played. All other outcomes are “jointly coherent”.
 - There must be at least one jointly coherent outcome, because there can be no *ex ante* arbitrage opportunity for the observer, as already noted.
- *Joint coherence (avoidance of ex post arbitrage by the group) is the minimal material standard of common knowledge of rationality: the players as a group should not deliberately throw money away, and they all know it, and all know that they all know it....*

Correlated equilibrium

- A *correlated equilibrium* is an equilibrium in pure or randomized strategies, in which randomization is permitted to be correlated (Aumann 1974, 1987).
- A correlated strategy can be implemented via a mediator who uses a correlated randomization device to make private strategy recommendations to players.
- A Nash equilibrium is a special case of a correlated equilibrium in which randomization, if any, must be independent between players.
- A third-party mediator is not necessarily required for correlation: a public device such as coin-flipping might suffice in a coordination game such as battle-of-the-sexes.

Correlated vs. Nash equilibrium

	Left	Center	Right
Top	0, 0	2, 1	1, 2
Middle	1, 2	0, 0	2, 1
Bottom	2, 1	1, 2	0, 0

Here's an example of a 3x3 game in which coordination is desirable: the players would rather not end up in one of the diagonal outcomes.

Correlated vs. Nash equilibrium

	Left	Center	Right
Top	0, 0	2, 1	1, 2
Middle	1, 2	0, 0	2, 1
Bottom	2, 1	1, 2	0, 0

In the unique Nash equilibrium of this game, all outcomes have probability $1/9$ and the expected payoff is 1 for both players. These are actually the *lowest possible expected payoffs* for both players within the larger set of all correlated equilibria.

Correlated vs. Nash equilibrium

	Left	Center	Right
Top	0, 0	2, 1	1, 2
Middle	1, 2	0, 0	2, 1
Bottom	2, 1	1, 2	0, 0

An obvious and efficient and equitable correlated equilibrium is one that assigns probability $1/6$ to all the non-zero payoff cells, yielding an expected payoff of 1.5 to both players. This could be implemented through a mediator who rolls a 6-sided die to randomly select a cell with non-zero payoffs and then tells each player privately what she should play.

Correlated vs. Nash equilibrium

	Left	Center	Right
Top	0, 0	2, 1	1, 2
Middle	1, 2	0, 0	2, 1
Bottom	2, 1	1, 2	0, 0

In the die-rolling equilibrium, if Row is privately told to play Bottom, she knows her opponent is equally likely to have been told to play Left or Center. Her expected payoff is 1.5 for playing Bottom as instructed, vs. 0.5 for playing Middle or 1.0 for playing Top, on the assumption that her opponent also follows the mediator's instructions.

- The set of all correlated equilibria is the set of probability distributions $\{\pi\}$ satisfying $G\pi \geq 0$, where G is the matrix of acceptable bets that also determines jointly coherent strategies
- These are precisely the incentive constraints that must be satisfied for players to follow a mediator who generates private strategy recommendations according to π .
- This system of linear inequalities defines a convex polytope.
- The Nash equilibria of a game always lie on the *surface* of the polytope, i.e., on a supporting hyperplane to it, although not necessarily at extreme points (Nau et al. 2004).
- The set of Nash equilibria can be quite weird: non-convex, disconnected, and/or consisting only of irrational probabilities
- The correlated equilibrium polytope is a “nice” object, but can have many vertices (e.g., $> 100,000$ in a 4x4 game).

Fundamental theorem of noncooperative games

- Theorem 2: An outcome of the game is jointly coherent if and only if there exists a correlated equilibrium in which it occurs with positive probability.
 - Proof: By a separating-hyperplane argument (linear programming duality again), the system of inequalities $\alpha \geq 0$, $\alpha \cdot G \leq 0$, and $[\alpha \cdot G](\mathbf{s}) < 0$ has no solution if and only if the system of inequalities $G\pi \geq 0$, $\pi \geq 0$, $\pi(\mathbf{s}) > 0$ has a solution. Here, π is the normal vector of a hyperplane separating the convex hull of the rows of G from the negative orthant.
- Corollary: There is at least one correlated equilibrium
 - Proof: This follows from the existence of at least one jointly coherent strategy. It is more elementary than the proof of existence of a Nash equilibrium, insofar as it requires only the duality result and the existence of a stationary distribution of a Markov chain, not the existence of a fixed point of a continuous mapping of a set into itself. It is also constructive: finding a jointly coherent strategy and finding a correlated equilibrium are a primal-dual pair of linear programs.

- So, common knowledge of the game's rules (revelation of G) determines a convex set of probability distributions concerning its outcome, namely the set of correlated equilibria, and the rational outcomes of the game are the ones with positive support in this set.
- A Nash equilibrium is a special case of a correlated equilibrium in which randomization, if any, is independent between players
- Correlated equilibrium, not Nash equilibrium, is the more fundamental concept.
- *The “solution” of the game that follows (only) from common knowledge of its rules and common knowledge of rationality is the entire set of correlated equilibria, i.e., a possibly-imprecise probability distribution.*

Example: battle of the sexes

- The privately-known payoff matrix:

	Left	Right
Top	2, 1	0, 0
Bottom	0, 0	1, 2

- The commonly-knowable rules matrix (G):


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

- The vertices of the correlated equilibrium polytope consist of two pure Nash equilibria, one mixed Nash equilibrium, and two non-Nash equilibria:


	TL	TR	BL	BR	Nash?
Vertex 1	1	0	0	0	Yes
Vertex 2	0	0	0	1	Yes
Vertex 3	$2/9$	$4/9$	$1/9$	$2/9$	Yes
Vertex 4	$2/5$	0	$1/5$	$2/5$	No
Vertex 5	$1/4$	$1/2$	0	$1/4$	No

- Absent any constraints other than common knowledge of the rules of the game (acceptability of bets that are the rows of G) and common knowledge of rationality (joint coherence), the “solution” of this game consists of the convex hull of these vectors, i.e., the entire set of correlated equilibria.


- The battle-of-sexes correlated equilibrium polytope has 5 vertices and 6 faces (the most possible in a 2x2 game):



- The 3 points at which it touches the “saddle” of independent distributions are the Nash equilibria



- Nash equilibria always lie on the *surface* of the correlated equilibrium polytope, but need not be vertices of it in games with more than 2 players—they can be isolated points in the middle of edges or lie along curves within faces.



An obvious and efficient and equitable solution of the battle-of-sexes game is for the players to flip a coin to choose between TL and BR, which is neither a Nash equilibrium nor an extremal correlated equilibrium. It yields an expected payoff of 1.5 to both players.

The mixed-strategy Nash equilibrium is the worst possible equilibrium solution, yielding an expected payoff of only $2/3$ to each player. If they are unable to coordinate on TL or BR or implement a correlated strategy, they would be better off flipping coins independently, which yields an expected payoff of $3/4$ to both, although it is not an equilibrium.

Games with risk averse players

- If players are risk averse, the small bets they accept will be distorted by the nonlinearity of their utility functions, which causes their marginal utility for money to be state-dependent.
- The result will be that they “hedge their bets,” i.e., accept bets that are less risky than if they were risk neutral.
 - This is true even for vanishingly small bets: it is a first-order effect, not a second-order effect, in this setting
- It introduces even more imprecision into the solution, and....
- It changes the interpretation of the parameters of the equilibrium distributions from “true” probabilities to “risk neutral” probabilities, as in financial markets.

- De Finetti's method of defining and eliciting subjective probabilities assumes state-independent marginal utility for money.
- What if marginal utility is state-dependent due to risk aversion and significant "background risk"?
- Suppose the agent has a strictly concave vNM utility function $U(x)$, with its derivative denoted by $U'(x)$.
- Suppose her background risk is represented by a payoff vector \mathbf{z} whose range of values is "large" in comparison to her risk tolerance.
- Then the acceptability for her of a "small" bet \mathbf{x} will not be based on whether its expected value is non-negative but rather on whether it has a non-negative impact on her overall expected utility when it is added to \mathbf{z} .

- If the agent's beliefs are represented by a precise probability distribution \mathbf{p} , then her status quo expected utility is $E_{\mathbf{p}}[U(\mathbf{z})]$, and \mathbf{x} is an acceptable bet if $E_{\mathbf{p}}[U(\mathbf{z}+\mathbf{x})] \geq E_{\mathbf{p}}[U(\mathbf{z})]$.
- 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 its *expected marginal utility* is non-negative, i.e., $E_{\mathbf{p}}[U'(\mathbf{z})\mathbf{x}] \geq 0$.
- This condition is equivalent to $E_{\pi}[\mathbf{x}] \geq 0$ 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, i.e.,

$$\pi(\mathbf{s}) \propto p(\mathbf{s})U'(z(\mathbf{s})).$$

- π is called the *risk neutral probability distribution of the agent* at \mathbf{z} , because she evaluates small bets risk-neutrally with respect to it.
- That is, she is willing to accept a small bet whose *expected value* is non-negative when it is computed with respect to π rather than \mathbf{p} .
- In an equilibrium of a financial market, agents must agree publicly on risk neutral probabilities (although their true probabilities may differ), because otherwise there is an arbitrage opportunity on-the-margin, i.e., there is a violation of the law-of-one-price for contingent claims.

Incorporating risk aversion into the model of coherent behavior in games

- Let U_1 and U_2 denote the von Neumann-Morgenstern utility functions of players 1 and 2, and assume they are strictly concave (risk averse).
- Let U_1' and U_2' denote their derivatives (marginal utilities of money).

- In units of *money* the privately-known payoff matrix of a 2x2 game is:

	Left	Right
Top	a_1, a_2	b_1, b_2
Bottom	c_1, c_2	d_1, d_2

- In units of *money* the privately-known payoff matrix of a 2x2 game is:

	Left	Right
Top	a_1, a_2	b_1, b_2
Bottom	c_1, c_2	d_1, d_2

- ...and in units of *utility* it is:

	Left	Right
Top	$U_1(a_1), U_2(a_2)$	$U_1(b_1), U_2(b_2)$
Bottom	$U_1(c_1), U_2(c_2)$	$U_1(d_1), U_2(d_2)$

- Now consider the effect of risk aversion on acceptable bets, which are the parameters of the situation that are commonly knowable.
- If player 1 chooses Top over Bottom, she evidently prefers the *utility-payoff* profile $(U_1(a_1), U_1(b_1))$ over $(U_1(c_1), U_1(d_1))$, based on her then-current beliefs about whether player 2 is going choose Left or Right.
- This means she would be willing to accept a small bet on Left-vs-Right whose payoffs in *utils* are proportional to the difference, i.e., $(U_1(a_1) - U_1(c_1), U_1(b_1) - U_1(d_1))$, conditional on choosing Top.

- In unobservable units of utility, the payoff vectors of acceptable small bets would therefore be proportional to:

	TL	TR	BL	BR
1TB	$U_1(a_1) - U_2(c_1)$	$U_1(b_1) - U_1(d_1)$		
1BT			$U_1(c_1) - U_1(a_1)$	$U_1(d_1) - U_1(b_1)$
2LR	$U_2(a_2) - U_2(b_2)$		$U_2(c_2) - U_2(d_2)$	
2RL		$U_2(b_2) - U_2(a_2)$		$U_2(d_2) - U_2(c_2)$


- In unobservable units of utility, the payoff vectors of acceptable small bets would therefore be proportional to:

	TL	TR	BL	BR
1TB	$U_1(a_1) - U_2(c_1)$	$U_1(b_1) - U_1(d_1)$		
1BT			$U_1(c_1) - U_1(a_1)$	$U_1(d_1) - U_1(b_1)$
2LR	$U_2(a_2) - U_2(b_2)$		$U_2(c_2) - U_2(d_2)$	
2RL		$U_2(b_2) - U_2(a_2)$		$U_2(d_2) - U_2(c_2)$

Player 1's marginal utility of money in outcome TL is $U_1'(a_1)$, because her monetary payoff is a_1 there.

- In unobservable units of utility, the payoff vectors of acceptable small bets would therefore be proportional to:

	TL	TR	BL	BR
1TB	$U_1(a_1) - U_2(c_1)$	$U_1(b_1) - U_1(d_1)$		
1BT				$U_1(d_1) - U_1(b_1)$
2LR	$U_2(a_2) - U_2(b_2)$		$U_2(c_2) - U_2(d_2)$	
2RL		$U_2(b_2) - U_2(a_2)$		$U_2(d_2) - U_2(c_2)$



Her marginal utility of money in outcome TR is $U_1'(b_1)$, because her monetary payoff is b_1 there.

- To convert the acceptable bets to observable and commonly-understood units of *money*, each utility payoff must be divided by the local marginal utility of money:

	TL	TR	BL	BR
1TB	$\frac{U_1(a_1) - U_2(c_1)}{U_1'(a_1)}$	$\frac{U_1(b_1) - U_1(d_1)}{U_1'(b_1)}$		
1BT			$\frac{U_1(c_1) - U_1(a_1)}{U_1'(c_1)}$	$\frac{U_1(d_1) - U_1(b_1)}{U_1'(d_1)}$
2LR	$\frac{U_2(a_2) - U_2(b_2)}{U_2'(a_2)}$		$\frac{U_2(c_2) - U_2(d_2)}{U_2'(c_2)}$	
2RL		$\frac{U_2(b_2) - U_2(a_2)}{U_2'(b_2)}$		$\frac{U_2(d_2) - U_2(c_2)}{U_2'(d_2)}$

Dividing the utility payoffs by the local marginal utilities for money yields the payoffs of acceptable monetary bets

- Let G^* denote this matrix: the revealed rules of the game among risk averse players
- Definition: A *risk neutral equilibrium* is a distribution π satisfying $G^*\pi \geq 0$
- Theorem 3: An outcome of the game among risk averse players is jointly coherent iff it occurs with positive probability in some risk neutral equilibrium.
 - The proof is the same as before.
- Theorem 4: 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, i.e., *risk aversion makes the solution even more imprecise.*

- An explicit risk neutral equilibrium can be implemented via a correlation device with respect to whose outputs the players have *common prior risk neutral probabilities*.
- For example, the correlation device could derive some of its inputs from outcomes of market events, sporting events, weather events, political events, etc., about which the players may have different subjective beliefs.
- Common prior risk neutral probabilities for the input events could be achieved through the players' ex ante participation in a public financial market or betting market.
- Participation in such markets generally does not reveal or reconcile anyone's "true" subjective probabilities when the players are risk averse and heavily invested (so what?)
- ..but *it changes the rules of the game* (more about this later)

Consistency of beliefs?

- Risk neutral equilibrium is a refinement of the concept of subjective correlated equilibrium (Aumann 1974).
- A subjective correlated equilibrium is a correlated equilibrium involving possibly-heterogeneous subjective probabilities, i.e., an uncommon prior.
- Aumann objects to the “conceptual inconsistency” that would be inherent in such a solution.
- A risk neutral equilibrium imposes the additional requirement that the players’ risk neutral probabilities should be mutually consistent.
- The inconsistency of the true probabilities is not relevant (just as it is not relevant in financial markets), because they are generally unobservable when players are significantly risk averse with unobservable true payoffs.

Example: the zero-sum game of “matching pennies”:

	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

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.

Suppose that both players are risk averse with exponential utility functions, $U(x) = 1 - \exp(-\rho x)$, where the risk aversion parameter is $\rho = \underline{\text{LN}}(\sqrt{2})$. In units of utility, the payoff matrix is then:

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

where $a = \underline{\text{LN}}(1/2) \approx 0.293$ and $b = 1 - \sqrt{2} \approx -0.414$.

The corresponding marginal utilities of money under the outcomes a and b are 0.245 and 0.49, respectively, which conveniently differ by a factor of exactly 2.

The “true” game with these utility payoffs still has a unique Nash/correlated equilibrium, namely the uniform distribution, but that fact is not observable.

The rules matrix of the corresponding revealed game, G^* , looks like this when scaled to a maximum value of 1:

	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 a tetrahedron with these vertices:


	TL	TR	BL	BR
Vertex 1	2/15	4/15	1/15	8/15
Vertex 2	8/15	1/15	4/15	2/15
Vertex 3	4/15	8/15	2/15	1/15
Vertex 4	1/15	2/15	8/15	4/15

Here is a picture of the polytope of risk neutral equilibria for the matching-pennies game between risk-averse players.

The polytope is of full dimension (a tetrahedron) and the saddle of independent distributions cuts through its interior, which would be impossible for a full-dimension correlated equilibrium polytope such as that of battle-of-the-sexes.

Its barycenter is the uniform distribution that is the unique correlated equilibrium of the “true” game.

In general, the set of risk neutral equilibria is a superset of the set of correlated equilibria, i.e., risk aversion introduces even more imprecision into the solution.



What do the players know, and what should they do?

- How should the players decide what to do here?
- If they knew each others' true utility payoffs, they would know to play the unique correlated/Nash equilibrium.
- But they don't have this much information, so their strategies are not (yet) uniquely determined.
- The set of risk neutral equilibria doesn't include any pure strategies, so some independent or correlated randomization is needed if the game is played deliberately within these rules.
- A priori, the players (and observer) can only put bounds on the risk neutral probabilities of game outcomes.
- They could further constrain the risk neutral probabilities by making additional bets, beyond those that reveal the rules...
- ...but this would *change* the rules to some extent.

Re-writing the rules

- In a game among risk averse players, unlike a game among risk neutral players, the players may be able to make themselves all better off (get a more Pareto efficient solution) by betting with each other to hedge their mutual risks.
- The concept of risk neutral equilibrium is thus a hybrid of the concepts of strategic equilibrium and competitive equilibrium.
- *Risk averse players would rather not play matching pennies at all:* they would find it mutually desirable to zero-out their payoffs in the game through side bets with each other.
- If risk averse players are in fact playing the Nash equilibrium in the original game, and if they accept additional side bets consistent with their equilibrium beliefs, there is an arbitrage opportunity for an observer, because the Nash equilibrium is not a competitive equilibrium of the betting market.

- Summary:
 - *Coherence* (no-arbitrage) can be extended in a natural way from subjective probability theory to game theory by applying it to the conditional bets that would be needed to make the rules of the game common knowledge.
 - This leads to the solution concept of *correlated equilibrium*, a generalization (coarsening) of Nash equilibrium.
 - The “solution” of a game based only on common knowledge of its rules and common knowledge of rationality is generically a *convex set of probability distributions* consisting of all the correlated equilibria.
 - Of course, in any particular game, the players are free to “refine” the solution as they wish (e.g., via flipping a coin in battle-of-the-sexes), which they could make common knowledge through acceptance of additional bets.

- Summary, continued:
 - When agents are risk averse and the ranges of game payoffs are large in comparison to their risk tolerances, the corresponding solution concept is *risk neutral equilibrium*
 - The parameters are *risk neutral probabilities* (products of probabilities and relative marginal utilities for money), as in models of asset pricing by arbitrage
 - Risk aversion tends to enlarge the set of equilibria, as illustrated by the matching-pennies game
 - Side bets with each other may be mutually desirable for risk averse players, and they may alter the rules of the game.
 - This approach to the modeling of rational behavior in games yields a unification of game theory with subjective probability theory (à la de Finetti) and asset pricing theory

Some final thoughts

- These results are admittedly very stylized: outside of financial markets, casinos, and sports gambling, individuals do not often speak to each other in the language of bets—particularly with respect to each others' actions—and observers are not hovering over them, looking for arbitrage opportunities.
- But... any model of a game that assumes commonly known payoffs and equilibrium behavior is implicitly assuming at least as much as is assumed here, and generally a lot more.
- Where do the commonly-known numbers come from in such models? In situations where money is not used as a yardstick, it may be problematic to assume common knowledge and common understanding of precise numbers that measure personal beliefs and values.

Some final thoughts

- Uniqueness and stochastic independence of equilibrium strategies are also problematic: games often have many equilibria, independent and correlated, and issues beyond risk aversion—e.g., indeterminacy of payoffs, incomplete preferences—may further enlarge the set.
- Why allow correlated strategies in the solution? In some situations (such as battle-of-the-sexes), deliberate correlation may be desirable, and it does not always require a mediator: coin-flipping or playing paper-scissors-rock or following a taking-turns convention in repeated play may suffice.
- In other situations, correlation may only be a property of the beliefs of an observer, e.g., if there are multiple equilibria and the observer believes the players have selected one but he does not know which one.

Some final thoughts

- Last but not least, noncooperative solution concepts use only half of the data about the players' stakes in the game!
- Information about how one player's payoffs depend on the moves of others, holding her own move fixed (which is subtracted out of G and G^*), is what determines *which side of the equilibrium polytope is the efficient frontier*.
- The left-out data may play a role in determining which solutions are attractive or focal, whether the game is a win-win situation or a prisoner's dilemma, and whether and how the rules ought to be changed by, e.g., by large side bets, multilateral contracts, or further mechanism design.