Rochester Center for

Economic Research

A Few Humble Observations on Overconfidence and Equilibrium

Francesco Squintani

Working Paper No. 480
April 2001

UNIVERSITY OF

ROCHESTER

A Few Humble Observations on Overconfidence and Equilibrium*

Francesco Squintani
University of Rochester
Department of Economics†

February, 2001

Abstract

This paper describes equilibrium in games where the informed players may be overconfident. Motivated by specific moral-hazard, signalling and screening problems, we first assume that the "uninformed" players know that "informed" players may be mistaken, but that the "informed" players are unaware of this. In standard Bayesian games, we identify a conflict between self-perception and equilibrium conjectures. We thus turn to population games and assume that while each player believes that her own perception is correct, she also knows that the other players in the population are on average overconfident. It is shown that in any equilibrium of any such game, players cannot be made better-off by being overconfident. Overconfidence may be beneficial only when comparing payoff across different games, or across different equilibria of the same game. The second part of the paper considers any description of high-order knowledge of overconfidence. We determine the descriptions that allow to construct an equilibrium concept immune to introspective conflicts. It is shown that overconfidence cannot make any player better off also in the case that she is aware that the opponents think that she is overconfident. The paper is concluded by showing how to translate our knowledge-based analysis in the language of Mertens and Zamir (1985) universal types.

JEL CLASSIFICATION: J31, D82, D83

*I thank Larry Blume, Eddie Dekel, Massimo Marinacci, Juuso Valimaki, and Asher Wolinsky.
†Harkness Hall, University of Rochester, Rochester, NY 14627.

"...and that's the news from Lake Wobegon, Minnesota, where all the women are strong, and all the men are good-looking, and all the children are above average !"

Garrison Keillor

"A Prairie Home Companion"

1 Introduction

In Bayesian games modelling adverse-selection, signalling or screening problems, some individuals' unobservable characteristics (such as ability, intelligence and so on) influence the players' payoff. It is customary to assume that informed players precisely assess their own characteristics. Well-settled experimental evidence from the psychology literature contradicts that assumption. Subjects on average overestimate their own personal characteristics,[1] even though they do not suffer from this bias when judging the characteristics of others.[2] The aim of this paper is to describe the equilibrium of Bayesian games with private information where the informed parties may be mistaken in their self-perception.[3]

Before proceeding with the presentation, it is important to underline a few qualifications on the object of this paper's analysis. We intend *confidence* to mean one's private information about her characteristics, and we will employ the terms self-confidence and self-perception interchangeably. It is important to separate overconfidence (a mistake in self-judgement) from "overconfident" behavior. One can easily describe strategic interactions where an overconfident player deceivingly chooses to take a very humble course of action. A second distinction is in place. It is conceivable that overconfidence has a direct psychologic effect on a player's utility, common wisdom deems in fact that confident people feel better about themselves.

[1] The experiment protocols of Einhorn and Hogarth (1978), Svenson (1981), and Segerstrom et al. (1993) ask each subject in the sample in which percentile of the experiment pool she belongs to, with respect to intelligence, ability etc... The typical finding is that most subject believe to be above average.

[2] Among the experimental studies showing that self-evaluations (presumably incorrect) are systematically higher than the evaluation made by others, see Lewinsohn et al (1980), Taylor and Brown (1988), Burger and Burns (1988).

[3] In order to present the different themes of this contribution in the most transparent manner, we will restrict attention to 2-player games with one-sided private information. This restriction does not entail any relevant loss of generality. It is easy to see how to generalize the construction and results to any Bayesian game.

However we assume that overconfidence does not directly influence the player's payoff, and can only have an indirect effect on utility by modifying the player's behavior. In fact, we would like to clearly separate self-perception, which is a payoff-irrelevant property of beliefs, from self-esteem, which has a direct effect on one's welfare.[4]

With respect to the game-theoretical structure of the interactions analyzed in this paper, we should say that our formulation is quite general, as it deals with all the equilibria of any Bayesian games. In particular, this framework does not apply only to one-shot games, rather it allows for any sequential structure of the player's moves. This paper thus applies indifferently to the case of one-shot interactions, dynamic or repeated interactions. Also, by explicitly including nature's moves in the framework, this paper allows for the option of experimentation and learning by the players, in any possible fashion. Since our results hold for any equilibrium of a given Bayesian game, they hold a fortiori for any equilibrium refinement that embeds considerations motivated by robustness, or by a particular sequence of choices in an underlying extensive form game.

While no constraints are imposed on the structure of the game, the first part of the paper imposes some restrictions on the players' high-order knowledge of overconfidence, motivated by the economic scenarios for which we want to account. Consider, for instance, a worker overconfident about her own ability who starts a position in a firm, or an overoptimistic entrepreneur who applies for a loan in a bank, or an individual with a mistaken perception of her riskiness applying for an insurance. By construction, it can never be possible that the informed player is aware that she may be mistaken, or else she would be able to revise her own self-judgement and correctly assess her own characteristics. In these economic scenarios, on the other hand, it does not seem reasonable to postulate that the uninformed player will be unaware of the possibility that her opponent may be overconfident.[5] However, since the informed player is unaware of being overconfident, it seems odd to postulate that she will be

[4] Note that if one understands the beneficial effect of overconfidence on welfare as the result of incorporating high expectations on future achievements, then such an effect is captured in our framework.

[5] After all, the experimental evidence that shows the tendency for overconfidence is readily available to the managers of any firm, bank, or insurance company.

able to anticipate that her counterpart knows that she is overconfident. Thus the first part of the paper will focus on instances where the "uninformed" player knows that "informed" player may be mistaken in her self-perception, but the "informed" player is not aware of this.

The first conceptual problem is how to construct equilibrium predictions in these instances. On the one hand, in equilibrium the players should correctly conjecture the opponents' strategy. On the other hand, the uninformed player's equilibrium strategies necessarily depend on the awareness that her opponent may be overconfident, a possibility of which the informed player is not aware. Thus one can describe games where the informed player would deem the opponent's "equilibrium" strategies inconsistent with her understanding of the game. To substantiate the issue, consider an adverse-selection model, where the informed player is overconfident. Suppose that she may take a costly course of action intended to signal that her ability is high. If the opponent is aware of the sender's overconfidence, she will discount this signal to incorporate the mistaken self-perception, and adopt a low-profile strategy which is appropriate vis-a-vis a low-ability sender. If the sender correctly conjectures the receiver's strategy, she may want to take a less costly course of action. But this conjecture is in conflict with the belief that the receiver knows that the sender knows her ability.

In order to construct a framework that rules out this type of conflicts, we imagine a large population, ideally a continuum, of informed players. The distribution of ability in the population is common knowledge. While the players do not know whether a given individual is overconfident, each player knows the distribution of perceived ability in the population of informed players. The key assumption is that, while each informed player believes that her perception is correct, she also acknowledges that (the other) informed players are on average overconfident. Before playing the game, a single informed player is randomly anonymously chosen to play against the opponent. The opponent's strategy does not depend on the characteristics of the player she is facing, but only on the distribution of characteristics in the population of informed players. Thus no conflict between her understanding of the game and the anticipation her opponent's equilibrium strategy can arise in the mind of any informed player. In practice, she rationalizes the equilibrium choice of her opponent with the following consideration: "My opponent is acting as if we were not as good as we am, because there are many overconfident individuals in my population, and my opponent does not know that I am not one of them..."

Once formally defined equilibrium in the population games, we proceed with comparing overconfident and unbiased players' utility. Common wisdom holds confidence advantageous, so we may believe that overconfident players would be successful. *On the contrary, we show that in any equilibrium of any population game, a player cannot be made better-off by being overconfident.* In fact, while a player's choice depends on her perceived ability, her actual utility depends on her actual ability. Each overconfident player plays *as if* her ability were higher. In equilibrium she correctly anticipates the opponent's strategy. The equilibrium condition implies that the player's actual utility cannot be larger when she plays the equilibrium strategy of any player with ability different than hers, regardless of whether she does so in good faith, or with fraudulent intentions.

This result is not valid when comparing the player's payoff across different games, or across different equilibria of the same game. One may compare a game where informed players are likely to be unbiased, with a game where they are likely to be overconfident. It may be then be possible that an overconfident player of the second game fares better than unbiased players in the first game, because the opponent's strategy is modified by the knowledge that the informed player is more likely to be overconfident. Such a comparison *across different games*, however does not allow to conclude that overconfidence is beneficial, because any unbiased player playing in the second game would fare better than the overconfident player.

These results may be related to the literature on the value of information initiated by Hirshleifer (1971), and studied game-theoretically by Kamien, Tauman and Zamir (1990), and Neyman (1991), among others. It is often argued that, while it is well known that in a decision problem having less information cannot make a person better off, it may be possible in games that less informed players are better off, as different information may modify the opponents' strategy. Neyman (1991) however underlines that such a result depends on the comparison of the equilibrium of different games, and that if one compares equilibria of interactions embedded in the same fully-specified game, a player whose information is unilaterally refined cannot be worse off in equilibrium. These contributions focus on the value of a refinement of the information structure, where no conceptual difficulty takes place in constructing equilibrium play. We show that a result in the spirit of Neyman (1991) holds also for the case of possibly incorrect information.

The second part of the paper studies overconfidence and equilibrium for any description of high-order beliefs of overconfidence. We identify the knowledge descriptions that allow to construct an equilibrium concept immune to introspective conflicts, and for which the players correctly anticipate the opponents' strategies. This may be accomplished only in two instances. The first one, obviously, is the case where both players are unaware of overconfidence, and share the common belief that the informed player knows her own ability. The second one consists of the situation where the informed player, while believing that her perception is correct, knows that the opponent thinks that her perception is mistaken, the opponent knows this and so on. In such a case, we show that the only equilibrium concept immune to introspective conflict is the Bayesian equilibrium of an associated game with subjective priors. Again, we show that in any equilibrium, the actual payoff of an overconfident player is non-larger than the payoff of an unbiased player. *This means that overconfidence makes any player worse off in equilibrium also in the case that she knows that her opponent thinks that she is overconfident, or in the case that the opponent is not aware of overconfidence.*

We conclude our analysis by reformulating the model of overconfidence in the language of Mertens and Zamir (1985) universal types. We introduce a version of universal type that identifies a player's payoff-relevant personal characteristic, as well as her beliefs on her own and the opponents' characteristics, her beliefs over her own and the opponents' beliefs and so on. While it is possible to construct an equilibrium concept where the players' conjectures of each opponent type's strategy is correct, this does not solve the conceptual problem introduced by our adverse-selection model, but rather restates the problem in a different language. As each player's belief about her opponent's type is embedded in her universal type, one can construct games and knowledge descriptions identifying profiles of types that think they are facing some completely fictitious types. As a result the players will not be able to anticipate the opponents' strategies.

In order to substantiate this observation, consider again a signalling game where the sender is overconfident, and unaware that the receiver is aware of overconfidence, the receiver knows this and so on. The sender's type embeds the belief to play against a receiver who is unaware of overconfidence. Such a type adopts a high-profile strategy in response to the costly signal. Thus in "equilibrium" the sender believes that the receiver plays the high-profile strategy in response to the costly signal. In fact, our description of overconfidence knowledge identifies a type of receiver who discounts the sender signal and adopt a low-profile strategy in response to the costly signal. The sender will not be able to anticipate the receiver's choice.

In the recent wave of research on behavioral economics, some contributions have explored the economic consequences of overconfidence. Camerer and Lovallo (1999) conduct an experimental study that suggests a relationship between excess entry of new companies and entrepreneurs' optimism with respect to their own ability, relative to the ability of competitors. Benabou and Tirole (1999) consider the problem of search of information of overconfident individuals with time-inconsistent preferences. In the game between present and future selves, they show that an overconfident individual may strategically prefer to ignore some opportunities to gather information about their uncertain payoff, and can selectively decide to forget bad news. Flam and Risa (1998) study a search-theoretical problem where an individual chooses to take tests whose outcome depends on her own ability, and she is allowed to override failed tests. Thus overconfident players will eventually hold a higher status that unbiased ones, but because of longer periods of testing, their ex-ante discounted utility is smaller than unbiased players' utility. Benabou and Tirole (2000) characterize some incentive schemes that a principal may use to manipulate an agent's self-confidence to her own benefit. They describe situations where people criticize the performance of their partners as an instance of a battle for authority in the relationship, and discuss some humble self-presentation strategies of underconfident agents. Heifetz and Spiegel (2000) give an evolutionary account to the persistence of overconfident genes in a large class of games that includes both some games with strategic substitutes and some games with strategic complements.

This paper is presented as follows. The second section presents the set-up and an introductory example. The third section gives a precise account of the players' knowledge of the game and of equilibrium. The fourth section studies population games. The fifth section derives our results on utility comparison across overconfident and unbiased players. The sixth section generalized the analysis to any description of player's overconfidence. The seventh section reformulates the problem in the language of universal types. Omitted proofs are in Appendix.

2 Private Information and Overconfidence

In this section we first review the games of private information that are used in information economics to represent adverse-selection, screening or signaling interactions. Secondly, we show how to modify these games to account for both private information and mistaken self-perception.

For simplicity, say that there are only two players. Player 1 has private information about her own individual characteristics, summarized as $\theta \in \Theta$ (which we will denote for short as *ability*).[6] We may think that player 1 is a worker whose performance is being evaluated by

[6] It is important to stress that the ability θ need not be a number, it may indicate a vector or a distribution over personal characteristic, the modeler may consider appropriate, and each player j's strategy s_j need not

the employer, an entrepreneur who is applying for a credit line, an individual applying for an insurance. Her counterpart, player 2, is not informed of θ. It is common knowledge among the players that θ is distributed according to the distribution $\phi \in \Delta(\Theta)$, which is assumed to have full support. For simplicity, we assume that the set Θ is finite. Each player j's plan of action is denoted by $s_j \in S_j$, and we explicit nature's acts s_0 in the game, to allow for the possibility that the players learn about θ while playing the game. The nature's choice is denoted by $\mu \in \Delta(S_0)$, and it is common knowledge.[7] The strategy space is $S = S_0 \times S_1 \times S_2$, for simplicity we assume it finite. The players' payoffs $u : S \times \Theta \to \mathbb{R}^2$ depend on the players' choices s_1 and s_2, on nature's choice s_0 and on player 1's ability. In the case where player 1 knows the value of θ, this situation is represented the 2-player Bayesian game with common prior $G = (\Theta, \phi, S, \mu, u)$, the associated equilibrium concept of Bayesian Equilibrium is well understood.[8]

In order to represent the case where it is possible that player 1's perception is mistaken, first of all, it is necessary to distinguish between player 1's actual ability, denoted by θ_1, and her perception, denoted by θ_2.[9] Whenever player 1's personal characteristics are (θ_1, θ_2), she is informed of θ_2.[10] For any appropriately assigned order $>$ on the set Θ, we can define player 1 overconfident when $\theta_2 > \theta_1$. The relevant space of player 1's personal characteristics is thus $\boldsymbol{\Theta} = \Theta^2$ and we let $\boldsymbol{\phi} \in \Delta(\boldsymbol{\Theta})$ denote the (full-support) distribution over the pairs (θ_1, θ_2). Since we want to separate the economic effect of overconfidence from the psychological benefit of self-esteem, we assume that the players' utility depends only on the ability θ_1, and is independent of the perception θ_2. We denote the players' utility by $u : S \times \Theta_1 \to \mathbb{R}^2$, where

be a single action, but may be a complicated strategy, or even an infinite horizon policy.

[7] While the informed player may have a mistaken perception of herself, we want to allow for the possibility that she learns about her ability if presented with clear evidence by nature. Thus it must be the case that the player cannot be mistaken about the move of nature μ.

[8] Bayesian games are first defined in Harsanyi (1967), see Fudenberg and Tirole (1991) for the textbook treatment.

[9] For the sake of simplicity, we say that the player's perception θ_2 belongs to the same space as the player's ability θ_1. Alternatively, we could say that the player's perception is a measure that belongs to $\Delta(\Theta)$. It is easy to see that all our results can be extended under this alternative formulation, but that they would be more difficult to state and interpret.

[10] Note that this formulation does not necessarily require that the informed player is overconfident. This will not be the case in fact, whenever the player's characteristics θ_1 and θ_2 coincide.

Θ_1 denotes the first component of space $\mathbf{\Theta}$. From the game $G = (\Theta, \phi, S, \mu, u)$, we have obtained the expanded game $\mathbf{G} = (\mathbf{\Theta}, \boldsymbol{\phi}, S, \mu, u)$. Let $\boldsymbol{\phi}_1$ denote the marginal of $\boldsymbol{\phi}$ on the ability component, and $\boldsymbol{\phi}_2$ the marginal of $\boldsymbol{\phi}$ on the perception component.

Player 1 believes that her perception is always correct. This is equivalent to say that whenever the game $\mathbf{G} = (\mathbf{\Theta}, \boldsymbol{\phi}, S, \mu, u)$ is played, player 1 believes that she is instead playing the game $\mathbf{G}^0 = (\mathbf{\Theta}, \boldsymbol{\phi}^0, S, \mu, u)$,[11] where the distribution $\boldsymbol{\phi}^0 \in \Delta(\mathbf{\Theta})$ is derived from $\boldsymbol{\phi}$ according to the rule that for any $\theta_2 \in \Theta$, $\boldsymbol{\phi}^0(\theta_2, \theta_2) = \boldsymbol{\phi}_2(\theta_2)$.[12] It is immediate to see that the operator $(\cdot)^0$ maps Bayesian games into Bayesian games, and that $(\mathbf{G}^0)^0 = \mathbf{G}^0$. Player 2 is aware that player 1 may be overconfident. Specifically, we assume that player 2 knows that she is playing game $\mathbf{G}$. Player 1 thinks that player 2 thinks that player 1's perception is correct. This is equivalent to say that player 1 thinks that player 2 thinks that she is playing game $\mathbf{G}^0$. At the same time, we say that player 2 knows that player 1 believes that she knows that the game is $\mathbf{G}^0$, and so on...

Before giving a precise account for the players' knowledge in the game, and showing the conflict between self-perception and equilibrium, we introduce the issue by presenting a simple moral hazard game.

Example 1 Player 1 's ability θ_1 may be either high (θ_H) or low (θ_L). Player 2 would prefer a low-profile policy (y_L) if the opponent's ability is low, and a high-profile policy (y_H) if the opponent's ability is high. Before player 2 chooses her policy, player 1 may either send a costly signal s_H, or a default signal s_L. The low-ability sender prefers to send s_L regardless of the receiver's choice. Player 1 may be overconfident (but not underconfident): we assume that $\boldsymbol{\phi}(\theta_L, \theta_H) > 0$, and $\boldsymbol{\phi}(\theta_H, \theta_L) = 0$.

This game, sequential in essence, may be represented by the Bayesian game $\mathbf{G} = (\mathbf{\Theta}, \boldsymbol{\phi}, S, \mu, u)$, where $\Theta = \{\theta_L, \theta_H\}$, S_0 is a singleton set, $S_1 = \{s_L, s_H\}$, and $S_2 = \{y_L y_L, y_H y_L,$

[11] It is not enough to say that each player with characteristics (θ_1, θ_2) believes that $\theta_1 = \theta_2$, because this allows for the possibility that she may think that, if her perception had been $\theta_2' \neq \theta_2$, then she would have believed that her ability would have been $\theta_1' \neq \theta_2'$.

[12] It should be immediately pointed out that this description is *not* equivalent to a Bayesian game with subjective priors, where player 1's prior is $\boldsymbol{\phi}^0$, and player 2's prior is $\boldsymbol{\phi}$. We will formally study Bayesian games with subjective priors in a later section.

$y_L y_H$, $y_H y_H\}$: the first (respectively second) component denotes the receiver's choice after receiving the signal s_L (respectively s_H). The probabilities $\boldsymbol{\phi}$ and μ, as well as the utility functions u_1 and u_2 are immediately constructed from the above description. We assume that player 2 knows that she is playing the game $\mathbf{G} = (\boldsymbol{\Theta}, \boldsymbol{\phi}, S, \mu, u)$, player 1 thinks that she is playing game $\mathbf{G}^0 = (\boldsymbol{\Theta}, \boldsymbol{\phi}^0, S, \mu, u)$ and that player 2 thinks that she is playing $\mathbf{G}^0$, and so on...

In order to demonstrate the conflict between mistaken self-perception and equilibrium, we make these specific assumptions on the players' payoffs:

$$u_1(s_L, y_H, \theta_H) < u_1(s_H, y_L, \theta_H) < u_1(s_L, y_L, \theta_H) < u_1(s_H, y_H, \theta_H), \tag{1}$$

$$u_1(s_H, y, \theta_L) < u_1(s_L, y, \theta_L), \quad \text{for any } y \in \{y_L, y_H\}; \tag{2}$$

$$\{y_H\} = \arg\max_y \boldsymbol{\phi}_2(\theta_H) u_2(s_L, y, \theta_H) + \boldsymbol{\phi}_2(\theta_L) u_2(s_L, y, \theta_L), \tag{3}$$

$$\{y_L\} = \arg\max_y \boldsymbol{\phi}(\theta_H, \theta_H) u_2(s, y, \theta_H) + \boldsymbol{\phi}(\theta_L, \theta_H) u_2(s, y, \theta_L), \text{ and} \tag{4}$$

$$u_2(s, y_L, \theta_H) < u_2(s, y_H, \theta_H), \; u_2(s, y_H, \theta_L) < u_2(s, y_L, \theta_L), \quad \text{for } s \in \{s_L, s_H\}. \tag{5}$$

By Condition (2) the low-perception sender (θ_L, θ_L) plays s_L. Given that, by Conditions (4) and (5), player 2 plays y_H upon observing s_H if and only if the overconfident (θ_L, θ_H) sender plays s_L, and the high-ability (θ_H, θ_H) sender plays s_H. But if the receiver plays y_H after s_H, the (θ_L, θ_H) sender plays s_H, by Condition (1). Thus, if the players correctly anticipate the opponents' choice, player 2 plays $y_L y_L$ (by Conditions (4) and (5)), and player 1, regardless of her perception plays s_L. In the mind of the overconfident sender, however the opponent's strategy $y_L y_L$ is in conflict with what she thinks about the game. By Condition (3), she thinks that player 2 should play y_H after receiving s_L, rather than y_L.[13] $\diamond$

3 Knowledge of the Game

In order to formally represent the players' knowledge of the game they are playing, we introduce an underlying (compact metric) state space Ω, and the associated Borel σ-algebra on Ω, denoted by $\mathcal{B}(\Omega)$. Let the nature's choice on the probability space $(\Omega, \mathcal{B}(\Omega))$ be denoted

[13] Notice that, given Conditions (4) and (5), Condition (3) is satisfied whenever $\phi(\theta_L, \theta_H)$ is large enough.

as p. The nature selects the game $\mathbf{G}(\omega)$ as a function of the state of the world. We thus introduce the measurable surjective relation $\tilde{\mathbf{G}} : \omega \mapsto \mathbf{G}$, and the event $[\mathbf{G}] = \{\omega | \tilde{\mathbf{G}}(\omega) = \mathbf{G}\}$. Given the game $\mathbf{G}(\omega)$, the nature then selects the individual characteristics (θ_1, θ_2) according to the distribution $\phi(\omega)$, and the strategy s_0 according to the distribution $\mu(\omega)$.[14] The players' information with respect to the game may be represented by the (non-necessarily truthful) information structures $P_j : \Omega \to \mathcal{B}(\Omega)$, $j = 1, 2$.[15] We denote by information model, the collection $\mathcal{I}= (\Omega, P_1, P_2, p)$. It will also be useful to introduce the knowledge operators $K_j : 2^\Omega \to 2^\Omega$, such that, for any $E \in 2^\Omega$, $K_j E = \{\omega | P_j(\omega) \subseteq E\}$ for $j = 1, 2$. The common knowledge operator $CK : 2^\Omega \to 2^\Omega$ is introduced by defining the sequence of operators $\{K^n\}_{n \geq 0}$, where for any n, $K^n : 2^\Omega \to 2^\Omega$, and specifically, $K^0 E = K_1 E \cap K_2 E$, and $K^n E = K_1(K^{n-1}E) \cap K_2(K^{n-1}E)$ for any $n \geq 1$. Let the event "E is common knowledge" be defined as $CKE = \cap_{n \geq 0} K^n E$. Let $(P_1 \wedge P_2)$ be the finest common coarsening of P_1 and P_2; it is known that $(P_1 \wedge P_2)(\omega) \subseteq E$ if and only if $\omega \in CKE$.[16]

Say that the players are playing an arbitrary game $\mathbf{G}$. In order to identify the description of players' knowledge of the game informally introduced in the previous section, we define the events $\kappa_1^0[\mathbf{G}] = K_1[\mathbf{G}^0] \cap K_2[\mathbf{G}^0]$, $\kappa_2^0[\mathbf{G}] = K_1[\mathbf{G}^0] \cap K_2[\mathbf{G}]$, and iteratively, for any $n \geq 1$, $\kappa_2^n[\mathbf{G}] = K_2 \kappa_2^{n-1}[\mathbf{G}]$, and $\kappa_1^n[\mathbf{G}] = K_1 \kappa_1^{n-1}[\mathbf{G}]$.[17] We are interested in describing the players' (equilibrium) strategies only for states ω that belongs to the event

$$E[\mathbf{G}] = [\mathbf{G}] \cap K_1[\mathbf{G}^0] \cap K_2[\mathbf{G}] \cap [\cap_{n \geq 2} (\kappa_1^n[\mathbf{G}] \cap \kappa_2^n[\mathbf{G}])].$$

We need to show that there exist information models such that the event $E[\mathbf{G}]$ is non-empty. If this were not the case, the task of describing equilibrium play on $E[\mathbf{G}]$ would be

[14] While this implies that the description of the state of the world Ω is incomplete, as it does not capture all uncertainty in the game, it is easy to see how to expand the state space to account for the nature's choice of θ_2 and s_0. We adopt this "reduced" formulation of the state space to simplify the analysis, and to underline that we are focusing our attention on the players' knowledge of the game.

[15] It is straightforward to show that if $P_j(\omega) \in \mathcal{B}$, then the restriction of $\mathcal{B}$ onto $P_j(\omega)$ is a σ-algebra (in fact the Borel σ-algebra on $P_j(\omega)$).

[16] An introduction to the formal representation of knowledge may be found in Dekel and Gul (1997).

[17] It should be noted that the information correspondence P_1 identifies the knowledge of player 1 of game $\mathbf{G}$ only in an ex-ante sense, i.e. before that she is assigned her characteristics (θ_1, θ_2). Her knowledge when she takes her choice is expressed by the conjuction of $P_1(\omega)$ with her perception θ_2.

meaningless.

Lemma 1 *There exist information models* $\mathcal{I}= (\Omega, P_1, P_2, p)$ *such that, for any game* $\mathbf{G}$, *the event* $E[\mathbf{G}]$ *is non-empty.*

The players' strategies in the game are a function of the underlying state of the world. Player 2's strategy consists of the function $\tilde{\sigma}_2 : \Omega \to \Delta(S_2)$, measurable with respect to the information structure P_2. When choosing her action, player 1 is informed of $\theta_2 \in \Theta$. Her strategy in the game is thus expressed by the function $\tilde{\sigma}_1 : \Omega \to \Delta(S_1)^\Theta$, measurable with respect to the information structure P_1. We denote by σ_1 any arbitrary element of $\Delta(S_1)^\Theta$ and by σ_2 any arbitrary element of $\Delta(S_2)$. The event that player 1, respectively player 2, play a given strategy σ_1, respectively σ_2, are denoted by the notations $[\sigma_1] = \{\omega | \ \tilde{\sigma}_1(\omega) = \sigma_1\}$ and $[\sigma_2] = \{\omega | \tilde{\sigma}_2(\omega) = \sigma_2\}$.[18] Player j is rational if she maximizes her utility on the basis of her information. Formally, we define the events:

$$[R_2] \;=\; \left\{ \omega \middle| \forall s_2, \ \tilde{\sigma}_2(s_2, \omega) > 0 \ \Rightarrow \ s_2 \in \arg\max_{s_2'} E\left[u_2(s_1, s_2', \theta_1) | P_2(\omega)\right] \right\} \tag{6}$$

$$[R_1] \;=\; \left\{ \omega \middle| \forall(\theta_2, s_1), \ \tilde{\sigma}_1(s_1|\theta_2, \omega) > 0 \ \Rightarrow \ s_1 \in \arg\max_{s_1'} E\left[u_2(s_1', s_2, \theta_1) | \theta_2, P_1(\omega)\right] \right\}. \tag{7}$$

We now turn to the construction of the equilibrium concept. A minimal requirement for an equilibrium construction is that the players are rational and correctly anticipate each other's strategies in the game (in an ex-ante sense). If one does not impose further restriction, she adopts the stand-point of declaring herself agnostic so as to how these anticipations are formed. Given the description of knowledge $E[\mathbf{G}]$, we will define as *naive equilibrium* any profile $\sigma = (\sigma_1, \sigma_2)$ such that, upon knowing that player 2 plays σ_2, player 1 rationally chooses σ_1, and viceversa. Define the events $[\sigma] = [\sigma_1] \cap [\sigma_2]$, and $[R] = [R_1] \cap [R_2]$.

Definition 1 *For any arbitrary information model* $\mathcal{I}$ *and game* $\mathbf{G}$, *the profile* σ *is a naive equilibrium for* $E[\mathbf{G}]$ *if the event* $E[\mathbf{G}] \cap [R] \cap [\sigma] \cap K^0[\sigma]$ *is non-empty.*

[18] Unlike Auman and Brandeburger (1995), in this formulation player i does not know the specific action a_i she takes at a certain state ω, but only the mixed strategy σ_i. It assumed that after choosing the state ω (which identifies which game $\mathbf{G}$ is played, and which mixed strategies σ are taken by the players), the nature moves again in the game $\mathbf{G}$, operating the randomizing device identified by σ. It will be seen that this formulation greatly simplifies our analysis.

We now show that it is possible to construct an information model $\mathcal{I}$, such that for any game $\mathbf{G}$, the naive equilibria for $E[\mathbf{G}]$ coincide with the subjective equilibria of the Bayesian game with subjective priors $\mathbf{G}' = (\boldsymbol{\Theta}, \boldsymbol{\phi}, \boldsymbol{\phi}^0, S, \mu, u)$, where $\boldsymbol{\phi}$ identifies both the move of nature and the prior of player 2, and $\boldsymbol{\phi}^0$ identifies the prior of player 1.

Definition 2 *The strategy profile σ is a subjective equilibrium of $\mathbf{G}' = (\boldsymbol{\Theta}, \boldsymbol{\phi}, \boldsymbol{\phi}^0, S, \mu, u)$ if for any θ_2 and s_1, it is the case that $\sigma_1(s_1|\theta_2) > 0$ only if*

$$s_1 \in \arg\max_{s_1' \in S_1} \sum_{\theta_1 \in \Theta} \sum_{s_2 \in S_2} \sum_{s_0 \in S_0} u_1(s_0, s_1', s_2, \theta_1)\mu(s_0)\sigma_2(s_2)\frac{\boldsymbol{\phi}^0(\theta_1, \theta_2)}{\phi_2(\theta_2)}, \tag{8}$$

and for any s_2, it is the case that $\sigma_2(s_2) > 0$ only if

$$s_2 \in \arg\max_{s_2' \in S_2} \sum_{(\theta_1,\theta_2) \in \boldsymbol{\Theta}} \sum_{s_1 \in S_1} \sum_{s_0 \in S_0} u_2(s_0, s_1, s_2', \theta_1)\mu\left(s_0\right)\sigma_1(s_1|\theta_2)\boldsymbol{\phi}(\theta_1, \theta_2). \tag{9}$$

Proposition 1 *There is an information model $\mathcal{I}= (\Omega, P_1, P_2)$, such that for any game $\mathbf{G} = (\boldsymbol{\Theta}, \boldsymbol{\phi}, S, \mu, u)$, any strategy profile σ is a naive equilibrium for $E[\mathbf{G}]$ if and only if σ is a subjective equilibrium of the game $\mathbf{G}' = (\boldsymbol{\Theta}, \boldsymbol{\phi}, \boldsymbol{\phi}^0, S, \mu, u)$.*[19]

It is well known that subjective equilibrium exists in all finite Bayesian games with subjective priors. It follows that there is an information model $\mathcal{I}$ such that for any finite game $\mathbf{G}$, there exists a naive equilibrium for $E[\mathbf{G}]$. The concept of naive equilibrium however is unappealing for the purposes of a theory of equilibrium in the presence of mistaken self-perception because the players cannot introspectively rationalize the choice that they impute to their opponents. Specifically, suppose that player 1 knows that player 2 is rational. When playing a strategy σ_1 and believing that player 2 plays strategy σ_2, player 1 may ask herself the question: "What would player 2 play if she knew that we am playing σ_1?" It may well

[19]If the underlying game $\mathbf{G}$ has complete information, it follows that σ is a naive equilibrium if and only if it is a Nash Equilibrium of $\mathbf{G}$. Aumann and Brandenburger (1995) show that, in 2-player games, Nash Equilibrium conjecture follow from public knowledge of payoffs, rationality, and conjectures, where a player j's conjecture is a conditional distribution on the actions of her opponent, player l. Our requirement that each player j knows that player l surely plays σ_l is stronger than just requiring that player j's belief over player l's action coincides with σ_l. This allows to obtain Nash Equilibrium without requiring public knowledge of rationality (see also the "preliminary observation" in Aumann and Brandenburger 1995 at page 1167).

be that the answer to this question is a strategy σ'_2 different than σ_2. It is then unclear why player 1 should maintain her belief that player 2 is playing σ_2. We define as "sophisticated" equilibrium, a strategy profile σ that is immune to this conflict due to introspective reasoning. For any game $\mathbf{G}$, and description of knowledge of the game $E[\mathbf{G}]$, we say that the profile σ is a sophisticated equilibrium if it is possible that the players' knowledge of the game is described by $E[\mathbf{G}]$, and at the same time it is common knowledge that the players are rational, and that the play is σ.

Definition 3 *For any arbitrary information model $\mathcal{I}$ and game $\mathbf{G}$, the profile σ is a sophisticated equilibrium for $E[\mathbf{G}]$ if the event $E[\mathbf{G}] \cap [R] \cap [\sigma] \cap CK[[R] \cap [\sigma]]$ is non-empty.*

We conclude the section by showing that one can construct games such that there does not exists any equilibrium immune to introspective conflict, one such game is presented in Example 1.

Proposition 2 *There are games $\mathbf{G}$ such that, for any information model $\mathcal{I} = (\Omega, P_1, P_2)$, there does not exist any sophisticated equilibrium for $E[\mathbf{G}]$.*

4 Population Games

Imagine a continuous population of informed players indexed in $i \in I = [0, 1]$. Each pair (θ_1, θ_2) is interpreted as the actual characteristics of some individuals in the population i. The characteristics are assigned by the (measurable) function $\zeta : I \to \mathbf{\Theta}$, where $\zeta_1(i)$ denotes the ability of sender i, and $\zeta_2(i)$ denotes her perception. Given the assignment ζ, the distribution $\boldsymbol{\phi}(\zeta) : \mathbf{\Theta} \to [0, 1]$ over the pairs (θ_1, θ_2) is derived according to the rule:

$$\boldsymbol{\phi}(\zeta)(\theta_1, \theta_2) = \nu\{i : \zeta(i) = (\theta_1, \theta_2)\}, \tag{10}$$

where ν denotes the Lebesgue measure. The sequence of moves in population games is as follows. At the first period, a single informed player $i \in I$ is randomly chosen by nature according to the uniform distribution on $[0, 1]$. At the second period, player i plays against nature and player 2 a Bayesian game $\mathbf{G} = (\mathbf{\Theta}, \phi, S, \mu, u)$. The strategy space of each player i coincides with S_1, her utility is denoted by $u^i(s, \zeta)$ and coincides with $u_1(s, \zeta_1(i))$ for each strategy profile $s \in S$. We have associated to each Bayesian game $\mathbf{G}$ a population game $\mathbf{\Gamma} = (\mathbf{\Theta}, I, \zeta, \mu, S, u)$.

As in the previous section, we represent the players' knowledge of the game by means of an information model $\mathcal{I}$, and a measurable surjective relation $\tilde{\mathbf{\Gamma}} : \omega \mapsto \mathbf{\Gamma}$, the relation $\tilde{\mathbf{\Gamma}}$ includes the relation $\tilde{\zeta} : \omega \mapsto \zeta$. For any $\mathbf{\Gamma}$ and ζ, we define the events $[\mathbf{\Gamma}] = \{\omega | \tilde{\mathbf{\Gamma}}(\omega) = \mathbf{\Gamma}\}$, and $[\boldsymbol{\phi}(\zeta)] = \{\omega | \boldsymbol{\phi}(\tilde{\zeta}(\omega)) = \boldsymbol{\phi}(\zeta)\}$. The nature first chooses the population game $\mathbf{\Gamma}$, then she selects player i according to ν, and finally she takes μ in the game $\mathbf{G}$. The knowledge and common knowledge operators are constructed in analogous way as in the previous section. For any game $\mathbf{\Gamma} = (\mathbf{\Theta}, I, \zeta, \mu, S, u)$, we restrict attention to information models $\mathcal{I}$ for which the collection $(\mathbf{\Theta}, I, \mu, S, u)$ is common knowledge on $[\mathbf{\Gamma}]$.

We want to represent instances where player 2 is not able to distinguish the identity of the players in the pool I, and has no information about the assignment of any single player $i \in I$, but at the same time, she knows the aggregate distribution of the individual characteristics (θ_1, θ_2). Thus we restrict attention to information models satisfying the following assumption.

Assumption 1 *The information model* $\mathcal{I}= (\Omega, (P^i)_{i \in I}, , P_2, p)$ *is such that for any game* $\mathbf{\Gamma}$ *and any state* $\omega \in [\mathbf{\Gamma}]$, $P_2(\omega) \subseteq [\boldsymbol{\phi}(\zeta)]$; *and such that for any* ν*-preserving isomorphism* $\iota : \mathcal{B}[0,1] \to \mathcal{B}[0,1]$, *any set* $B \in \mathcal{B}[0,1]$, *and any set* $\Theta' \subseteq \mathbf{\Theta}$, $p(\zeta(\iota(B)) \in \Theta' | P_2(\omega)) = p(\zeta(B) \in \Theta' | P_2(\omega))$.

The key assumption of our construction is that, while each informed player i believes that her ability $\zeta_1(i)$ coincides with perception $\zeta_2(i)$, she also acknowledges that (the other) informed players are on average overconfident. Specifically we restrict attention to information models that satisfy the following assumption.

Assumption 2 *For any game* $\mathbf{\Gamma}$, *the information model* $\mathcal{I}= (\Omega, P^i, P_2, p)$ *is such that for any* $\omega \in [\mathbf{\Gamma}]$, *and any* $i \in I$, $P^i(\omega) \subseteq \{\omega' : \tilde{\zeta}(\omega')(i) = (\zeta_2(i), \zeta_2(i))\} \cap [\boldsymbol{\phi}(\zeta^i)]$, *where the function* $\zeta^i : [0,1] \to \mathbf{\Theta}$ *is such that* $\zeta^i(j) = \zeta(j)$ *for any* $j \neq i$, *and* $\zeta^i(i) = (\zeta_2(i), \zeta_2(i))$.

Under these assumptions, by construction, any player i is overconfident (and unaware of this) whenever $\zeta_1(i) \neq \zeta_2(i)$. It is immediate to see, in fact, that for any perception θ_2, at any state ω such that $\tilde{\zeta}_2(\omega')(i) = \theta_2$, it is the case that $\omega \in K_i\{\omega|\tilde{\zeta}_1(\omega')(i) = \theta_2\}$, regardless of player i's actual ability $\tilde{\zeta}_1(\omega')(i)$. Nevertheless, the following Lemma verifies that there are information models satisfying our assumptions, where the players always share common knowledge of the aggregate distribution of individual characteristics.

Lemma 2 *There is an information model* $\mathcal{I}= (\Omega, (P^i)_{i\in I}, P_2, p)$ *satisfying Assumptions 1 and 2 such that for any game* $\mathbf{\Gamma} = (\mathbf{\Theta}, I, \zeta, \mu, S, u)$, *it is the case that* $[\mathbf{\Gamma}] \subseteq [\boldsymbol{\phi}(\zeta)] \cap CK[\boldsymbol{\phi}(\zeta)]$.

For any $i \in I$, we let the function $\tilde{\sigma}^i : \Omega \to \Delta(S_1)$, measurable with respect to the information structure P^i, be the strategy of player i. As in the previous section, player 2's strategy is described by $\tilde{\sigma}_2 : \Omega \to \Delta(S_2)$ measurable with respect to P_2. We denote by $\boldsymbol{\sigma}_1$ any arbitrary element of $\Delta(S_1)^I$ and by σ_2 any arbitrary element of $\Delta(S_2)$. Since the players in population I are indistinguishable, it is natural to restrict attention to symmetric strategy profiles, where all players with the same ability and assessment take the same strategy. Formally, a strategy profile $\boldsymbol{\sigma} \in \Delta(S_1)^I \times \Delta(S_2)$ is symmetric if for any pair $(i, j) \in I^2$, it is the case that $\sigma^i = \sigma^j$ whenever $\zeta_2(i) = \zeta_2(j)$. The definitions of the events $[\boldsymbol{\sigma}_1]$, $[\sigma_2]$, $[\sigma]$, $[R_2]$, $[R_s^i]$ for any i, and $[R]$ are immediately extended from the analogous definition in the previous section, and so are the definitions of naive and sophisticated equilibrium. Because of Lemma 2, for each game $\mathbf{\Gamma}$, we are interested in the naive and sophisticated equilibria associated with the event $[\mathbf{\Gamma}]$, under information models satisfying Assumptions 1 and 2.

Given any population game $\mathbf{\Gamma} = (\mathbf{\Theta}, I, \zeta, \mu, S, u)$, we derive the 2-player Bayesian game $\mathbf{G}' = (\mathbf{\Theta}, \boldsymbol{\phi}, \boldsymbol{\phi}^0, S, \mu, u)$, where for any pair (θ_1, θ_2), the nature's choice (and player 2's prior) is $\boldsymbol{\phi}(\theta_1, \theta_2) = \nu\{i : \zeta(i) = (\theta_1, \theta_2)\}$, and the prior of player 1 is $\boldsymbol{\phi}^0(\theta_2, \theta_2) = \nu\{i : \zeta_2(i) = \theta_2\}$. Each symmetric strategy profile $\boldsymbol{\sigma}$ of a game $\mathbf{\Gamma}$ identifies a unique strategy profile σ' of $\mathbf{G}'$ according to the rule $\sigma'_2 = \sigma_2$, and $\sigma'_1(\cdot|\theta_2) = \sigma^i$ if $\zeta_2(i) = \theta_2$. Up to equivalence classes, each strategy profile σ' of $\mathbf{G}'$ identifies a symmetric strategy profile $\boldsymbol{\sigma}$ of $\mathbf{\Gamma}$. Proposition 3 below shows that, within the restrictions imposed by Assumptions 1 and 2, it is possible to construct information models, such that the symmetric naive equilibria of any game $\mathbf{\Gamma}$ coincide with the subjective Bayesian equilibria of the associated game $\mathbf{G}'$. Moreover, the symmetric naive equilibria of any game $\mathbf{\Gamma}$ coincide with the symmetric sophisticated equilibria of $\mathbf{\Gamma}$. This immediately implies the existence of sophisticated symmetric equilibrium in all population games with finite characteristics and strategy spaces. Thus Proposition 3 allows us to conclude that the framework of population games resolves the conflict between overconfidence and equilibrium identified by Example 1.

Proposition 3 *There is an information model* $\mathcal{I}= (\Omega, (P^i)_{i\in I}, P_2, p)$ *satisfying Assumptions 1 and 2, such that for any game* $\mathbf{\Gamma} = (\mathbf{\Theta}, I, \zeta, S, \mu, u)$, *the set of symmetric sophisticated equilibria for* $[\mathbf{\Gamma}]$ *coincides with the set of symmetric naive equilibria for* $[\mathbf{\Gamma}]$*, which is isomorphic (up to equivalence classes) to the set of subjective equilibria of the game* $\mathbf{G}' = (\mathbf{\Theta}, \boldsymbol{\phi}(\zeta), (\boldsymbol{\phi}(\zeta))^0, S, \mu, u)$.

5 Utility Comparisons

Now that we are endowed with a set-up that settles the conflict between mistaken beliefs and equilibrium, we can compare the equilibrium payoffs of overconfident players with unbiased ones. We will show that in equilibrium, the payoff of an overconfident player cannot be larger than the payoff of an unbiased one.

Intuitively, consider any symmetric equilibrium $\boldsymbol{\sigma}$ of any population game. Take two players i and j whose actual ability is θ. Say that player i correctly perceives her ability, and that player j is overconfident and perceives her ability to be $\theta' > \theta$. In equilibrium, both players correctly anticipate player 2's strategy σ_2. Each player's expected utility when choosing her strategy depends on her perceived utility, and therefore on her perceived ability; but her *actual* equilibrium utility depends on her actual ability, not on the perceived one. The overconfident player j chooses a strategy σ^j that maximizes the expected equilibrium payoff of a player whose actual ability is θ', while the unbiased player i chooses a strategy σ^i that maximizes the actual expected equilibrium payoff of a player whose ability is θ. Since both player's actual ability is in fact θ, it cannot be the case that the overconfident player j fares strictly better than the unbiased player i. As this result holds for any equilibrium of any population game, it holds a fortiori for any refinement motivated by robustness or by a specific sequential structure of the underlying game.

In order to formalize the result, for any game $\mathbf{\Gamma} = (\mathbf{\Theta}, I, \zeta, S, \mu, u)$, and any symmetric equilibrium $\boldsymbol{\sigma}$, we introduce the notation $u^i(\boldsymbol{\sigma})$ which identifies player i's actual payoff (in ex-ante terms) at the equilibrium $\boldsymbol{\sigma}$. It is easy to see that

$$\text{for any } i,\ u^i(\boldsymbol{\sigma}) = \sum_{s_2\in S_2}\sum_{s_1\in S_1}\sum_{s_0\in S_0} u^i(s_0, s_1, s_2, \zeta)\mu(s_0)\sigma^i(s_1)\sigma_2(s_2).$$

Proposition 4 *In any symmetric equilibrium $\boldsymbol{\sigma}$ of any population game $\mathbf{\Gamma} = (\mathbf{\Theta}, I, \zeta, S, \mu, u)$, for any level of ability $\theta \in \Theta$, and any pair of players (i, j) such that $\zeta(i) = (\theta, \theta)$ and $\zeta(j) = (\theta, \theta')$ with $\theta' \neq \theta$, it must be the case that $u^i(\boldsymbol{\sigma}) \geq u^j(\boldsymbol{\sigma})$.*

Proof. Consider any symmetric equilibrium $\boldsymbol{\sigma}$ of any population game $\mathbf{\Gamma} = (\mathbf{\Theta}, I, \zeta, S, \mu, u)$. By Proposition 3, $\boldsymbol{\sigma}$ identifies a naive equilibrium σ for $E[\mathbf{G}]$. For any $\omega \in [\mathbf{\Gamma}] \cap [R_1] \cap K^i[\sigma_2]$, player i plays strategy $\tilde{\sigma}^i(\cdot|\omega)$, such that $\tilde{\sigma}^i(s_1|\omega) > 0$ only if

$$s_1 \in \arg\max_{s_1'\in S_1} \sum_{s_2\in S_2}\sum_{s_0\in S_0} u_1(s_0, s_1', s_2, \theta_2)\mu(s_0)\sigma_2(s_2). \tag{11}$$

It follows that any $s_1 \in Supp(\sigma^i)$ must satisfy Condition (11).

For any arbitrary level of activity θ, pick any pair of players $(i, j) \in I^2$ such that

$$\zeta_1(i) = \zeta_2(i) = \zeta_1(j) = \theta, \text{ and } \zeta_2(j) = \theta', \text{ where } \theta' \neq \theta.$$

Since $\zeta_1(i) = \zeta_1(j) = \theta$, it follows that for any profile of pure strategies s, $u^i(s, \zeta) = u^j(s, \zeta) = u_1(s, \theta)$.

Condition (11) implies that for any $s_1 \in Supp(\sigma^i)$,

$$\sum_{s_2\in S_2}\sum_{s_0\in S_0} u_1(s_0, s_1, s_2, \theta)\mu(s_0)\sigma_2(s_2) \geq \sum_{s_2\in S_2}\sum_{s_0\in S_0} u_1(s_0, s_1', s_2, \theta)\mu(s_0)\sigma_2(s_2), \text{ for any } s_1' \in S_1,$$

this condition holds a fortiori for any $s_1' \in Supp(\sigma^j)$. It follows that

$$\begin{aligned} u^i(\boldsymbol{\sigma}) &= \sum_{s_1\in Supp(\sigma^i)} \sigma^i(s_1) \sum_{s_2\in S_2}\sum_{s_0\in S_0} u_1(s_0, s_1, s_2, \theta)\mu(s_0)\sigma_2(s_2) \\ &\geq \sum_{s_1'\in Supp(\sigma^j)} \sigma^j(s_1') \sum_{s_2\in S_2}\sum_{s_0\in S_0} u_1(s_0, s_1', s_2, \theta)\mu(s_0)\sigma_2(s_2) = u^j(\boldsymbol{\sigma}). \end{aligned}$$

■

Proposition 4 compares the utility of overconfident and unbiased players for any fixed game and equilibrium. When comparing payoffs across different games, one may specify a game where informed players are likely to be unbiased, and a game where they are likely to be overconfident. It may be then be possible that an overconfident player of the second game fares better than unbiased players in the first game, because the opponent's strategy is modified by the knowledge that the informed player is more likely to be overconfident. Such a comparison across different games, however, does not allow to conclude that overconfidence is beneficial, because any unbiased player playing in the second game would fare better than this overconfident player. Similar considerations hold also for utility comparison across different equilibria of the same game.[20] Nevertheless it is interesting to notice that an increment of the likelihood that informed players are overconfident may make all players in the population i (overconfident and unbiased) better off. To substantiate this point, we present the following example.

Example 2 Consider a family of population games $\mathbf{\Gamma}_\alpha$ indexed in $\alpha \in (0,1)$, $\alpha \neq 1/2$. For each α, say that $\Theta = \{\theta_L, \theta_H\}$, and that for any $i \in I$, it is the case that $\zeta_2(i) = \theta_L$. At the same time however $\zeta_2(i) = \theta_H$ for any $i \geq \alpha$, and $\zeta_2(i) = \theta_L$ for any $i < \alpha$. Thus each player $i \geq \alpha$ is overoptimistic. The strategy set S_0 is a singleton set, $S_1 = \{A, B\}$ and $S_2 = \{C, D\}$. The payoffs are as represented below.

θ_L	C	D
A	1,0	0,2
B	0,2	2,0

θ_H	C	D
A	2,0	0,2
B	0,2	2,0

[20] When comparing the outcome of a Bayesian game involving only a single unbiased player with the outcome of a game involving a single overconfident player, as we have already pointed out, the meaning of equilibrium analysis is unclear. But if one accepts the concept of naive equilibrium, she obtains that, while the overconfident player of the second game may fare better than the unbiased player of the first game, it is also true that a (hypothetical) unbiased player called to play the second game would fare even better (cf. Proposition 8).

For any fixed α, the game $\mathbf{\Gamma}_\alpha$ has a unique equilibrium. For $\alpha < 1/2$, in equilibrium, the overoptimistic players from population i must mix between A and B. Thus the unique equilibrium $\boldsymbol{\sigma}$, is such that $\sigma_2(C) = 1/2$, that $\sigma^i(A) = 0$ whenever $i \in [0, \alpha)$, and $\sigma^i(A) = 1/[2(1-\alpha)]$ if $i \in (\alpha, 1]$. The players' payoffs are $u_2(\boldsymbol{\sigma}) = 1$, $u^i(\boldsymbol{\sigma}) = 1$ if $i \in [0, \alpha)$, and $u^i(\boldsymbol{\sigma}) = [3-4\alpha]/[4(1-\alpha)]$ when $i \in (\alpha, 1]$. For $\alpha > 1/2$ the unbiased players from population i must play a mixed strategy in equilibrium. Thus the unique equilibrium $\boldsymbol{\sigma}$, is such that $\sigma_2(C) = 2/3$, that $\sigma^i(B) = 1/[2\alpha]$ whenever $i \in [0, \alpha)$, and that $\sigma^i(B) = 0$ if $i \in (\alpha, 1]$. The players' payoffs are $u_2(\boldsymbol{\sigma}) = 1$, and $u^i(\boldsymbol{\sigma}) = 2/3$ if $i \in [0, \alpha)$, $u^i(\boldsymbol{\sigma}) = 1/3$ when $i \in (\alpha, 1]$.

For any arbitrary $\alpha' \in (0, 1/4)$ and $\alpha'' \in (1/2, 1)$, the overoptimistic players of game $\mathbf{\Gamma}_{\alpha'}$ achieve a higher utility than the unbiased players of game $\mathbf{\Gamma}_{\alpha''}$. Also, the payoff of overoptimistic players is increasing in $1-\alpha$, the fraction of overoptimistic players in the population i. In this game, an increment of the likelihood that the players in population i are mistakenly too optimistic (as long as the resulting likelihood is above $3/4$) makes all players better off. $\diamond$

6 General Knowledge Descriptions

This section studies self-perception and equilibrium in 2-player games, for general descriptions of the players' knowledge of the game. In any such a scenario, the players play a game $\mathbf{G} = (\boldsymbol{\Theta}, \boldsymbol{\phi}, S, \mu, u)$, but, in order to account for player 1's possibly mistaken perception, we say that player 1 believes to play the game $\mathbf{G}^0$. A description of the player's knowledge of the game is generated by the events $[\mathbf{G}]$, and $[\mathbf{G}^0]$, and by the iterated application of the operators K_1, K_2, as well as complementation and intersection. Given the space Ω, and the relation $\tilde{\mathbf{G}} : \omega \mapsto \mathbf{G}$, for any game $\mathbf{G} = (\boldsymbol{\Theta}, \boldsymbol{\phi}, S, \mu, u)$, we consider the space $\Omega_{\mathbf{G}} = [\mathbf{G}] \cup [\mathbf{G}^0]$. We introduce the algebra[21] $\mathcal{A}^1_{\mathbf{G}} = \{\emptyset, [\mathbf{G}], [\mathbf{G}^0], \Omega_{\mathbf{G}}\}$, and for any $n \geq 1$, the algebra $\mathcal{A}^n_{\mathbf{G}}$ generated by $\mathcal{A}^{n-1}_{\mathbf{G}} \cup \{K_i E |\ E \in \mathcal{A}^{n-1}_{\mathbf{G}},\ i = 1, 2\}$. The algebra that includes

[21] An algebra of Ω is a collection of subsets of Ω that contains Ω, that is closed under complementation and finite intersection.

all the descriptions of players' knowledge of the game $\mathbf{G}$ is $\mathcal{A}_{\mathbf{G}} = \cup_{n=1}^{\infty} \mathcal{A}_{\mathbf{G}}^{n}$.[22] It is known (cf. Aumann 1999, or Hart Heifetz and Samet 1996) that not all the lists in $\mathcal{A}_{\mathbf{G}}$ are consistent: there are lists of events whose intersection is empty for any information model $\mathcal{I}$. To avoid triviality, for any game $\mathbf{G}$, we restrict attention to lists $l_{\mathbf{G}} \in \mathcal{A}_{\mathbf{G}}$, such that there is a model $\mathcal{I}$ for which the event $E_l(\mathbf{G}) = \cap_{n=0}^{\infty} l_{\mathbf{G}}^{n}$ is non-empty. Since it is also known that whether a list $l_{\mathbf{G}}$ generated by the event $[\mathbf{G}]$ is consistent or not depends only on the combinations of knowledge and logic operators, and is independent of the generating event, we will drop the subscript from the notation $l_{\mathbf{G}}$, with the understanding that the notation l identifies the list $l_{\mathbf{G}}$ when in conjunction with a specific game $\mathbf{G}$.

First we extend Proposition 1 to any instance where the informed player is overconfident, and her opponent is aware of this, regardless of the players' high-order knowledge of overconfidence. We introduce the collection of lists $A = \{l \in \mathcal{A} |$ for any $\mathbf{G}$, $\emptyset \neq E_l(\mathbf{G}) \subseteq [\mathbf{G}] \cap K_1[\mathbf{G}^0] \cap K_2[\mathbf{G}]\}$. Recall that for any game $\mathbf{G}$, the game $\mathbf{G}'$ denotes the associated game with subjective priors.

Proposition 5 *For any list $l \in A$, there is an information model $\mathcal{I}$ such that for any game $\mathbf{G}$, the profile σ is a naive equilibrium for $E_l(\mathbf{G})$ if and only if σ is a subjective equilibrium of $\mathbf{G}'$.*

Secondly we show that if player 2 is unaware that player 1 may be overconfident, then the naive equilibria of any game $\mathbf{G}$ coincide with the Bayesian equilibria of the game $G = (\Theta, \boldsymbol{\phi}_2, S, \mu, u)$. We let $U = \{l \in \mathcal{A} |$ for any $\mathbf{G}$, $\emptyset \neq E_l(\mathbf{G}) \subseteq [\mathbf{G}] \cap K_1[\mathbf{G}^0] \cap K_2[\mathbf{G}^0]\}$.

Proposition 6 *For any list $l \in U$, there is an information model $\mathcal{I}$ such that for any game $\mathbf{G}$, the profile σ is a naive equilibrium for $E_l(\mathbf{G})$ if and only if σ is a Bayesian equilibrium of G.*

The final and most important result of this section identifies under which conditions sophisticated and naive equilibrium coincide. This may occur in two instances. First, it may

[22] An example of a list of events in $\mathcal{A}_{\mathbf{G}}$ is the list $l_{\mathbf{G}} = \{[\mathbf{G}],\ K_1[\mathbf{G}^0],\ K_2[\mathbf{G}],\ \left(\kappa_1^2[\mathbf{G}], \kappa_2^2[\mathbf{G}]\right),\ ...,\ \left(\kappa_1^n[\mathbf{G}], \kappa_2^n[\mathbf{G}]\right), ...\}$, which represent the knowledge description studied in the third section.

be the case that, while they are truly playing game $\mathbf{G}$, the players share common knowledge that they are playing the game $\mathbf{G}^0$, so that not only the informed player is unaware to be overconfident, but also her opponent is unaware that she could be overconfident, and so on. In this case, sophisticated, naive and Bayesian equilibrium all coincide. Secondly, and more interestingly, it may be the case that the players "agree to disagree" on the game that they are playing. I.e. player 1 is overconfident and unaware of that, player 2 knows that player 1 is overconfident, player 1 thinks that player 2 thinks that player 1 is overconfident, and so on. In this case, naive and sophisticated equilibria of game $\mathbf{G}$ coincide with the subjective equilibria of the associated game with subjective priors $\mathbf{G}'$. For any other description of the players' knowledge of overconfidence, there are games that do not have any sophisticated equilibrium. This result shows that a model of Bayesian equilibrium with subjective priors is appropriate to represent overconfidence if and only if the overconfident player, while unaware of being overconfident, is aware that the opponent thinks that she is overconfident, and so on.

We denote by l^0 the list $l \in \mathcal{A}$ such that for any game $\mathbf{G}$, $E_l(\mathbf{G}) = [\mathbf{G}] \cap CK[\mathbf{G}^0]$, and by l^* the list l such that for any game $\mathbf{G}$, $E_l(\mathbf{G}) = [\mathbf{G}] \cap \left(\cap_{n \geq 0} \bar{K}^n[\mathbf{G}] \right)$, where $\bar{K}^0[\mathbf{G}] = K_1[\mathbf{G}^0] \cap K_2[\mathbf{G}]$, and for any $n > 0$, $\bar{K}^n[\mathbf{G}] = K_1 \bar{K}^{n-1}[\mathbf{G}] \cap K_2 \bar{K}^{n-1}[\mathbf{G}]$.[23]

Proposition 7 *For $l \in \{l^0, l^*\}$ there is a model $\mathcal{I}$ such that for any game $\mathbf{G}$, the profile σ is a sophisticated equilibrium for $E_l(\mathbf{G})$ if and only if σ is a naive equilibrium for $E_l(\mathbf{G})$. For any other list l, and model $\mathcal{I}$, there exist games $\mathbf{G}$ where a sophisticated equilibrium does not exist for $E_l(\mathbf{G})$.*

One may be interested in utility comparisons between overconfident and unbiased players in this general environment. It is easy to show that in any naive equilibrium of any game, regardless of the players' knowledge of the game, if player 1 is overconfident, she obtains a non-larger payoff than if she is unbiased. As shown by Proposition 7, the concept of naive equilibrium coincides with the concept of sophisticated equilibrium when both players

[23] It is immediate to show that l^0 and l^* are consistent.

are unaware of overconfidence, or when they "agree to disagree" on overconfidence. This means that overconfidence makes any player worse off in equilibrium also in the case that the opponent is not aware of overconfidence, and in the case that the opponent is aware of overconfidence, but the informed player knows this. For any strategy profile σ, any ability θ_1, and any perception θ_2, let the actual utility of the informed player with characteristics (θ_1, θ_2) be

$$u_1(\sigma, \theta_1, \theta_2) = \sum_{s_2 \in S_2} \sum_{s_1 \in S_1} \sum_{s_0 \in S_0} u_1(s_0, s_1, s_2, \theta_1)\mu(s_0)\sigma(s_1|\theta_2)\sigma_2(s_2).$$

Proposition 8 *For any list $l \in A$, any information model $\mathcal{I}$, and any game $\mathbf{G}$, in any naive equilibrium for $E_l(\mathbf{G})$, for any level of ability θ_1, and any perception θ_2, it must be the case that $u_1(\sigma, \theta_1, \theta_1) \geq u_1(\sigma, \theta_1, \theta_2)$.*

7 Games with Universal Types

This section shows how universal types can be used to describe self-perception.[24] We propose a straightforward extension of Brandenburger and Dekel (1993) construction that includes in a player's type also individual objective characteristics (i.e. ability), as well as beliefs over own ability and opponents' ones, and higher-order beliefs. Since the construction of universal types is well understood, proofs and unnecessary calculations are omitted, and made available upon request to the author.

For any player $j = 1, 2$, let the space of j's ability be a complete separable metric space Θ_j. Iteratively set $X_1 = \Theta_1 \times \Theta_2$, and for any $n \geq 1$, $X_{n+1} = X_n \times [\Delta(X_n)]^2$. Let a type t_j be a hierarchy $(\delta_{0j}, \delta_{1j}, \delta_{2j}...) \in \Theta_j \times (\times_{n=1}^{\infty} \Delta(X_n))$, and define $T_{j0} = \Theta_j \times (\times_{n=1}^{\infty} \Delta(X_n))$. For each space X_n, the above definition of type includes more than one possible specification of beliefs, the first one is identified by δ_{nj}, another is the marginal of δ_{n+1j} on X_n. In order to avoid a conflicting definition, we restrict attention to coherent types. Formally, a type t_j is coherent if for any $n \geq 1$, the marginal distribution projected by δ_{n+1j} on X_n coincides with δ_{nj}. Let

[24] The concept of universal types has been first introduced by Mertens and Zamir (1985), see also Brandeburger and Dekel (1993) and Epstein and Wang (1996).

the set of coherent types be T_{j1}. One can show that each coherent type uniquely identifies a system of beliefs with respect to her own and the opponent's type. The sequence $(\delta_{1j}, \delta_{2j}, ...)$, in fact, identifies consistent probability measures on the denumerable sequence $\{X_n\}$ of finite cylinders. Since each cylinder is a Polish space, t_j identifies through Kolmogoroff Extension Theorem a unique (canonical) probability measure on the space $\Delta(T_{j0} \times T_{-j0})$.[25] Completing the relation with the identity on the space Θ_j, this construction identifies a (unique) homeomorphism $f_j : T_{j1} \to \Theta_j \times \Delta(T_{j0} \times T_{-j0})$.[26] This result, however, permits any coherent type to identify a belief that she or her opponent is not coherent. To avoid this inconsistency, we restrict attention to types that satisfy "common knowledge of coherency." Formally, for any $n \geq 1$, let $T_{jn+1} = \{t \in T_{j1} : f_j(t)(T_{jn} \times T_{-jn}) = 1\}$, and let the universal type space of player j be $T_j = \cap_{n=1}^{\infty} T_{jn}$. In order to verify that each universal type identifies a unique belief over the state of nature, and an opponent's universal type, one can show that there is a (unique) homeomorphism $g_j : T_j \to \Theta_j \times \Delta(T_j \times T_{-j})$, generated by f_j.[27]

Given the space of individual abilities $\Theta = \Theta_1 \times \Theta_2$, and the space of universal types $T = T_1 \times T_2$, in order to obtain a fully specified game, we need to specify a nature's prior $p \in \Delta(T)$, a strategy space $S = S_0 \times S_1 \times S_2$, a move of nature μ, and payoffs $u : S \times \Theta \to \mathbb{R}^2$. A strategy in the game $\Gamma = \{\Theta, T, p, S, u\}$ is an profile $\sigma = (\sigma_1, \sigma_2)$, where for each j, the function $\sigma_j : T_j \to \Delta(S_j)$ is measurable. The actual utility of the players is expressed by the function $u : \Theta \times S \to \mathbb{R}^2$; for any mixed strategy σ, and any player j, the actual utility of type t_j when playing against t_{-j} is:

$$u_j(t, \sigma) = \sum_{s_1 \in S_1} \sum_{s_2 \in S_2} u_j(\delta_0, s) \sigma_1(s_1|t_1) \sigma_2(s_2|t_2) \mu(s_0).$$

Some types t_j of player j include a mistaken belief about their ability, and their perceived payoff may differ from their actual payoff. Specifically, for any mixed strategy σ, any player

[25] See for instance Dudley (1999).

[26] This result is an extension of Proposition 1 in Brandenburger and Dekel (1993), where the reader can find additional details.

[27] This result is a simple extension of Proposition 2 in Brandenburger and Dekel (1993).

j of type t_j perceives that her utility is:

$$\tilde{u}_j(t_j, \sigma) = \int \sum_{s_1 \in S_1} \sum_{s_2 \in S_2} u_j(g_j(t_j)(\theta), s)\sigma_1(s_1|t_1)\sigma_2(s_2|t_2)\mu(s_0)dg_j(t).$$

For this construction, the most appropriate definition of equilibrium requires that all types choose a payoff-maximizing strategy, and not only those are selected with positive prior probability.[28]

Definition 4 *An equilibrium of game Γ is a profile $\sigma = (\sigma_1, \sigma_2)$, where for each player $j \in \{1, 2\}$, the strategy σ_j is such that $\sigma_j(s_j|t_j) > 0$ only if*

$$s_j \in \arg\max_{s'_j \in S_j} \int \sum_{s_{-j} \in S_{-j}} \sum_{s_0 \in S_0} u_j(g_j(t_j)(\theta), s'_j, s_{-j}, s_0)\sigma_{-j}(s_{-j}|t_{-j})\mu(s_0)dg_j(t).$$

In order to show how this construction relates to our analysis of overconfidence, we reconsider the game $\mathbf{G} = (\Theta, \boldsymbol{\phi}, S, \mu, u)$ presented in Example 1. Recall that player 1's ability belongs to the set $\Theta = \{\theta_L, \theta_H\}$, that player 2's ability is irrelevant, that $\boldsymbol{\phi}(\theta_L, \theta_H) = 0$, that $S_1 = \{s_L, s_H\}$, that $S_2 = \{y_L y_L, y_H y_L, y_L y_H, y_H y_H\}$, and that we are interested to the equilibrium play associated with the knowledge description identified by the event $E[\mathbf{G}] =[\mathbf{G}] \cap K_1[\mathbf{G}^0] \cap K_2[\mathbf{G}] \cap [\cap_{n\geq 2} (\kappa_1^n[\mathbf{G}] \cap \kappa_2^n[\mathbf{G}])]$.

In the language of universal types, one can show that the event $E[\mathbf{G}]$ identifies a type-distribution $p = (p_1, p_2) \in \Delta(T)$ such that for any $(\theta_1, \theta_2) \in \boldsymbol{\Theta}$, p assigns probability $\boldsymbol{\phi}(\theta_1, \theta_2)$ to the pair of types (t_1, t_2) such that $\delta_{01} = \theta_1$, $\delta_{11} = \delta(\theta_2)$, $\delta_{12} = \boldsymbol{\phi}_1$, where the notation $\delta(\cdot)$ identifies the distribution degenerate on $\cdot$, and that the high-order beliefs are recursively defined as follows. Let $\delta_{21} = \delta_{11} \cdot \delta(\delta_{11}) \cdot \delta(\boldsymbol{\phi}_2)$, $\delta_{22} = \delta(\delta_{12}) \cdot \boldsymbol{\phi}(\theta_1, \delta_{11})$, where the last term assigns probability $\boldsymbol{\phi}(\theta_1, \theta_2)$ to the state $\{\theta_1, \delta(\theta_2)\}$, and for any $n \geq 2$, $\delta_{n+1,1} = \delta_{n1} \cdot \delta(\delta_{n1}) \cdot \delta(\delta_{n2} \cdot \boldsymbol{\phi}^0(\theta_2, \delta_{11}, ..., \delta_{n1}))$, $\delta_{n+1,2} = \delta(\delta_{n2}) \cdot \boldsymbol{\phi}(\theta_1, \delta_{11}, ..., \delta_{n-1,1})$, where the terms $\boldsymbol{\phi}^0$ and $\boldsymbol{\phi}$ are derived as before. For simplicity, say that $\boldsymbol{\phi}(\theta_L, \theta_H) = 1$, and for

[28] In the context of correlated equilibrium, Brandenburger and Dekel 1987 introduce the distinction between ex-ante equilibrium (which requires that each player maximizes ex-ante payoff), and a-posteriori equilibrium (which requires that also null- probability types choose payoff-maximizing strategies).

any given $\theta_2 \in \Theta$, let $t_1[\theta_2]$ denote the type $t_1 \in Supp(p_1)$ such that $\delta_{11} = \delta(\theta_2)$, note that p_2 is degenerate.

The key observation is that for any $\theta_2 \in \Theta$, type $t_1[\theta_2]$ identifies through g_1 the belief that player 1 is of type $t'_1[\theta_2]$ and that player 2 is of type t'_2, where $\delta'_{01} = \theta_2$, $\delta'_{11} = \delta(\theta_2)$, $\delta'_{12} = \boldsymbol{\phi}_2$, $\delta'_{21} = \delta_{11} \cdot \delta(\delta_{11}) \cdot \delta(\boldsymbol{\phi}_2)$, $\delta'_{22} = \delta(\delta_{12}) \cdot \boldsymbol{\phi}_2(\theta_2, \delta_{11})$, and for any $n \geq 2$, $\delta'_{n+1,1} = \delta'_{n1} \cdot \delta(\delta_{n1}) \cdot \delta(\delta'_{n2} \cdot \boldsymbol{\phi}_2(\theta_2, \delta'_{11}, ..., \delta'_{n1}))$, $\delta'_{n+1,2} = \delta(\delta'_{n2}) \cdot \boldsymbol{\phi}_2(\theta_2, \delta'_{11}, ..., \delta'_{n-1,1})$. Since the types $t'_1[\theta_L]$, $t'_1[\theta_H]$, and t'_2 identify Bayesian game with common prior $G = (\Theta, \boldsymbol{\phi}_2, S, u)$, any Bayesian equilibrium σ of game Γ must be such that $\sigma(t'_2)(y_L y_H) = 1$, that $\sigma_1(t'_1[\theta_L])(s_L) = 1$, and that $\sigma_1(t'_1[\theta_H])(s_H) = 1$. Player 2, on the other hand, is of type t_2 which identifies the belief that player 1 is of type $t_1[\theta_L]$ with probability $\phi(\theta_L, \theta_L)$, and of type $t_1[\theta_H]$ with probability $\phi(\theta_L, \theta_H)$. Since she believes that $\theta_1 = \theta_L$ with probability 1, in any Bayesian Equilibrium she must play $\sigma(t_2)(y_L y_L) = 1$.

The formulation of the game with universal types allows us to construct a Bayesian equilibrium that predicts that at any state $\omega \in E[G]$, player 1 plays s_H when believing her ability to be high, regardless of the fact that player 2 will respond to that choice by playing y_L. This occurs because state ω identifies a type of player 1 that believes to play against a type of player 2 which is different than the type of player 2 which is identified by state ω. It follows that player 1 cannot anticipate the strategy played by player 2. In this sense, this reformulation does not solve the problem introduced by Example 1 and formalized by Proposition 2. While it is true that in any Bayesian Equilibrium of the game with universal type, the assignment of strategies to types is common knowledge among the players, it is the case that player 1 cannot anticipate player 2's choice because she believes to play against a completely fictitious type of player 2.

8 Appendix

Proof of Lemma 1. Pick an information model $\mathcal{I} = (\Omega, P_1, P_2, p)$ such that for any game G, and any $\omega \in [G]$, $P_2(\omega) \subseteq [G]$, and $P_1(\omega) \subseteq [G^0]$. It is immediate to see that $[G] \subseteq K_2[G] \cap K_1[G^0]$. To show the remaining part of the claim, we proceed by induction.

First notice that $[G] \subseteq K_2[G]$ and $[G] \subseteq K_1[G^0]$, imply $[G] \subseteq K_2[K_1[G^0] \cap K_2[G]] = \kappa_2^1[G]$; while $[G] \subseteq K_1[G^0]$, and $[G^0] \subseteq K_2[G^0]$ imply $[G] \subseteq K_1[K_1[G^0] \cap K_2[G^0]] = \kappa_2^1[G]$. Secondly, notice that $[G^0] \subseteq K_1[G^0]$. Thus for any $n \geq 1$, it follows that $[G] \subseteq \kappa_1^{n-1}[G]$ implies $[G] \subseteq \kappa_1^n[G]$. Also $[G^0] \subseteq K_2[G^0]$, and $[G] \subseteq K_2[G]$, thus for any $n \geq 1$, it follows that $[G] \subseteq \kappa_2^{n-1}[G^0]$ implies $[G] \subseteq \kappa_2^n[G^0]$. The result is then obtained by induction. ■

Proof of Proposition 1. Take the information model I such for any game $G = (\Theta, \phi, S, \mu, u)$, and any strategy profile σ of G, the event $[G] \cap [\sigma]$ is non-empty, and such that for any $\omega \in [G] \cap [\sigma]$, $P_1(\omega) = [\sigma] \cap [G^0]$, and that $P_2(\omega) = [\sigma] \cap [G]$. By Lemma 1, $E[G]$ is non-empty. Also, $[G] \cap [\sigma] \subseteq K^0[\sigma]$, and thus the profile σ is a naive equilibrium if and only if the event $[G] \cap [\sigma] \cap [R]$ is non-empty. Suppose that this is the case: there is a ω such that $\tilde{G}(\omega) = G$, $\tilde{\sigma}(\omega) = \sigma$, and that for every s_2, $\tilde{\sigma}_2(s_2, \omega) > 0$ only if

$$s_2 \in \arg\max_{s_2'} E\left[u_2(s_1, s_2', \theta_1) | P_2(\omega)\right], \tag{12}$$

and for every θ_2, and s_1, $\tilde{\sigma}_1(s_1|\theta_2, \omega) > 0$ only if

$$s_1 \in \arg\max_{s_1'} E\left[u_2(s_1', s_2, \theta_1) | \theta_2, P_1(\omega)\right]. \tag{13}$$

Since $\tilde{\sigma}(\omega) = \sigma$, it follows that, by plugging the expression $P_2(\omega) = [\sigma] \cap [G]$ in Condition (12) we obtain that σ must satisfy Condition (8), and by plugging the expression $P_1(\omega) = [\sigma] \cap [G^0]$ in Condition (13), we obtain that σ must satisfy Condition (9). It follows that σ is a subjective equilibrium of $G' = (\Theta, \phi, \phi^0, S, \mu, u)$. Conversely, if σ is a subjective equilibrium of G', then it must satisfy Conditions (8) and (9). It follows that for any $\omega \in [G] \cap [\sigma]$, it is the case that $\omega \in [R]$. Thus σ is a naive equilibrium of G. ■

Proof of Proposition 2. Consider the game $G = (\Theta, \phi, S, \mu, u)$ presented in Example 1. For any $\omega \in K_2[G] \cap [R_2]$, for any distribution $\xi_2(\omega) \in \Delta(S_1)^\Theta$, where the conjecture ξ_2 is supposed to be measurable with respect to P_2, it is the case that $\tilde{\sigma}_2(s_2, \omega) > 0$ only if

$$s_2 \in \arg\max_{s_2'} \sum_{(\theta_1, \theta_2) \in \boldsymbol{\Theta}} \sum_{s_1 \in S_1} u_2(s_1, s_2, \theta_1) \xi_2(s_1(\theta_2)|\omega) \boldsymbol{\phi}(\theta_1, \theta_2).$$

It follows that for any $\omega \in K_2[G] \cap [R_2]$, it is the case that $\tilde{\sigma}_2(y_L y_L|\omega) = 1$. So that letting $\sigma_2(y_L y_L) = 1$, we obtain that $K_2[G] \cap [R_2] \subseteq [\sigma_2]$. So for any strategy $\sigma_2' \in \Delta(S_2)$, $\sigma_2' \neq \sigma_2$, it follows that $K_2[G] \cap [R_2] \cap [\sigma_2'] = \emptyset$.

For any $\omega \in K_1[G^0] \cap [R_1]$, for any distribution $\xi_1(\omega) \in \Delta(S_2)$, where the conjecture ξ_1 is supposed to be measurable with respect to P_1, for any $\theta_2 \in \Theta$, it is the case that

$\tilde{\sigma}_1(s_1(\theta_2), \omega) > 0$ only if

$$\begin{aligned} s_1 &\in \arg\max_{s_1'} \sum_{\theta_1 \in \Theta} \sum_{s_2 \in S_2} u_1(s_1', s_2, \theta_1) \xi_1(s_2|\omega) \frac{\boldsymbol{\phi}^0(\theta_1, \theta_2)}{\phi_2(\theta_2)} \\ &= \arg\max_{s_1'} \sum_{s_2 \in S_2} u_1(s_1', s_2, \theta_1) \xi_1(s_2|\omega). \end{aligned}$$

It follows that for any $\omega \in K_1[G^0] \cap [R_1] \cap K_1[\sigma_2]$, for any $\theta_2 \in \{\theta_L, \theta_H\}$, it is the case that $\tilde{\sigma}_1(s_L|\theta_2, \omega) = 1$. So letting $\sigma_1(s_L|\theta_2) = 1$, for any $\theta_2 \in \{\theta_L, \theta_H\}$, we obtain that $K_1[G^0] \cap [R_1] \cap K_1[\sigma_2] \subseteq [\sigma_1]$. So for any strategy $\sigma_1' \in \Delta(S)^\Theta$, $\sigma_1' \neq \sigma_1$, it follows that $K_1[G^0] \cap [R_1] \cap K_1[\sigma_2] \cap [\sigma_1'] = \emptyset$.

For any $\omega \in K_2[G^0] \cap [R_2] \cap K_2[\sigma_1]$, instead, it is the case that $\sigma_2(s_2, \omega) > 0$ only if

$$s_2 \in \arg\max_{s_2'} \boldsymbol{\phi}_2(\theta_H) u_2(s_L, s_2', \theta_H) + \boldsymbol{\phi}_2(\theta_L) u_2(s_L, s_2', \theta_L).$$

It follows that $\tilde{\sigma}_2(y_L y_L, \omega) = 0$ for any $\omega \in K_2[G^0] \cap [R_2] \cap K_2[\sigma_1]$. Therefore, it must be the case that $K_1 K_2[G^0] \cap K_1[R_2] \cap K_1 K_2[\sigma_1] \cap K_1[\sigma_2] = \emptyset$. It follows that for any strategy pair (σ_1, σ_2) it is the case that

$$K_2[\mathbf{G}] \cap K_1[\mathbf{G}^0] \cap K_1 K_2[\mathbf{G}^0] \cap [R] \cap [\sigma_2] \cap [\sigma_1] \cap K_1[\sigma_2] \cap K_1[R_2] \cap K_1 K_2[\sigma_1] = \emptyset.$$

Thus, for any strategy pair σ, the event $E[G] \cap [R] \cap [\sigma] \cap CK([R] \cap [\sigma])$ is empty. ∎

Proof of Proposition 2. By Condition (10), and by Assumption 2, it follows that for any pair $(\theta_1, \theta_2) \in \Theta$,

$$\begin{aligned} \boldsymbol{\phi}(\zeta^i)(\theta_1, \theta_2) &= \nu\{j \in [0,1] : \zeta^i(j) = (\theta_1, \theta_2)\} \\ &= \nu\{j \in [0,i) \cup (i,1] : \zeta^i(j) = (\theta_1, \theta_2)\} \\ &= \nu\{j \in [0,i) \cup (i,1] : \zeta(j) = (\theta_1, \theta_2)\} = \boldsymbol{\phi}(\zeta)(\theta_1, \theta_2). \end{aligned}$$

Now pick an information model I such that P_2 satisfies Assumption 1, and such that for any $i \in I$, it is the case that $P^i(\omega) = P_2(\omega) \cap \{\omega' : \tilde{\zeta}(\omega')(i) = (\zeta_2(i), \zeta_2(i))\}$; since this set is non-empty I is well-defined, and satisfies Assumption 2. We have shown that for any $\omega \in [\Gamma]$, for any i, $P^i(\omega) \subseteq P_2(\omega)$, it follows that $P_2(\omega) \wedge (\wedge_{i \in I} P^i(\omega)) = P_2(\omega) \subseteq [\phi(\zeta)]$, where the latter relation follows by Assumption 1. This implies that for any $\omega \in [\Gamma]$, it is the case that $\omega \in CK[\phi(\zeta)]$. ∎

Proof of Proposition 3. Take the information model I such for any game $\Gamma = (\Theta, I, \zeta, S, \mu, u)$, and any symmetric strategy profile σ of Γ, the event $[\Gamma] \cap [\sigma]$ is non-empty, such that for any $\omega \in [\Gamma] \cap [\sigma]$, $P_2(\omega) = [\sigma] \cap [\phi(\zeta)]$, such that for any ν-measure preserving isomorphism $\iota : B[0,1] \rightarrow B[0,1]$, for any set $B \in B[0,1]$, and for any set $\Theta' \subseteq \Theta$, it is the case that $p(\zeta(\iota(B)) \in \Theta' | P_2(\omega)) = p(\zeta(B) \in \Theta' | P_2(\omega))$, and such that for any i, $P_s^i(\omega) = \{\omega' : \tilde{\zeta}(\omega')(i) = (\zeta_2(i), \zeta_2(i))\} \cap P_2(\omega)$. Since for any i, $[\phi(\zeta)] = [\phi(\zeta^i)]$, it follows that I satisfies Assumptions (1), and (2).

Fix a game $\Gamma = (\Theta, I, \zeta, S, \mu, u)$, and say that σ' is a subjective equilibrium of the game $G' = (\Theta, \phi(\zeta), (\phi(\zeta))^0, S, \mu, u)$. Up to equivalence classes σ' identifies a symmetric strategy profile σ of Γ. Say that $\omega \in [\Gamma] \cap [\sigma]$. It is the case that $\omega \in [R_2]$ if and only if the strategy $\tilde{\sigma}_2(\cdot, \omega) = \sigma_2(\cdot)$ is such that $\tilde{\sigma}_2(s_2, \omega) > 0$ only if:

$$s_2 \in \arg\max_{s_2' \in S_2} E_\zeta \left[\int_I \sum_{s_1 \in S_1} \sum_{s_0 \in S_0} u_2(s_0, s_1, s_2', \zeta_1(i)) \mu(s_0) \sigma^i(s_1|\omega) d\nu(i) \middle| P_2(\omega) \right]. \tag{14}$$

By construction, the strategy $\sigma_1' \in \Delta(S_1)^\Theta$ satisfies the rule $\sigma_1(s_1|\theta_2) = \sigma^i(s_1|\omega)$ whenever $\zeta_2(i) = \theta_2$. Since player i is selected from pool i according to the uniform distribution on $[0,1]$, since for any isometry $\iota : B[0,1] \rightarrow B[0,1]$, $p(\zeta(\iota(B)) \in \Theta' | P_2(\omega)) = p(\zeta(B) \in \Theta' | P_2(\omega))$ for any set $B \in B[0,1]$, and for any set $\Theta' \subseteq \Theta$, and since $P_2(\omega) \subseteq [\phi(\zeta)]$, it follows that the pair distribution $\phi(\zeta)$ is a sufficient statistic of player 2's information on the assignment ζ. Thus Equation (14) can be summarized by aggregating the players in i across the characteristics (θ_1, θ_2). Substituting in the Equation, the expressions for $\phi(\zeta)$ and for σ_1', we obtain:

$$s_2 \in \arg\max_{s_2' \in S_2} \sum_{(\theta_1, \theta_2) \in \boldsymbol{\Theta}} \sum_{s_1 \in S_1} \sum_{s_0 \in S_0} u_2(s_0, s_1, s_2', \theta_1) \mu(s_0) \sigma_1'(s_1|\theta_2) \boldsymbol{\phi}(\theta_1, \theta_2). \tag{15}$$

This condition coincides with Condition (8), which σ' satisfies by definition. It follows that $[\Gamma] \cap [\sigma] \subseteq [R_2]$. With respect to the players in i, we can say that for any i, and any $\omega \in [\Gamma] \cap [\sigma]$, it is the case that $\omega \in [R^i]$ if and only if player i plays $\tilde{\sigma}^i(\cdot|\omega)$ such that $\tilde{\sigma}^i(s_1|\omega) > 0$ only if

$$\begin{aligned} s_1 &\in \arg\max_{s_1' \in S_1} \sum_{\theta_1 \in \Theta} \sum_{s_2 \in S_2} \sum_{s_0 \in S_0} u^i(s_0, s_1', s_2, \theta_1) \mu(s_0) \sigma_2(s_2|\omega) p(\zeta_1(i) = \theta_1 | \zeta_2(i)) \\ &= \arg\max_{s_1' \in S_1} \sum_{s_2 \in S_2} \sum_{s_0 \in S_0} u_1(s_0, s_1', s_2, \zeta_2(i)) \mu(s_0) \sigma_2(s_2|\omega). \end{aligned} \tag{16}$$

By construction, $\tilde{\sigma}^i(\cdot|\omega) = \sigma^i(\cdot) = \sigma'(\cdot|\zeta_2(i))$, thus σ_1 satisfies Condition (16) for every i if and only if σ' satisfies Condition (9) for every θ_2, which is the case by definition. It follows that $[\Gamma] \cap [\sigma] \subseteq [R_1]$.

The above arguments have shown that $[\Gamma] \cap [\sigma] \subseteq [R]$. Since it is also the case that $[\Gamma] \subseteq [\phi(\zeta)]$, it follows that for any $\omega \in [\Gamma] \cap [\sigma]$, $P_2(\omega) \cap [R] \neq \emptyset$. To show that for any i, it is also the case that $P_s^i(\omega) \cap [R] \neq \emptyset$, pick an arbitrary i, and any arbitrary state ω' such that $\tilde{\zeta}(\omega') = \zeta^i$. By construction, $\omega' \in P_s^i(\omega)$. Moreover, by construction, the 2-player subjective-priors Bayesian game associated with the game $\tilde{\Gamma}(\omega')$, coincides with G'. It follows that $[\tilde{\Gamma}(\omega')] \cap [\sigma] \subseteq [R]$. Given our construction of I, the event $[\tilde{\Gamma}(\omega')] \cap [\sigma]$ is non-empty, it follows that there is a state $\omega' \in P_s^i(\omega) \cap [R]$.

Because of the above two results, we can refine the information structures $(P^i)_{i\in I}$ and P_2, by defining $\hat{P}_2(\omega) = P_2(\omega) \cap [R]$, and $\hat{P}_s^i(\omega) = P_s^i(\omega) \cap [R]$, for every $i \in I$. The event $[R]$ incorporates in its definition the information structures $(P^i)_{i\in I}$ and P_2, we define the event $[\hat{R}]$ which is to be understood as the event that the players are rational relative to the structures $(\hat{P}^i)_{i\in I}$ and $\hat{P}_2$. We want to show that $[\Gamma] \cap [\sigma] \cap [R] = [\Gamma] \cap [\sigma] \cap [\hat{R}]$. Since $[\Gamma] \cap [\sigma] \subseteq [R_2]$, it is rational for player 2 to play σ_2 on the event $[\Gamma] \cap [\sigma]$. It follows that for any player i, the information that player 2 is rational does not add anything to the belief that she plays σ_2, and so for any $\omega \in [\Gamma] \cap [\sigma]$, $E\left[u^i(s_1, s_2', \theta_1)|P^i(\omega)\right] = E\left[u^i(s_1, s_2', \theta_1)|\hat{P}^i(\omega)\right]$. Conversely, since for any i, $[\Gamma] \cap [\sigma] \subseteq [R^i]$, it is the case that for any $\omega \in [\Gamma] \cap [\sigma]$, $E\left[u_2(s_1, s_2', \theta_1)|P_2(\omega)\right] = E\left[u_2(s_1, s_2', \theta_1)|\hat{P}_2(\omega)\right]$.

Since by construction, for any i, $\hat{P}^i(\omega) \subseteq \hat{P}_2(\omega)$, it follows that $\hat{P}_2(\omega) \wedge \left(\wedge_{i\in I}\hat{P}^i(\omega)\right) = \hat{P}_2(\omega) \subseteq [\hat{R}] \cap [\sigma]$ that is to say $\omega \in C\hat{K}[[\hat{R}] \cap [\sigma]]$. Wrapping up, we have shown that $[\Gamma] \cap [\sigma] \subseteq [\hat{R}] \cap C\hat{K}[[\hat{R}] \cap [\sigma]]$, since by construction, $[\Gamma] \cap [\sigma]$ is non-empty, we conclude that if σ' is a subjective equilibrium of the game associated with Γ, then σ is a sophisticated equilibrium of Γ.

The fact that if the profile σ is a sophisticated equilibrium of Γ under any information model I (and in particular information model $\hat{I}$), then it is also a naive equilibrium of Γ under I trivially follows from the fact that $[\Gamma] \cap [R] \cap [\sigma] \cap CK\left([R] \cap [\boldsymbol{\sigma}]\right) \subseteq [\Gamma] \cap [R] \cap [\sigma] \cap K^0[\sigma]$.

We are left to show that if σ is a naive equilibrium of Γ under the information model $\hat{I}$, then σ' is a Bayesian equilibrium of the associated game G'. Since $[\Gamma] \cap [\sigma] \subseteq \hat{K}^0[\sigma]$, the profile σ is a naive equilibrium if and only if the event $[\Gamma] \cap [\sigma] \cap [\hat{R}]$ is non-empty. Suppose that this is the case: there is a ω such that $\tilde{\Gamma}(\omega) = \Gamma$, $\tilde{\sigma}(\omega) = \sigma$, and that for every s_2, $\tilde{\sigma}_2(s_2, \omega) > 0$ only if $s_2 \in \arg\max_{s_2'} E\left[u_2(s_1, s_2', \theta_1)|\hat{P}_2(\omega)\right] = E\left[u_2(s_1, s_2', \theta_1)|P_2(\omega)\right]$, and for every i, and s_1, $\sigma^i(s_1, \omega) > 0$ only if $s_1 \in \arg\max_{s_1'} E\left[u^i(s_1', s_2, \theta_1)|\hat{P}^i(\omega)\right] = E\left[u^i(s_1', s_2, \theta_1)|P^i(\omega)\right]$. Since $\tilde{\sigma}(\omega) = \sigma$, it follows that, by plugging in the expressions for $P_2(\omega)$ and $P^i(\omega)$ we obtain Conditions (16) and (15). This implies that σ' must satisfy Conditions (8) and (9), i.e. σ' is a subjective equilibrium of the game G'. ■

The proofs of Proposition 5 and 6 are analogous to the proof of Proposition 1. Similarly the proof of Proposition 8 is easily derived from the proof of Proposition 4. These proofs, available upon request to the author, are omitted.

Proof of Proposition 7. The proof of the first part is analogous to the proof of Proposition 3, and is thus omitted. For the second part, consider the following game $G = (\Theta,$

ϕ, S, u, μ), such that $\Theta = \{\theta_L, \theta_H\}$, $\phi(\theta_L, \theta_H) = 1$, μ is degenerate, S and u are represented below (assume that $x \geq 2$).

θ_L	C	D
A	1,0	0,1
B	0,1	1,0

θ_H	C	D
A	x,0	0,x
B	0,1	1,0

We want to show that for any list l other than l^0 or l^*, and any information model I, there is a game G that does not possess any sophisticated equilibrium for $E_l(G)$. Also, note that while the pair distribution ϕ does not have full support, it is easy to see how to extend this result to the case of games with generic full support distribution.

By definition, the distribution ϕ^0 is such that $\phi^0(\theta_H, \theta_H) = 1$. For any $\sigma_1 \in \Delta(\{A, B\})$, and any $\sigma_2 \in \Delta(\{C, D\})$, define the event

$$E(\sigma_1, \sigma_2) = [\sigma_1] \cap [R_1] \cap K_1[\mathbf{G}^0] \cap K_1[K_2[\mathbf{G}^0]] \cap K_1[[R_2] \cap [\sigma_2] \cap K_2[\sigma_1]].$$

The first step is to show that for any $\sigma_1(A) \neq x/[x+1]$ and any $\sigma_2(C) \neq x/[x+1]$, $E(\sigma_1, \sigma_2) = \emptyset$. Pick in fact any $\sigma_1(A) > x/[x+1]$, then $K_2[G^0] \cap K_2[\sigma_1] \cap [R_2] \subseteq [\sigma_2(C) = 0]$. But $K_1[K_2[G^0] \cap K_2[\sigma_1] \cap [R_2]] \subseteq K_2[G^0] \cap K_2[\sigma_1] \cap [R_2]$, and $K_1[G^0] \cap K_1[\sigma_2(C) = 0] \cap [R_1] \subseteq [\sigma_1(A) = 0]$. Thus $E(\sigma_1, \sigma_2) = \emptyset$, for any $\sigma_1(A) > x/[x+1]$. Conversely, for $\sigma_1(A) < x/[x+1]$, then $K_2[G^0] \cap K_2[\sigma_1] \cap [R_2] \subseteq [\sigma_2(C) = 1]$, but $K_1[G^0] \cap K_1[\sigma_2(C) = 1] \cap [R_1] \subseteq [\sigma_1(A) = 1]$, and $E(\sigma_1, \sigma_2) = \emptyset$, for any $\sigma_1(A) < x/[x+1]$. The proof that $E(\sigma_1, \sigma_2) = \emptyset$ for any $\sigma_2(C) \neq x/[x+1]$ is identical. It follows that $K_2[E(\sigma_1, \sigma_2)] \subseteq K_2[\sigma_1(A) = x/[x+1]]$. For any $\omega \in [G] \cap K_2[G] \cap K_2[E(\sigma_1, \sigma_2)]$, it follows that $\omega \in [\sigma_2(C) = 0]$. This proves the claim for any list l such that $\{[G], K_1[G^0], K_2[G], K_1[K_2[G^0]]\} \subseteq l$. To show the claim for any list l such that $\{[G], K_1[G^0], K_2[G], \sim K_1[K_2[G]]\} \subseteq l$, note that for any prior p, any state ω, and following posterior $p\{K_2[G^0]|P_1(\omega)\}$, we can set x arbitrarily large so as to get again $\omega \in [\sigma_2(C) = 0]$.

Letting $E^1 = K_2[E(\sigma_1(A) = x/[x+1], \sigma_2(C) = x/[x+1])] \cap K_2[G] \cap [R_2]$, the above passages also show that for any $\omega \in [\phi] \cap K_1[E^1] \cap K_1[G^0] \cap [R_1] \cap K_1[\sigma_2(C) = 0] \cap K_1[G^0] \cap [R_1]$ implies that $\omega \in [\sigma_1(A) = 0]$. This shows the claim for any list l such that $\{[G], K_1[G^0], K_2[G], K_1[K_2[G]], K_2[K_1[G^0]], K_2[K_1[K_2[G]]]\} \subseteq l$. Letting $E^2 = K_1[E^1] \cap K_1[G^0] \cap [R_1]$, one obtains that $\omega \in [\phi] \cap K_2[E^2] \cap K_2[G] \cap [R_2]$ implies that $\omega \in [\sigma_1(C) = 1]$, and shows that the claim holds for any list l such that $\{[G], K_1([\mathbf{G}^0] \cap K_2([\mathbf{G}] \cap K_1[\mathbf{G}^0]))$, $K_2([\mathbf{G}] \cap K_1([\mathbf{G}^0] \cap K_2[\mathbf{G}]))$, $K_1[K_2[K_1[K_2[\mathbf{G}^0]]]]\} \subseteq l$.

By repeating the construction ad infinitum, the claim is proved. ■

References

[1] Aumann R. [1999]: "Interactive Epistemology I: Knowledge" International Journal of Game Theory 28: 263-300

[2] Aumann R. and A. Brandenburger [1995] "Epistemic Conditions for Nash Equilibrium" Econometrica 63: 1161-80

[3] Benabou R. and J. Tirole [1999]: "Self-Confidence: Intrapersonal Strategies" mimeo, IDEI, Toulouse and NYU

[4] Benabou R. and J. Tirole [2000]: "Self-Confidence and Social Interactions" NBER working paper 7585

[5] Brandenburger A. and E. Dekel [1987]: "Rationalizability and Correlated Equilibrium" Econometrica 55: 1391-402

[6] Brandenburger A. and E. Dekel [1993]: "Hierarchies of Beliefs and Common Knowledge" Journal of Economic Theory 59: 189-98

[7] Burger J. and L. Burns [1988]: "The Illusion of Unique Invulnerability and the Use of Effective Contraception" Personality and Social Psychology Bulletin 14: 264-70

[8] Camerer C. and D. Lovallo [1999]: "Overconfidence and Excess Entry: and Experimental Approach" American Economic Review 89: 306-318

[9] Dekel E. and F. Gul [1997]: "Rationality and Knowledge in Game Theory", in Advances in economics and econometrics: Theory and applications: Seventh World Congress, Volume 1, Kreps D. M. and Wallis K. F. eds, Econometric Society Monographs 26, Cambridge University Press, Cambridge

[10] Dudley R. M. [1999]: Real Analysis and Probability Chapman and Hall, New York

[11] Einhorn H.J. and R.M. Hogarth [1978]: "Confidence in Judgement: Persistence in the Illusion of Validity", Psychological Review 85: 395-416

[12] Epstein L. and T. Wang [1996]: "Beliefs about Beliefs without Probabilities" Econometrica 64: 1343-73

[13] Flam S.D. and A.E. Risa [1998]: "Status, Self-Confidence, and Search", University of Bergen mimeo

[14] Fudenberg D. and Tirole J. [1991]: Game Theory, Cambridge MA, MIT press

[15] Harsanyi J. [1967]: "Games with Incomplete Information Played by 'Bayesian' Players, part I, II, III", Management Science 14: 159-182, 320-334, 486-502

[16] Heifetz A. and Y. Spiegel [2000] "On the Evolutionary Emergence of Optimism," Math Center WP 1304, Northwestern University

[17] Hirshleifer J. [1971]: "The Private and Social Value of Information and the Reward to Inventive Activity" American Economic Review 61: 561-573

[18] Kamien M. We, Y. Tauman and S. Zamir [1990]: "On the Value of Information in a Strategic Conflict" Games and Economic Behavior 2: 129-153

[19] Lewinsohn P, W. Mischel, W. Chaplin, and R. Barton [1980]: "Social Competence and Depression: the Role of Illusory Self-Perceptions" Journal of Abnormal Psychology 898: 203-212

[20] Mertens F. and S. Zamir [1985]: "Formulation of Bayesian Analysis for Games with Incomplete Information" International Journal of Game Theory 14: 1-29

[21] Neyman A. [1991]: "The Positive Value of Information" Games and Economic Behavior 3: 330-355

[22] Segerstrom S, W. McCarthy, N. Caskey, T. Gross and M. Jarvik [1993]: "Optimistic Bias among Cigarette Smokers" Journal of Applied Social Psychology 23: 1606-18

[23] Svenson O. [1981]: "Are we all less Risky Drivers and more Skillful than our Fellow Drivers?" Acta Psychologica 47: 143-148