Belief-invariant equilibria in games of incomplete information*

Vincenzo Auletta[1], Diodato Ferraioli[1], Ashutosh Rai[2], Giannicola Scarpa[3], and Andreas Winter[4]

[1] DIEM, Università degli Studi di Salerno, Italy
{auletta, dferraioli}@unisa.it
[2] Faculty of Computing, University of Latvia
ashutosh.rai@lu.lv
[3] Facultad de Ciencias Matemáticas, Universidad Complutense de Madrid, Spain
gscarpa@ucm.es
[4] ICREA and Departament de Física: Grup d'Informació Quàntica, Universitat Autònoma de Barcelona, Spain
andreas.winter@uab.cat

Abstract. Drawing on ideas from game theory and quantum physics, we investigate nonlocal correlations from the point of view of equilibria in games of incomplete information. These equilibria can be classified in decreasing power as general *communication equilibria*, *belief-invariant equilibria* and *correlated equilibria*, all of which contain *Nash equilibria*. The notion of belief-invariant equilibrium has appeared in game theory before (in the 1990s). However, the class of *non-signalling correlations* associated to belief-invariance arose naturally already in the 1980s in the foundations of quantum mechanics.

In the present work, we explain and unify these two origins of the idea and study the above classes of equilibria. We present a general framework of belief-invariant communication equilibria, which contains correlated equilibria as special cases. We then use our framework to show new results related to the *social welfare* of games. Namely, we exhibit a game where belief-invariance is socially better than any correlated equilibrium, and a game where all non-belief-invariant communication equilibria have a suboptimal social welfare. We also show that optimal social welfare can in certain cases be achieved by quantum mechanical correlations, which do not need an informed mediator to be implemented, and go beyond the classical "sunspot" or shared randomness approach.

1 Introduction

The notion of equilibrium of a strategic game and the mathematical formulation of rational behaviour are among the most fruitful ideas of the last century. The topic was initiated by the classic treatment of von Neumann and Morgenstern [13], and one of the fundamental milestones has been the definition of

* A full version of this paper is available as [1].

Nash equilibrium [12] and Nash's proof that finite games always have such an equilibrium. These pioneering results were followed by a multitude of further investigations into other concepts of equilibrium and their properties. For example, great attention has been devoted to the question of how the players, knowing the game, can find an equilibrium [2,11]. The realization that Nash equilibria sometimes can be "bad" both individually and collectively for the players, has motivated a major direction in game theory, i.e., to explore how players can be induced to a more beneficial equilibrium [14]. One important idea is that giving the players some *advice*, in the form of a random variable generated by a correlation device, can change the landscape of equilibria. This generalizes the concept of Nash equilibrium to *correlated equilibria* [2].

The present paper deals with advice in the setting of *games of incomplete information*. As it turns out, this is a subject of considerable complexity, since correlation devices can be far more general than in the complete information setting. In games of incomplete information, or *Bayesian games*, each player has a *type* which is not perfectly known to, but only estimated by, the other players. Depending on what the game models, a type can represent different properties: e.g., a characteristic of the player (strong, weak, rich, poor, etc.) or a secret objective of the player (interest in one particular outcome). For these games, a relevant solution concept is the *communication equilibrium* [6]. Here, the players privately communicate their type to a mediator, who implements a correlation and gives each player advice for a convenient action. It is reasonable to assume that players are comfortable with revealing their private information to a trusted mediator if this gives them an advantage. However, the advices of the mediator can reveal a player's private information even to other players, and there are situations where it is crucial for players that this never occurs (e.g., trade secrets). Thus, it would be interesting to study correlation devices that do not allow that the information about the private type of one player is leaked by the other players. These devices have been already introduced in game theory: e.g., in [7] equilibria based on these devices are called "belief-invariant". However, the property of these devices is usually adopted to make the analysis of the equilibria more convenient, and is not highlighted as interesting in its own right. From a completely different angle, belief-invariance has been a topic of research in physics (motivated by questions in the foundations of quantum mechanics [15]) and theoretical computer science (motivated by multi-prover interactive proof systems [10] and parallel repetition of games [3]), under the name of *non-signalling correlations*. Here, belief-invariance is relevant because it describes the largest class of correlations that obey relativistic causality.

Here, we bring together the strands of thought coming from these two backgrounds. On the one hand, this results in a more general and much richer picture of non-locality as a resource and, on the other hand, allows us to import findings from literature in physics and theoretical computer science about non-locality and non-signalling to game theory.

In this paper we study the class of *belief-invariant communication equilibria* and compare it with the classes of communication equilibria and correlated equilibria. In particular, we will evaluate these classes with respect to privacy, computational complexity, and *social welfare* of games. Specifically, we highlight that these three equilibrium concepts have different requirements about who can leak information about players' private type. Moreover, we report the known hardness results for these equilibria, by highlighting some interesting open problems that may be of independent interest to the TCS community. Finally, we exhibit a game where belief-invariance is socially better than any correlated equilibrium, and a game where all non-belief-invariant communication equilibria have a suboptimal social welfare. That is, belief-invariant equilibria, even if they are more constrained than communication equilibria, still they can perform better than the latter for what concerns social welfare.

Next we formally introduce the concepts of interest of this paper.

2 Definitions

A general *correlation* is a joint conditional probability distribution $Q(\mathbf{s} \mid \mathbf{r})$, where $\mathbf{r} = (r_1, \dots, r_n)$ is a tuple of *inputs* r_i for each player i, with r_i drawn from an alphabet R_i, and $\mathbf{s} = (s_1, \dots, s_n)$ is a tuple of *outputs* s_i for each player i, with s_i drawn from an alphabet S_i. A joint conditional probability distribution Q is *belief-invariant* (also called *non-signalling*) if the distribution of the output variable s_i given r_i does not give any information about r_j, with $j \neq i$. This class is easily seen to be strictly contained in the general class of correlations. A joint conditional probability distribution Q is called *local* if it can be simulated locally by each party i, by observing (their part of) a random variable $\boldsymbol{\gamma} = (\gamma_1, \dots, \gamma_n)$ (with distribution $V(\boldsymbol{\gamma})$) independent of $\mathbf{r}$, and doing local operations depending only on r_i and γ_i.

A *game with incomplete information* G is defined by the following objects: a finite set of *players* N of size n; a finite set of *type profiles* $T := \times_i T_i$; a finite set of *action profiles* $A := \times_i A_i$; A prior probability distribution $P(t)$ on the types $t \in T$; For each player $i \in N$, a *payoff function* $v_i \colon T \times A \to \mathbb{R}$. A *strategy* g_i for the player i is a map from the information known to i to an action $a_i \in A_i$.

The game goes as follows. The types $\mathbf{t} = (t_1, \dots, t_n)$ are sampled according to P. Each player i learns his type t_i, uses his strategy g_i to select an action $a_i \in A_i$, and is awarded according to his payoff function v_i (which can depend on the other players' actions and types). The expected utility of player i is $\langle v_i \rangle = \sum_{\mathbf{t},\mathbf{a}} P(\mathbf{t}) v_i(\mathbf{t}, \mathbf{a}) \prod_{i=1}^{n} g_i(a_i \mid t_i)$, where $\mathbf{g} = (g_1, \dots, g_n)$ and $\mathbf{a} = (a_1, \dots, a_n)$.

A *solution* is a family of strategies $\mathbf{g} = (g_1, \dots, g_n)$, one for each player. A solution is then said to be an *equilibrium* (more precisely, a *Nash equilibrium*) if no player has an incentive to change the adopted strategy. I.e., $\langle v_i \rangle = \mathbb{E}_{\mathbf{t},\mathbf{g}_{-i}} v_i\big(\mathbf{t}, \mathbf{g}(\mathbf{t}) g_i(t_i)\big) \geq \mathbb{E}_{\mathbf{t}} \mathbb{E}_{\mathbf{g}_{-i}} v_i\big(\mathbf{t}, \mathbf{g}_{-i}(\mathbf{t}_{-i}) \chi_i(t_i)\big)$, for all i and $\chi_i \in A_i^{T_i}$.

A *solution with communication* for G studies the behaviour of players who have access to a *correlation device* that depends on inputs communicated by the players during the game. The most common operational interpretation of this setting is that a *trusted mediator*, who has private communication channels with all the players, collects from each player i the input r_i, samples $\mathbf{s}$ according to

$Q(\mathbf{s} \mid \mathbf{r})$ and sends to each i the output s_i. Formally, we add to the strategies of the players the use of a *correlation* $Q(\mathbf{s} \mid \mathbf{r})$. In this setting, a *pure strategy* for each player i is a pair of functions, $f_i \colon T_i \to R_i$ and $g_i \colon T_i \times S_i \to A_i$; and a *mixed strategy* is a pair of jointly distributed random functions $(f_i, g_i) \in R_i^{T_i} \times A_i^{T_i \times S_i}$.

The game now goes as follows. The types $\mathbf{t} = (t_1, \ldots, t_n)$ are sampled according to P. Each player i learns his type t_i, and sends the input $r_i = f_i(t_i)$ to the correlation device. He then gets the correlation output s_i and plays the action $a_i = g_i(t_i, s_i)$. The expected payoff of i is: $\langle v_i \rangle = \sum_{\mathbf{t},\mathbf{s}} P(\mathbf{t}) Q\big(\mathbf{s} \mid f_1(t_1), \ldots, f_n(t_n)\big) v_i\big(\mathbf{t}, g_1(t_1, s_1), \ldots, g_n(t_n, s_n)\big)$.

The most general class we consider here is the class of *communication equilibria*. Formally, a solution $(\mathbf{f}, \mathbf{g}, Q)$ is a *communication equilibrium* of G if for each i we have $\sum_{\mathbf{t},\mathbf{s}} P(\mathbf{t}) Q(\mathbf{s} \mid f_i(t_i)\mathbf{f}_{-i}(\mathbf{t}_{-i})) v_i(t, g_i(t_i, s_i)\mathbf{g}_{-i}(\mathbf{t}_{-i}, \mathbf{s}_{-i})) \geq \sum_{\mathbf{t},\mathbf{s}} P(\mathbf{t}) Q(\mathbf{s} \mid \varphi_i(t_i)\mathbf{f}_{-i}(\mathbf{t}_{-i})) v_i(t, \chi_i(t_i, s_i)\mathbf{g}_{-i}(\mathbf{t}_{-i}, \mathbf{s}_{-i}))$, for all random functions $\varphi_i \colon T_i \to R_i$ and $\chi_i \colon S_i \to A_i$

We obtain some subclasses of communication equilibria by restricting the kind of correlation used in the equilibrium. A solution $(\mathbf{f}, \mathbf{g}, Q)$ is called *belief-invariant* if Q is a belief-invariant correlation. If $(\mathbf{f}, \mathbf{g}, Q)$ is a communication equilibrium, we call it a *belief-invariant (communication) equilibrium.* A solution $(\mathbf{f}, \mathbf{g}, Q)$ is, instead, called *correlated* if the output distribution of Q is independent of the input: $Q(\mathbf{s} \mid \mathbf{r}) = Q(\mathbf{s})$ for all $\mathbf{r}$ and $\mathbf{s}$. If it is a communication equilibrium, we speak of a *correlated (communication) equilibrium.*

3 Properties of equilibria

Privacy. Clearly, in order to implement a correlated equilibrium, no player except i needs to learn the type t_i. The belief-invariant class allows for a larger set of correlations at the price that a trusted mediator might learn something about the types. The use of a belief-invariant correlation guarantees however that the mediator will be the only one learning the types and no player except i can learn t_i. It is not always possible to respect this requirement in the more general class of communication equilibria.

*Computational complexity. N*ash equilibria of complete information games are hard to find: in fact, it is known that the problem is PPAD-hard even for two-player games [5,4]. Since games of incomplete information contain complete-information games as a special case, they are at least as hard. *C*orrelated equilibria of complete information games can be found in time that is polynomial in the size of the game specification through linear programming [9]. It is left open to understand if the above result can be extended to games of incomplete information. *B*elief-invariant equilibria of full coordination games of incomplete information can be instead found in time that is polynomial in the size of game description via linear programming. This is because the set of non-signalling correlations is defined by polynomially many non-negative variables subject to polynomially many linear inequalities. (See, for example, the LP in [3, page 8].) It is left open to understand if this extends to other games.

Social welfare. Finally, we consider the expected *social welfare* $\mathtt{SW}(\mathbf{g})$ of a solution $\mathbf{g}$, i.e., the sum of the expected payoffs of all players, $\mathtt{SW}(\mathbf{g}) = \sum_i \langle v_i \rangle$. We show that no-signalling correlation can have a positive impact on the social welfare of a game.

Indeed, we present a n-player game with conflict of interests in which a belief-invariant equilibrium exists that is better than any correlated equilibrium. Interestingly, our game is a variant of the GHZ game, a game motivated from quantum mechanics [8].

Since the class of belief-invariant equilibria strictly contains the class of correlated equilibria, it may be expected that the former contains equilibria that are better than the ones in the latter class. It is instead surprising that a correlated equilibrium can perform better than any other equilibrium in the class, even unrestricted ones. However, we show that there is a game (a special generalization of the Prisoners' Dilemma), for which this is the case. In other word, we prove that locality is not only a desirable requirement, but it is sometimes necessary in order to achieve high social welfare.

References

1. Auletta, V., Ferraioli, D., Rai, A., Scarpa, G., Winter, A.: Belief-invariant equilibria in games with incomplete information. CoRR abs/1605.07896 (2016)
2. Aumann, R.J.: Subjectivity and correlation in randomized strategies. Journal of mathematical Economics 1(1), 67–96 (1974)
3. Buhrman, H., Fehr, S., Schaffner, C.: On the parallel repetition of multi-player games: The no-signaling case. In: TQC '14 (2014)
4. Chen, X., Deng, X., Teng, S.: Settling the complexity of computing two-player Nash equilibria. J. ACM 56(3), 14:1–14:57 (2009)
5. Daskalakis, C., Goldberg, P.W., Papadimitriou, C.H.: The complexity of computing a Nash equilibrium. SIAM Journal on Computing 39(1), 195–259 (2009)
6. Forges, F.: A first study of correlated equilibria in repeated games with incomplete information. Core discussion paper 8218 (1982)
7. Forges, F.: Five legitimate definitions of correlated equilibrium in games with incomplete information. Theory and Decision 35(3), 277–310 (1993)
8. Greenberger, D.M., Horne, M.A., Shimony, A., Zeilinger, A.: Bell's theorem without inequalities. American Journal of Physics 58(12), 1131–1143 (1990)
9. Hart, S., Schmeidler, D.: Existence of correlated equilibria. Mathematics of Operations Research 14(1), 18–25 (1989)
10. Kalai, Y.T., Raz, R., Rothblum, R.D.: How to delegate computations: The power of no-signaling proofs. In: STOC '14. pp. 485–494 (2014)
11. Lemke, C.E., Howson, J.T.: Equilibrium points of bimatrix games. Journal of the Society for Industrial & Applied Mathematics 12(2), 413–423 (1964)
12. Nash, J.F.: Equilibrium points in n-person games. Proceedings of the National Academy of Sciences 36(1), 48–49 (1950)
13. von Neumann, J., Morgenstern, O.: Theory of Games and Economic Behavior. Princeton University Press (1944)
14. Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V.V.: Algorithmic Game Theory. Cambridge University Press (2007)
15. Tsirelson, B.S.: Quantum generalizations of Bell's inequality. Lett. Math. Phys. 4(2), 93–100 (1980)