Properties of Dual Reduction*

Yannick Viossat†

This version: July 2003

Abstract

Dual reduction, introduced by Myerson, allows to reduce games in a way that selects among correlated equilibrium distributions. Myerson's results are first recalled, then new properties of dual reduction are established. We show that generic two-player games have a unique sequence of iterative full dual reductions. We compare dual reduction to other correlated equilibrium refinements. Finally, we review and connect the linear programming proofs of existence of correlated equilibria.

1 Introduction

The first direct proofs of existence of correlated equilibrium distributions, based on the duality theorems of linear programming, were developed independently by Hart and Schmeidler [3] and Nau and McCardle [10]. These proofs are essentially identical, as shown in appendix A. They laid the mathematical foundations of dual reduction [7]. Dual reduction is a method to reduce finite games into games with fewer strategies in a way that selects among correlated equilibrium distributions. That is, any correlated equilibrium distribution of the reduced game induces a correlated equilibrium distribution in the original game. Myerson [7] shows that dual reduction includes elimination of weakly dominated strategies as a subprocess, and that, by iterative dual reduction, any game is eventually reduced to a game which has a strict correlated equilibrium distribution with full support. We see dual reduction as a powerful tool to study correlated equilibrium distributions. The aim of this paper is to investigate further the properties of dual reduction.

After introducing the basic notations and definitions in section 2, we recall the key-points of the direct proofs of existence of correlated equilibrium distributions, in section 3, and review the existing results on dual reduction in section 4. New results are established in sections 5 and 6. They are summed up at the beginning of section 5. In section 7, we briefly compare dual reduction to another correlated equilibrium refinement introduced by Myerson [8]: elimination of unacceptable pure strategies. Long proofs are gathered in section 8. Finally, in the appendix, we review and connect the proofs of existence of correlated equilibria given in [3], [10] and [7].

2 Notations and definitions

2.1 Basic notations

The analysis in this paper is restricted to finite games in strategic forms. The notations are taken from [7]. Let $\Gamma = \{N, (C_i)_{i\in N}, (U_i)_{i\in N}\}$ denote a finite game in strategic form: N is the nonempty finite set of players, C_i the

*I gratefully acknowledge the support and helpful comments of Bernhard von Stengel, Thomas Boulogne, Sylvain Sorin and Galit Ashkenazi.
†Laboratoire d'économétrie de l'Ecole polytechnique, 1 rue Descartes, 75005 Paris. E-mail: viossat@poly.polytechnique.fr

nonempty finite set of pure strategies of player i and $U_i : \times_{i\in N} C_i \to \mathbb{R}$ the utility function of player i. The set of (pure) strategy profiles is $C = \times_{i\in N} C_i$; the set of strategy profiles for the players other than i is $C_{-i} = \times_{j\in N-i} C_j$. Pure strategies of player i (strategy profiles; strategy profiles of the players other than i) are denoted c_i or d_i (c; c_{-i}). We may write (c_{-i}, d_i) to denote the strategy profile that differs from c only in that its $i-$component is d_i. For any finite set S, $\Delta(S)$ denotes the set of probability distributions over S. Thus $\Delta(C_i)$ is the set of mixed strategies of player i, which are denoted by σ_i or τ_i.

2.2 Correlated equilibrium distributions and deviation vectors

A *correlated strategy* of the players in N is is an element of $\Delta(C)$. Thus $\mu = (\mu(c))_{c\in C}$ is a correlated strategy if:

$$\mu(c) \geq 0 \qquad \forall c \in C$$
$$\sum_{c\in C} \mu(c) = 1$$

A correlated strategy is a *correlated equilibrium distribution* [Aumann, 1974] if it satisfies the following *incentive constraints*:

$$\sum_{c_{-i}\in C_{-i}} \mu(c)[U_i(c) - U_i(c_{-i}, d_i)] \geq 0 \qquad \forall i \in N, \forall c_i \in C_i, \forall d_i \in C_i \tag{1}$$

The following interpretation and the vocabulary introduced below will be useful for the next sections. Let $\mu \in \Delta(C)$ and consider the following extended game Γ_μ, based on Γ: before Γ is played, a strategy profile $c \in C$ is drawn at random, with probability $\mu(c)$, and c_i is privately announced to player i; then Γ is played. The players can thus condition their strategy in Γ on their private signal. A strategy of player i in this extended game is a *deviation plan*, i.e. a mapping $\alpha_i : C_i \to \Delta(C_i)$. Denoting by $\alpha_i(d_i|c_i)$ the probability that player i will play d_i when announced c_i we have:

$$\alpha_i(d_i|c_i) \geq 0 \qquad \forall c_i \in C_i, \forall d_i \in C_i, \forall i \in N \tag{2}$$
$$\sum_{d_i\in C_i} \alpha_i(d_i|c_i) = 1 \qquad \forall c_i \in C_i, \forall i \in N \tag{3}$$

A strategy profile is a *deviation vector*, i.e. a vector $\alpha = (\alpha_i)_{i\in N}$ of deviation plans. Such a deviation vector is *trivial* if, for all i in N, α_i is the identity mapping. The incentive constraints (1) mean that μ is a correlated equilibrium distribution of Γ if and only if the trivial deviation vector is a Nash equilibrium of Γ_μ.

3 Existence of correlated equilibrium distributions

This section is a variation on [3], [10] and [7]. Consider the following two-player, zero-sum auxiliary game G: the maximizer chooses a correlated strategy μ in $\Delta(C)$; the minimizer chooses a deviation vector α. The payoff is:

$$g(\mu, \alpha) = \sum_{c\in C} \mu(c) \sum_{i\in N} \sum_{d_i\in C_i} \alpha_i(d_i|c_i)[U_i(c) - U_i(c_{-i}, d_i)]^1 \tag{4}$$

It is clear from section 2.2 that μ guarantees 0 if and only if μ is a correlated equilibrium distribution of Γ. Thus Γ has a correlated equilibrium distribution if and only if the value of G is nonnegative. The remaining of this section is devoted to a proof of the following theorem:

[1] It is clear that G has a value. Indeed, G is the extension in behavioral strategies of the following two-player, zero-sum game: first the maximizer privately chooses c in C; then a player i in N is selected (with probability $1/n$, where n is the number of players) and c_i announced to the minimizer; the minimizer then chooses a deviation d_i from c_i. The payoff is: $n \times [U_i(c) - U_i(c_{-i}, d_i)]$.

Theorem 3.1 *The value of G is zero. Therefore correlated equilibrium distributions exists.*

A deviation plan $\alpha_i : C_i \to \Delta(C_i)$ induces a Markov chain on C_i. This Markov chain maps the distribution $\sigma_i \in \Delta(C_i)$ to the distribution $\alpha_i * \sigma_i$ given by:

$$\alpha_i * \sigma_i(d_i) = \sum_{c_i \in I} \alpha_i(d_i|c_i)\sigma_i(c_i) \forall d_i \in C_i$$

Similarly, if a mediator tries to implement μ[2] but player i deviates (unilaterally) according to α_i, this generates a new distribution on strategy profiles $\alpha_i * \mu$:

$$\alpha_i * \mu(c_{-i}, d_i) = \sum_{c_i \in C_i} \alpha_i(d_i|c_i)\mu_i(c) \qquad \forall d_i \in C_i, \forall c_{-i} \in C_{-i}$$

Definition 3.2 *Let $\alpha = (\alpha_i)_{i \in N}$ be a deviation vector. A mixed strategy $\sigma_i \in \Delta(C_i)$ is α_i-*invariant *if $\alpha_i * \sigma_i = \sigma_i$. A correlated strategy $\mu \in \Delta(C)$ is α_i-*invariant *(α-*invariant*) if (if for all $i \in N$) $\alpha_i * \mu = \mu$.*

Note that, by the basic theory of Markov chains, there exists at least one α_i-invariant strategy.

Let $U_i(\mu) = \sum_{c \in C} \mu(c) U_i(c)$ denote the average payoff of player i if μ is implemented. Myerson shows that:

$$g(\mu, \alpha) = \sum_{i \in N} [U_i(\mu) - U_i(\alpha_i * \mu)] \tag{5}$$

We can now prove theorem 3.1: first note that the minimizer can guarantee 0 by choosing the trivial deviation vector. Thus we only need to show that the maximizer can defend 0. Let α denote a deviation vector; for each i, let $\sigma_i \in \Delta(C_i)$ be α_i-invariant. The correlated strategy $\sigma = \prod_{i \in N} \sigma_i$ is α-invariant; hence, by (5), $g(\sigma, \alpha) = 0$. Therefore, the maximizer can defend 0.

4 Dual reduction

All results of this section are proved in [7].

4.1 Definition

The Markov chain on C_i induced by α_i partitions C_i into transient states and disjoint minimal absorbing sets[3]. For any minimal absorbing set B_i, there exists a unique α_i-invariant strategy with support in B_i[4]. Let C_i/α_i denote the set of (randomized) α_i-invariant strategies with support in some minimal α_i-absorbing set. It may be shown that the set of α_i-invariant strategies is the set of random mixture of the strategies in C_i/α_i; that is, the simplex $\Delta(C_i/\alpha_i)$.

Let $\alpha = (\alpha_i)_{i \in N}$ be a deviation vector. The α-reduced game $\Gamma/\alpha = \{N, (C_i/\alpha_i)_{i \in N}, (U_i)_{i \in N}\}$ is the game obtained from Γ by restricting the players to α-invariant strategies. That is, the set of players and the payoff functions are the same than in Γ but, for all i in N, the pure strategy set of player i is now C_i/α_i.[5]

[2] That is, the mediator draws a strategy profile c in C with probability $\mu(c)$ and then privately recommends c_i to player i.

[3] A subset B_i of C_i is α_i-absorbing if $\alpha(d_i|c_i) = 0$ for all c_i in B_i and all d_i in $C_i - B_i$. An α_i-absorbing set is minimal if it contains no proper α_i-absorbing subset.

[4] Actually its support is exactly B_i.

[5] Strictly speaking the payoff function of the reduced game is the function *induced* by the original game's payoff function on the reduced strategy space.

Before turning to dual reduction and their properties, let us make our vocabulary precise: let $c_i, d_i \in C_i$ ($c \in C$). The pure strategy c_i (strategy profile c) is *eliminated* in the α-reduced game Γ/α if $\sigma_i(c_i) = 0$ for all σ_i in C_i/α_i (if $\sigma(c) = 0$ for all σ in C/α). Thus c_i (resp. c) is eliminated if and only if (if and only if for some i in N) c_i is transient under α_i. The strategies c_i and d_i are *grouped together* if there exists σ_i in C_i/α_i such that $\sigma_i(c_i)$ and $\sigma_i(d_i)$ are positive. Thus, c_i and d_i are grouped together if and only if they are recurrent under α_i and belong to the same minimal α_i-absorbing set.

Definition 4.1 *A* dual vector *is an optimal strategy of the minimizer in the auxiliary game of section 3. Thus a deviation vector α is a dual vector if:*

$$-g(c,\alpha) = \sum_{i \in N}[U_i(\alpha_i * c) - U_i(c)] = \sum_{i \in N}\sum_{d_i \in C_i} \alpha_i(d_i|c_i)[U_i(c_{-i}, d_i) - U_i(c)] \geq 0 \qquad \forall c \in C \tag{6}$$

(The above equalities merely repeat the definition of $g(c,\alpha)$.)

Definition 4.2 *A* dual reduction *of Γ is an α-reduced game Γ/α where α is a dual vector. An* iterative dual reduction *of Γ is a reduced game $\Gamma/\alpha^1/\alpha^2/.../\alpha^m$, where m is a positive integer and, for all k in $\{1,2,...,m\}$, α^k is a dual vector of $\Gamma/\alpha^1/\alpha^2/.../\alpha^{k-1}$.*

Many examples can be found in [7, section 6]. Henceforth, unless stated otherwise, α is a dual vector.

4.2 Main properties

First, dual reduction generalizes elimination of weakly dominated strategies in the following sense:

Proposition 4.3 *Let $c_i \in C_i$; assume that there exists $\sigma_i \in \Delta(C_i), \sigma_i \neq c_i$, such that $U_i(c_{-i}, \sigma_i) \geq U_i(c)$ for all c_{-i} in C_{-i}. Then there exists a dual vector α such that $C_i/\alpha_i = C_i - \{c_i\}$ and $C_j/\alpha_j = C_j$ for $j \neq i$.*

Proof. Take for α: $\alpha_i(d_i|c_i) = \sigma_i(d_i)$ for all $d_i \in C_i$, and $\alpha_j(c_j|c_j) = 1$ if $j \neq i$ or $c_j \neq c_i$ ■

The main property of dual reduction is that it selects among correlated equilibrium distributions: let Γ/α denote a dual reduction of Γ; let $C/\alpha = \times_{i \in N} C_i/\alpha_i$ denote the set of strategy profiles of Γ/α. Let $\lambda \in \Delta(C/\alpha)$; the Γ-*equivalent correlated strategy* $\bar{\lambda}$ is the distribution on C induced by λ:

$$\bar{\lambda}(c) = \sum_{\sigma \in C/\alpha} \lambda(\sigma) \left(\prod_{i \in N} \sigma_i(c_i) \right) \tag{7}$$

Theorem 4.4 *If λ is a correlated equilibrium distribution of Γ/α, then $\bar{\lambda}$ is a correlated equilibrium distribution of Γ.*

By induction, theorem 4.4 extends to iterative dual reductions. That is, any correlated equilibrium distribution of an iterative dual reduction of Γ induces a correlated equilibrium distribution of Γ. A side product of the proof of theorem 4.4 is that, against any strategy of the other players in the reduced game, player i is indifferent between his strategies within a minimal absorbing set:

Proposition 4.5 *Let B_i denote a minimal α_i-absorbing set. For $j \neq i$, let $\sigma_j \in C_j/\alpha_j$ and let $\sigma_{-i} = \times_{j \in N-i} \sigma_j$. For any c_i, d_i in B_i, $U_i(\sigma_{-i}, c_i) = U_i(\sigma_{-i}, d_i)$.*

4.3 Jeopardization and Elementary Games

Let us say that a dual vector is trivial if it is the trivial deviation vector. A game may be reduced if and only if there exists a nontrivial dual vector[6]. So we are led to the question: when do nontrivial dual vectors exist ? A first step to answer this question is to introduce the notions of jeopardization and elementary games:

Definition 4.6 *Let $c_i, d_i \in C_i$. The strategy d_i* jeopardizes *c_i if for all correlated equilibrium distributions μ:*

$$\sum_{c_{-i} \in C_{-i}} \mu(c)[U_i(c) - U_i(c_{-i}, d_i)] = 0$$

That is, in all correlated equilibrium distributions in which c_i is played, d_i is an alternative best response to the conditional probabilities on C_{-i} given c_i. Note that if c_i has zero probability in all correlated equilibrium distributions, then any d_i in C_i jeopardizes c_i. Using complementary slackness properties allows to prove that:

Proposition 4.7 *The strategy d_i jeopardizes c_i if and only if there exists a dual vector α such that $\alpha_i(d_i|c_i) > 0$.*

Thus, there exists a nontrivial dual vector if and only if some strategy is jeopardized by some other strategy.

Definition 4.8 *A correlated equilibrium distribution μ is strict if*

$$\mu(c_i \times C_{-i}) > 0 \Rightarrow \sum_{c_{-i} \in C_{-i}} \mu(c)[U_i(c) - U_i(c_{-i}, d_i)] > 0 \qquad \forall i \in N, \forall c_i \in C_i, \forall d_i \neq c_i$$

A game is *elementary* if it has a strict correlated equilibrium distribution with full support. Myerson [7] shows that a game is elementary if and only if there exists no i, c_i and $d_i \neq c_i$ such that d_i jeopardizes c_i. Thus proposition 4.7 implies:

Corollary 4.9 *A game may be reduced if and only if it is not elementary. By iterative dual reduction, any game is eventually reduced to an elementary game.*

4.4 Full dual reduction

Let us say that two dual reductions Γ/α and Γ/β of the same game are different if $C/\alpha \neq C/\beta$. A game may admit different dual reductions (for instance, if several strategies are weakly dominated). A tentative way to restore uniqueness is to consider only reductions by some special dual vectors, which minimize the number of pure strategies remaining in the reduced game:

Definition 4.10 *A dual vector α is* full *if $\alpha(d_i|c_i) > 0$ for all i in N, and all c_i, d_i in C_i such that d_i jeopardizes c_i.*

Full dual vectors always exist [7]. Actually, almost all dual vectors are full[7].

Definition 4.11 *A* full dual reduction *of Γ is an α-reduced game Γ/α where α is a full dual vector. An* iterative full dual reduction *of depth m of Γ is a game $\Gamma/\alpha^1/\alpha^2/.../\alpha^m$ where m is a positive integer and, for all k in $\{1, 2, ..., m\}$, α^k is a full dual vector of $\Gamma/\alpha^1/\alpha^2/.../\alpha^{k-1}$.*

All full dual vectors α define, for all i, the same minimal α_i-absorbing sets. Thus in all full dual reductions, the same strategies are eliminated and the same strategies are grouped together. A game may nonetheless admit different full dual reductions, because the way these strategies are grouped together may differ quantitatively. We will return to this point in section 6.

[6]This is clear from the basic theory of Markov chains. See for instance [4] and references therein.

[7]The set of dual vectors is a polytope, whose relative interior is non empty if G is not elementary. All dual vectors in the relative interior of this polytope are full. If G is elementary, the only dual vector is trivially full.

5 Other properties of dual reduction

A basic desirable property for a decision-theoretic concept is that it be independent of the specific utility functions chosen to represent the preferences of the agents. So we begin by showing that dual reduction meets this requirement; that is, the ways in which a game may be reduced are unaffected by positive affine transformations of the utility functions. We then extends theorem 4.4 to other equilibrium concepts, including Nash one's, and prove its converse: if a correlated strategy λ of a reduced game induces a correlated (Nash, etc.) equilibrium distribution in the original game, then λ is an equilibrium distribution of the reduced game. We then investigate eliminations of strategies and equilibria. We show that strategies that are weakly dominated (are never played in correlated equilibria; have positive probability in some strict correlated equilibrium) need not be (are always; cannot be) eliminated in full dual reductions. Finally we study some specific classes of games. We show that games that are best-response equivalent to zero-sum games, as well as games with a unique correlated equilibrium distribution are reduced in games with a single strategy profile by full dual reduction. Symmetric games are shown to have symmetric full dual reductions (but possibly also asymmetric ones) and generic 2×2 games are analysed.

In section 6, we show that, even if only full dual reductions are used, there might still be multiple ways to reduce a game. This typically happens when some player is indifferent between some of his strategies: a nongeneric event. We show that generic two-players games have a unique sequence of iterative full dual reductions.

Both in sections 5 and 6, other, minor results are given. The proofs which are neither trivial nor given in the text are gathered in section 8.

5.1 Independence from the choice of utility functions

Proposition 5.1 *Let Γ and Γ' be best response equivalent [11]. Let c_i, d_i be pure strategies of player i in Γ and c_i', d_i' the corresponding strategies of player i in Γ'. The following holds: (i) d_i jeopardizes c_i if and only if d_i' jeopardizes c_i'; (ii) the strategies grouped together (eliminated) in full dual reductions of Γ correspond to the strategies grouped together (eliminated) in full dual reductions of Γ'.*

Proof. (i) is clear from the definitions; (ii) follows immediately from (i) ■

If Γ and Γ' are not only best response equivalent, but rescalings of each other (as defined below), then there is a canonical, one to one correspondence between dual reductions of Γ and dual reductions of Γ':

Proposition 5.2 *For each i in N, let $\phi_i : \mathbb{R} \to \mathbb{R}$ denote a positive affine transformation. That is, such that there exists real numbers $a_i > 0$ and b_i such that $\phi_i(x) = a_i x + b_i$ for all x in $\mathbb{R}$. Let $\phi(\Gamma)$ denote the rescaling of Γ obtained by changing the utility functions from U_i to $\phi_i \circ U_i$:*

$$\phi(\Gamma) = \{N, (C_i)_{i \in N}, (\phi_i \circ U_i)_{i \in N})$$

If Γ/α is a dual reduction of Γ, then $\phi(\Gamma/\alpha)$ is a dual reduction of $\phi(\Gamma)$.

Proposition 5.2 is not trivial because a game and its rescalings need not have the same dual vectors. Indeed, consider a game such as Matching-Pennies, which is nonelementary and in which all pure strategies are undominated:

$$\forall i \in N, \forall c_i \in C_i, \forall \sigma_i \in \Delta(C_i), \sigma_i \neq c_i \Rightarrow \exists c_{-i} \in C_{-i}, U_i(c) > U_i(c_{-i}, \sigma_i)$$

Let α be a nontrivial dual vector: there exist i and c_i such that $\alpha_i * c_i \neq c_i$. Since c_i is not weakly dominated, there exists c_{-i} such that $U_i(\alpha_i * c) - U_i(c) < 0$. Multiplying the payoff of player i by $a_i > 0$ yields a rescaled game Γ' such that:

$$\sum_{j \in N} [U_j'(\alpha_j * c) - U_j'(c)] = a_i[U_i(\alpha_i * c) - U_i(c)] + \sum_{j \neq i} [U_j(\alpha_j * c) - U_j(c)]$$

If a_i is high enough, this expression is negative and α cannot be a dual vector of Γ'. The key is that different deviation vectors may induce the same dual reductions:

Lemma 5.3 *Let α_i (α_i^{id}) be a (the trivial) deviation plan for player i. For any $0 \leq \epsilon \leq 1$, let $\alpha^\epsilon = \epsilon\alpha_i + (1-\epsilon)\alpha_i^{id}$. If ϵ is positive then $C_i/\alpha_i = C_i/\alpha_i^\epsilon$.*

Proof. For any mixed strategy σ_i in $\Delta(C_i)$, $\alpha_i^\epsilon * \sigma_i - \sigma_i = \epsilon(\alpha_i * \sigma_i - \sigma_i)$. ■

5.2 Extension and converse of theorem 4.4

In this section, we first define three equilibrium concepts introduced in [12] and [13]. We then show that theorem 4.4 extends to Nash equilibrium distributions[8], and to these other equilibrium concepts. We illustrate this by an example. Finally, we prove a converse of theorem 4.4.

Let $\mu \in \Delta(C)$ and $c_i \in C_i$. If $\mu(c_i \times C_{-i}) > 0$, let $\mu(.|c_i)$ denote the conditional probability on C_{-i} given c_i:

$$\mu(c_{-i}|c_i) = \mu(c_{-i}, c_i)/\mu(c_i \times C_{-i})$$

Definition 5.4 *The correlated strategy $\mu \in \Delta(C)$ is an equalizing distribution [13] if*

$$\mu(c_i \times C_{-i}) > 0 \Rightarrow \sum_{c_{-i} \in C_{-i}} \mu(c_{-i}|c_i)U_i(c) = U_i(\mu) \quad \forall i \in N, \forall c_i \in C_i, \quad \textit{where } U_i(\mu) = \sum_{c \in C} \mu(c)U_i(c)$$

That is, in an equalizing distribution, the expected payoff given a pure strategy is independent of this strategy.

Definition 5.5 *The correlated strategy $\mu \in \Delta(C)$ is an* equalizing correlated equilibrium distribution[9] *[12] (henceforth* equalizing c.e.d.*) if μ is both an equalizing and a correlated equilibrium distribution*[10].

Definition 5.6 *The correlated strategy $\mu \in \Delta(C)$ is a* stable matching distribution[11] *[12],[13] if*

$$\mu_i(c_i \times C_{-i})\mu_i(d_i \times C_{-i}) > 0 \Rightarrow \sum_{c_{-i} \in C_{-i}} [\mu(c_{-i}|c_i) - \mu(c_{-i}|d_i)]U_i(c) \geq 0 \qquad \forall i \in N, \forall c_i \in C_i, \forall d_i \in C_i$$

That is, c_i yields a higher expected payoff against the correlated strategy $\mu(.|c_i)$ of the players other than i than against $\mu(.|d_i)$.

Proposition 5.7 *Let λ be a correlated strategy of an iterative dual reduction Γ_r of Γ. If λ is an* equilibrium distribution *of Γ_r then the Γ-equivalent correlated strategy is an* equilibrium distribution *of Γ, where* equilibrium distribution *may stand for: Nash equilibrium distribution, equalizing distribution, equalizing c.e.d. or stable matching distribution.*

The following example illustrates proposition 5.7:

Example 5.8

	x_2	y_2	z_2
x_1	$2,0$	$0,2$	$0,-3$
y_1	$0,1$	$1,0$	$0,0$
z_1	$-3,0$	$0,0$	$1,1$

	σ_{B_2}	z_2
σ_{B_1}	$2/3, 2/3$	$0,-1$
z_1	$-1,0$	$1,1$

[8] The extension to Nash equilibrium distributions has been independently noted by Myerson

[9] Sorin [12] uses the expression *distribution equilibrium*

[10] Any Nash equilibrium distribution is an equalizing c.e.d. but the converse is false. See example 5.8.

[11] Sorin [12] uses the expression *dual correlated equilibrium*

Let Γ denote the game on the left. Consider the deviation vector α such that for $i = 1, 2$:

$$\alpha_i(x_i|x_i) = 2/3, \ \ \alpha_i(y_i|x_i) = 1/3; \ \ \alpha_i(x_i|y_i) = 1/6, \ \ \alpha_i(y_i|y_i) = 5/6; \ \ \alpha_i(z_i|z_i) = 1,$$

and all other $\alpha_i(d_i|c_i)$ are zero. α is a dual vector. The minimal α_i-absorbing sets are $B_i = \{x_i, y_i\}$ and $B_i' = \{z_i\}$. The α-reduced game Γ/α is the game on the right, where the α_i-invariant strategy σ_{B_i} is $(\frac{1}{3};\ \frac{2}{3};\ 0)$. Consider the distribution λ on C/α (below, right).[12] *This is an equalizing c.e.d. of Γ/α. Therefore, the Γ-equivalent distribution $\bar{\lambda}$ (below, left) is an equalizing c.e.d. of Γ.*

$\bar{\lambda} =$

$1/24$	$1/12$	$1/24$
$1/12$	$1/6$	$1/12$
$1/24$	$1/12$	$3/8$

$\lambda =$

$3/8$	$1/8$
$1/8$	$3/8$

Theorem 4.4 states that correlated equilibrium distributions of Γ/α induce correlated equilibrium distributions in Γ. We may wonder whether a correlated strategy of Γ/α, which is not a correlated equilibrium distribution, might nonetheless induce a correlated equilibrium distribution in Γ. The answer is negative:

Lemma 5.9 *Given any deviation vector α, a distribution $\bar{\lambda} \in \Delta(C)$ is α-invariant if and only if it is* Γ-equivalent *to a distribution $\lambda \in \Delta(C/\alpha)$. Such a λ is then unique.*

Proposition 5.10 *Let α denote a dual vector. Let $\bar{\lambda}$ denote an α-invariant distribution on C and λ the corresponding distribution on C/α. Then $\bar{\lambda}$ is an* equilibrium distribution *of Γ if and only if λ is an* equilibrium distribution *of Γ/α, where* equilibrium distribution *may stand for: Nash equilibrium distribution, correlated equilibrium distribution, equalizing distribution, equalizing c.e.d. or stable matching distribution.*

5.3 Elimination of strategies and equilibria

A first result is a converse of proposition 4.3:

Proposition 5.11 *Let $c_i \in C_i$; assume that there exists a dual vector α such that $c_i \notin C_i/\alpha_i$ and $C_j/\alpha_j = C_j$ for all j in $N - i$. Then there exists $\sigma_i \neq c_i$ in $\Delta(C_i)$ such that $U_i(c_{-i}, \sigma_i) \geq U_i(c)$ for all c_{-i} in C_{-i}.*

Proof. Let $\sigma_i = \alpha_i * c_i$. For all $j \neq i$, all strategies c_j in C_j are α_j-invariant. Thus (6) yields $U_i(c_{-i}, \sigma_i) \geq U_i(c)\ \forall c_{-i} \in C_{-i}$. Furthermore $c_i \notin C_i/\alpha_i$ hence c_i cannot be α_i-invariant and $\sigma_i \neq c_i$ ∎

Thus, only if a strategy is dominated does there exists a dual reduction that simply consists in eliminating this strategy. Note that if a strategy is weakly dominated it is eliminated in some dual reductions (proposition 4.3), but not necessarily in full dual reductions:

Example 5.12

	x_2	y_2
x_1	$1,1$	$1,0$
y_1	$1,0$	$0,0$

In the above game, μ is a correlated equilibrium distribution if and only if y_2 is not played in μ. That is, $\mu(x_1, y_2) = \mu(y_1, y_2) = 0$. Therefore y_1 jeopardizes x_1, and reciprocally. Thus, in all full dual reductions, x_1 and y_1 must be grouped together hence y_1 is not eliminated.

[12] We represent correlated strategies in tables. For instance, $\lambda(\sigma_{B_1}, z_2) = 1/8$.

This raises the following questions: except strictly dominated strategies, are there other classes of strategies that are always eliminated in full dual reductions ? A partial answer is the following:

Proposition 5.13 *(i) Let $c \in C$. Assume that c has probability zero in all correlated equilibrium distributions. In full dual reductions c is eliminated; hence there exists i in N such that, in all full dual reductions, c_i is eliminated. (ii) Let $i \in N, c_i \in C_i$. Assume that c_i has marginal probability zero in all correlated equilibrium distributions. Then c_i is eliminated in all full dual reductions.*

Proof. First note that (i) implies (ii). Indeed, let $\sigma_i \in C_i/\alpha_i$ and $\sigma_{-i} \in (C/\alpha)_{-i}$. If $\mu(c) = 0$ for all correlated equilibrium distributions μ and all c_{-i} in C_{-i} then, by (i), $\sigma(c) = \sigma_i(c_i)\sigma_{-i}(c_{-i}) = 0$ for all $c_{-i} \in C_{-i}$ implying $\sigma_i(c_i) = 0$. Point (i) is proved in section 8 ■

Let Γ^* denote the game obtained from Γ by deleting all pure strategies that have marginal probability zero in all correlated equilibrium distributions. Proposition 5.13 suggests that Γ and Γ^* have the same full dual reductions, but this is not so:

Example 5.14

	x_2	y_2
x_1	$1,1$	$0,1$
y_1	$0,1$	$1,0$

	x_2	y_2
x_1	$1,1$	$0,1$

Let Γ denote the left game. Then Γ^* is the game on the right. In Γ^* any mixed strategy profile is a Nash equilibrium. In Γ, a mixed strategy profile σ is a Nash equilibrium if and only if $\sigma_1(y_1) = 0$ and $\sigma_2(y_2) \leq 1/2$. In any full dual reduction of Γ or Γ^* there is a single strategy profile. If σ is a Nash equilibrium of Γ (Γ^*) then there exists a full dual vector α of Γ (Γ^*) such that $C/\alpha = \sigma$ ($C^*/\alpha = \sigma$) if and only if $\sigma(y_2)$ and $\sigma(x_2)$ are positive. Thus the set of full dual reductions of Γ is strictly included in the set of full dual reductions of Γ^*.

We now shift our attention to elimination of equilibria. Since dual reduction includes elimination of dominated strategies as a subprocess, it is clear that dual reduction may eliminate Nash equilibria. Nash equilibria may also be eliminated as strategies are grouped together (see for instance [7, fig. 7]). We show in section 7 that completely mixed, hence perfect Nash equilibria may be eliminated in full dual reductions. In contrast:

Proposition 5.15 *Strict correlated equilibrium distributions cannot be eliminated, not even in an iterative dual reduction.*

Proof. If μ is a strict correlated equilibrium distribution, a strategy that has positive marginal probability in μ cannot be jeopardized by another strategy. Thus, in any dual reduction Γ/α of Γ all the strategies used in μ must be available. Furthermore, as the player's options are more limited in Γ/α than in Γ, μ is a fortiori a strict correlated equilibrium distribution of Γ. Inductively, in any iterative dual reduction $\Gamma/\alpha^1/.../\alpha^m$ of Γ, all strategies used in μ are available and μ is still a strict correlated equilibrium distribution ■

The proof shows that a pure strategy that has positive marginal probability in some strict correlated equilibrium distribution can never be eliminated nor grouped with other strategies.

5.4 Some classes of games

5.4.1 Games with a unique correlated equilibrium distribution

If Γ has a unique Nash equilibrium σ, then any iterative dual reduction of Γ has a unique Nash equilibrium, which induces σ in Γ; but the strategy space need not be reducible to σ: counterexamples are [5, p.204] and [10, example 4]. In contrast,

Proposition 5.16 *Assume that Γ has a unique correlated equilibrium distribution σ. Then σ is a Nash equilibrium distribution, hence it may be seen as a mixed strategy profile. Let Γ_r be the reduced game in which the only strategy profile is σ and the payoff for player i is $U_i(\sigma)$. Any full (resp. elementary iterative) dual reduction of Γ is equal to Γ_r. In particular, Γ has a unique full dual reduction.*

5.4.2 Zero-sum games

Claim 5.17 *Any iterative dual reduction of a zero-sum game is a zero-sum game with the same value.*

Proof. Conservation of the zero-sum property is immediate. Conservation of the value comes from theorem 4.4 and the fact that in a two-player zero-sum game, any correlated equilibrium payoff equals the value of the game ■

Proposition 5.18 *Let Γ denote a two-player zero-sum game and α a deviation vector. (i) If for all $i = 1, 2$ and for all c_i in C_i, $\alpha_i * c_i$ is an optimal strategy of player i, then α is a dual vector; (ii) If furthermore, $\alpha_i * c_i$ is the same optimal strategy σ_i for all c_i in C_i, then $C_i/\alpha_i = \sigma_i$ (iii) in any elementary iterative reduction of Γ there is a unique strategy profile, which is a product of optimal strategies of Γ.*

Proposition 5.19 *If Γ is best response equivalent to a two-player zero-sum game then: (i) for any i in N, any (pure) strategy c_i which has positive marginal probability under some correlated equilibrium distribution jeopardizes all other strategies of player i; (ii) in all full dual reductions of Γ all the strategies of player i that have positive probability in some correlated equilibrium distribution are grouped together and his other strategies are eliminated hence (iii) there is a unique strategy profile σ. (iv) This strategy profile corresponds to a product of optimal strategies in the underlying zero-sum game.*

Proof. σ must be equivalent to a Nash equilibrium of Γ. This allows to prove (iv). Point (iii) follows from (ii) and proposition 5.13; (ii) follows from (i) and (i) is proved in [14]. ■

If Γ is zero-sum with value v, then the payoffs in any full dual reduction of Γ must be $(v, -v)$. In contrast, if Γ is only best response equivalent to a zero sum game, then the payoffs in a full dual reduction of Γ may depend on the full dual reduction:

Example 5.20

	x_2	y_2	z_2
x_1	$0, 0$	$0, 0$	$0, 0$
y_1	$0, 0$	$1, -1$	$-1, 1$
z_1	$0, 0$	$-1, 1$	$1, -1$

	x_2	y_2	z_2
x_1	$1, 1$	$0, 1$	$0, 1$
y_1	$1, 0$	$1, -1$	$-1, 1$
z_1	$1, 0$	$-1, 1$	$1, -1$

Let Γ (Γ') denote the game on the left (right). Γ is zero-sum and Γ' is best response equivalent to Γ. The proof of proposition 5.2 shows that Γ and Γ' have the same dual vectors. For $0 \leq \epsilon \leq 1$, let σ_i^ϵ denote the optimal strategy of player i such that: $\sigma_i^\epsilon(x_i) = \epsilon$ and $\sigma_i^\epsilon(y_i) = \sigma_i^\epsilon(z_i) = (1-\epsilon)/2$. Let $\alpha^{\epsilon,\eta}$ denote the deviation vector such that: $\alpha_1 * x_1 = \alpha_1 * y_1 = \alpha_1 * z_1 = \sigma_1^\epsilon$ and $\alpha_2 * x_2 = \alpha_2 * y_2 = \alpha_2 * z_2 = \sigma_2^\eta$. By proposition 5.18 α is a dual vector of Γ, hence of Γ'. If $0 < \epsilon < 1$ and $0 < \eta < 1$, α is full, the reduced strategy space $C'/\alpha^{\epsilon,\eta}$ is the singleton $(\sigma_1^\epsilon, \sigma_2^\eta)$ and the associated payoff is (η, ϵ).

5.4.3 Symmetric Games

In section 8 we recall the definition of a symmetric game and prove the following:

Proposition 5.21 *Let Γ be a symmetric game. There exists a full dual vector α such that Γ/α is symmetric.*

Example 5.8 shows that a nonsymmetric game may also have symmetric full dual reductions, even if all strategies are undominated. The following example shows that a symmetric game may have nonsymmetric full dual reductions:

Example 5.22

	x_2	y_2
x_1	$1,1$	$0,1$
y_1	$1,0$	$0,0$

In the above symmetric game Γ, any deviation vector is a dual vector. In any full dual reduction, the reduced strategy space is a singleton. For any $0 < \epsilon < 1$, $0 < \eta < 1$, there exists a full dual reduction in which the payoff is (ϵ, η). If $\epsilon \neq \eta$, this full dual reduction is nonsymmetric.

5.4.4 Generic 2×2 games

Proposition 5.23 *Let Γ be a 2×2 game such that a player is never indifferent between two different strategy profiles. That is, for all c, c' in C and all $i = 1, 2$: $c \neq c' \Rightarrow U_i(c) \neq U_i(c')$. Then either Γ is elementary or Γ has a unique correlated equilibrium distribution (in which case proposition 5.16 apply).*

Proof. Straightforward computations. The first case corresponds to games with three Nash equilibria: two pure and one completely mixed; the second case to games with either a dominating strategy or a unique, completely mixed Nash equilibrium. ■

6 The issue of uniqueness

As shown by example 5.22, a game may have several full dual reductions. This ambiguity arises naturally when a player is indifferent between some of his strategies:

Proposition 6.1 *Assume that player i is indifferent between c_i and d_i, i.e. $U_i(c) = U_i(c_{-i}, d_i)$ for all c_{-i} in C_{-i}. Then (i) for any $0 \leq \epsilon \leq 1$ there exists a dual reduction that simply consists in grouping c_i and d_i in the strategy σ_i such that $\sigma_i(c_i) = \epsilon$ and $\sigma_i(d_i) = 1 - \epsilon$; (ii) if c_i is not eliminated in full dual reductions, then there exists an infinity of full dual reductions.*

Proof. To prove (i) take as dual vector α: $\alpha_i(c_i|c_i) = \alpha_i(c_i|d_i) = \epsilon$, $\alpha_i(d_i|c_i) = \alpha_i(d_i|d_i) = 1 - \epsilon$ and all the other $\alpha_j(d_j|c_j)$ as in the trivial deviation vector. (ii) is proved in section 8 ■

A similar difficulty may arise if a player is indifferent between a pure and a mixed strategy (example 5.20) or if a player *becomes* indifferent between some of his strategies, after strategies of some other player have been eliminated (example 5.14). These are non-generic phenomena. We prove in this section that, for any positive integer m, two-player games generically have a unique iterative full dual reduction of depth m. We first show that there are severe restrictions on the ways strategies may be grouped together in dual reductions:

Notation: for all i in N, let $B_i \subset C_i$ and let $B = \times_{i \in N} B_i$. Then $\Gamma_B = (N, (B_i)_{i \in N}, (U_i)_{i \in N})$ denote the game obtained from Γ by reducing player i's pure strategy set to B_i, for all i in N.

Proposition 6.2 *Let α be a dual vector. For each i in N, let $B_i \subset C_i$ denote a minimal α_i-absorbing set and $B = \times_{i \in N} B_i$. Let σ_{B_i} denote the unique α_i-invariant strategy of player i with support in B_i and $\sigma_B = (\sigma_{B_i})_{i \in N}$. We have: σ_B is a completely mixed Nash equilibrium of Γ_B.*

Proof. First, the support of $\sigma_{B_{-i}}$ is exactly B_i so σ_B is completely mixed. Second, let $\sigma_{B_{-i}} = \times_{j \in N-i} \sigma_{B_j}$. Against $\sigma_{B_{-i}}$, player i is indifferent between the strategies of the minimal absorbing set B_i (proposition 4.5). Therefore, if player i is restricted to the strategies of B_i, σ_{B_i} is a best response to $\sigma_{B_{-i}}$ ■

Assume now that α is full. If Γ_B has a unique completely mixed Nash equilibrium, then for any full dual vector β, the β_i-invariant strategy with support in B_i must be σ_{B_i}. So proposition 6.2 has the following corollary:

Corollary 6.3 *If for any product $B = \times_{i \in N} B_i$ of subsets B_i of C_i, Γ_B has at most one completely mixed Nash equilibrium, then there exists a unique full dual reduction.*

In the remaining of this section, Γ is a two-player, bimatrix game. To show that, generically, two-player games have a unique sequence of iterative full dual reductions, we need to introduce some suitable notions of genericity:

Definition 6.4 Γ *is* generic *if for all Nash equilibria σ the supports of σ_1 and σ_2 have same cardinal*[13]. Γ *is locally generic if it is generic and if any game obtained from Γ by deleting some pure strategies is generic.*

Definition 6.5 Γ *is 2-generic if for any subset B_1 of C_1 and for any disjoint subsets B_2 and B_2' of C_2: if σ and σ' are respectively completely mixed Nash equilibria of $\Gamma_{B_1 \times B_2}$ and $\Gamma_{B_1 \times B_2'}$ then $\sigma_1 \neq \sigma_1'$. That is, the same mixed strategy cannot be a completely mixed Nash equilibrium strategy of player 1 both on $B_1 \times B_2$ and on $B_1 \times B_2'$. The notion of 1-genericity is defined similarly. A bimatrix game is $*$-generic if it is both 1-generic and 2-generic.*

A bimatrix game in which players 1 and 2 have respectively p and q pure strategies is given by two $p \times q$ payoff matrices, thus it may be viewed as a point in $\mathbb{R}^{pq} \times \mathbb{R}^{pq}$. It may be shown that the set of $p \times q$ bimatrix games which are both locally generic and $*$-generic contains an open, dense subset of $\mathbb{R}^{pq} \times \mathbb{R}^{pq}$. The two next propositions follow from proposition 6.2:

Proposition 6.6 *A locally generic bimatrix game has a unique full dual reduction.*

Proof. Locally generic bimatrix games check the conditions of corollary 6.3 ■

Proposition 6.7 *If Γ is both locally generic and $*$-generic, there are only three possibilities:*

1 Γ is elementary

2 In all dual reductions of Γ, some strategies are eliminated, but no strategies are grouped together.

3 In any full dual reduction of Γ the reduced strategy space C/α is a singleton.

As an immediate corollary of proposition 6.7 and definitions 6.4 and 6.5 we get:

Corollary 6.8 *If Γ is both locally generic and $*$-generic then any dual reduction of Γ is both locally generic and $*$-generic.*

As an immediate corollary of proposition 6.6 and corollary 6.8 we get:

Theorem 6.9 *If Γ is both locally generic and $*$-generic, then for any positive integer m, Γ has a unique iterative full dual reduction of depth m.*

[13] Any game which is nondegenerate in the sense of [15, def. 2.6 and thm 2.10] is generic in this sense

7 Dual reduction and elimination of unacceptable pure strategies

Dual reduction and elimination of unacceptable pure strategies [8] both generalize elimination of dominated strategies. Furthermore, there are similarities in the ways these concepts are defined.[14] Comparing dual reduction and elimination of unacceptable pure strategies is thus quite natural. In this section we show by means of example that none of these refinement concepts is more stringent than the other.[15]

Lemma 7.1 *If there exists a correlated equilibrium distribution with full support then all pure strategies are acceptable and predominant.*

Lemma 7.1 implies that the class of games in which all pure strategies are acceptable is strictly larger than the class of elementary games. This is not only due to the fact that in a game in which all strategy profiles are coherent, as in Matching-Pennies, dual reduction can still group strategies together. Indeed, consider the following game:

Example 7.2

	x_2	y_2	z_2
x_1	$0,0$	$0,0$	$0,0$
y_1	$0,0$	$1,1$	$-1,-1$
z_1	$0,0$	$-1,-1$	$1,1$

In this game, playing each strategy with equal probability is a completely mixed Nash equilibrium. Thus all strategies are acceptable and predominant. However, x_1 and x_2 are eliminated in any full dual reduction, and in any nontrivial dual reduction at least one of x_1 and x_2 is eliminated (to prove this, note that (i) by proposition 4.3, x_i must be either eliminated or grouped with other strategies in all full dual reductions; (ii) y_i and z_i must be invariant under any dual vector because they have positive probability in some strict correlated equilibrium distribution).

This example shows that dual reduction may eliminate acceptable and even predominant pure strategies. It also shows that dual reduction can eliminate completely mixed, hence perfect Nash equilibria. Since any perfect Nash equilibrium is a perfect direct correlated equilibrium [2], it shows that dual reduction may eliminate perfect direct correlated equilibrium distributions.

The next example shows that there may be unacceptable pure strategies that no dual reduction eliminates: let Γ denote the following three person game, where player 1 chooses the matrix (x_1 or y_1), player 2 the row, and player 3 the column.

Example 7.3 (taken from [8])

x_1

	x_3	y_3	z_3
x_2	$2,1,1$	$0,2,0$	$0,2,0$
y_2	$0,0,2$	$0,3,0$	$0,0,3$
z_2	$0,0,2$	$0,0,3$	$0,3,0$

y_1

	x_3	y_3	z_3
x_2	$1,3,3$	$1,3,3$	$1,3,3$
y_2	$1,3,3$	$1,3,3$	$1,3,3$
z_2	$1,3,3$	$1,3,3$	$1,3,3$

Myerson [8] shows that the only acceptable strategies for player i is x_i, for all i in $\{1,2,3\}$. However, y_1 cannot be eliminated by one-shot dual reduction. Indeed, since $c = (y_1, y_2, y_3)$ is a Nash equilibrium and that all unilateral deviations from c by player 1 are strictly detrimental for him, y_1 must be invariant under any dual vector.

[14] In particular, the *aggregate incentive value* of c for the set of players N, $V_N(c,\alpha)$, defined in [8, p.141, (3.3)], is exactly the payoff $g(c,\alpha)$ defined in section 3.

[15] For definitions and properties of acceptable (predominant) pure strategies, acceptable (predominant) correlated equilibria and codomination systems, see [8], [9] or [2].

Note that y_1 may be eliminated by *iterative* dual reduction. Actually, to prove that y_2, z_2, y_3, z_3 and y_1 are unacceptable, Myerson uses the codomination system (α^1, α^2) where α^1 and α^2 are the deviation vectors such that:

$$\alpha_i^1(x_i|y_i) = \alpha_i^1(x_i|z_i) = 1 \; \forall i \in \{2,3\}, \qquad \alpha_1^2(x_1|y_1) = 1,$$

and all other $\alpha_i^k(d_i|c_i)$ are as in the corresponding trivial deviation vectors. It is easy to check that α^1 is a dual vector of Γ and α^2 a dual vector of Γ/α^1. The only strategy profile remaining in $\Gamma/\alpha^1/\alpha^2$ is the strict Nash equilibrium (x_1, x_2, x_3), thus y_1 has been eliminated. Whether some unacceptable (or non predominant) pure strategies cannot be eliminated by any iterative dual reduction is still an open problem.

8 Proofs

In the proofs we may write c.e.d. for correlated equilibrium distribution.

Proof of proposition 5.2: Let α be a dual vector of Γ. Let $a_k = \min_{i \in N} a_i$ and, for each i in N, let $\epsilon_i = a_k/a_i$. Let $\phi(\alpha)$ denote the deviation vector whose i^{th} component is $\alpha_i^{\epsilon_i}$, defined in lemma 5.3. Let g and g_ϕ denote the payoff functions in the auxiliary games associated respectively to Γ and $\phi(\Gamma)$. We have:

$$g_\phi(c, \phi(\alpha)) = a_k \times g(c, \alpha) \geq 0 \quad \forall c \in C$$

Thus $\phi(\alpha)$ is a dual vector of $\phi(\Gamma)$. Furthermore lemma 5.3 implies: $\phi(\Gamma)/\phi(\alpha) = \phi(\Gamma/\alpha)$. Thus $\phi(\Gamma/\alpha)$ is a dual reduction of $\phi(\Gamma)$. The result still holds if we allow the constants b_i to depend on c_{-i}. Indeed, if the payoff functions $(U_i^\phi)_{i \in N}$ in the rescaled game $\phi(\Gamma)$ are of the slightly more general form: $U_i^\phi(c) = a_i \times U_i(c) + b_i(c_{-i})$ with $a_i > 0$ and $b_i : C_{-i} \to \mathbb{R}$, then the same proof shows that for any dual vector α of Γ, $\phi(\Gamma/\alpha)$ is a dual reduction of $\phi(\Gamma)$.

Proof of proposition 5.7 Notations and preliminary remarks: let $\lambda \in \Delta(C/\alpha)$ and let $\overline{\lambda} \in \Delta(C)$ be Γ-equivalent to λ. Let $c_i, d_i \in C_i$ check $\overline{\lambda}(c_i \times C_{-i})\overline{\lambda}(d_i \times C_{-i}) > 0$. There exist minimal α_i-absorbing sets B_i and B_i' such that c_i belongs to B_i and d_i to B_i'. Let σ_{c_i} (σ_{d_i}) be the α_i-invariant strategy with support in B_i (B_i'). Necessarily, $\lambda(\sigma_{c_i} \times (C/\alpha)_{-i})$ and $\lambda(\sigma_{d_i} \times (C/\alpha)_{-i})$ are positive. Note that: (i) $U_i(\lambda) = U_i(\overline{\lambda})$ and (ii) $\overline{\lambda}(.|c_i)$ is the conditional probability induced on C_{-i} by $\lambda(.|\sigma_{c_i})$. That is, if, for all j in $N - i$, c_j is α_j-recurrent, then:

$$\overline{\lambda}(c_{-i}|c_i) = \lambda(\sigma_{c_{-i}}|\sigma_{c_i}) \left(\prod_{j \in N-i} \sigma_{c_j}(c_j) \right) \quad \text{where } \sigma_{c_{-i}} = \times_{j \in N-i} \sigma_{c_j}$$

Otherwise, $\lambda(c_{-i}|c_i) = 0$. The proofs are now easy:

Nash equilibrium: it follows from (7) that if λ is an independent distribution, then so is $\overline{\lambda}$. This and theorem 4.4 imply that if λ is both an independent and a correlated equilibrium distribution, i.e. a Nash equilibrium distribution, then so is $\overline{\lambda}$.

Equalizing distributions and equalizing c.e.d.: Using (i) and (ii) we get:

$$\sum_{\sigma_{-i} \in (C/\alpha)_{-i}} \lambda(\sigma_{-i}|\sigma_{c_i}) U_i(\sigma_{-i}, \sigma_{c_i}) = U_i(\lambda) \Rightarrow \sum_{c_{-i} \in C_{-i}} \overline{\lambda}(c_{-i}|c_i) U_i(c) = U_i(\overline{\lambda})$$

Thus if λ is an equalizing distribution, then so is $\overline{\lambda}$. This and theorem 4.4 imply that if λ is an both an equalizing and a correlated equilibrium distribution, then so is $\overline{\lambda}$.

Stable matching distributions: Using (ii) we get:

$$\sum_{\sigma_{-i}\in(C/\alpha)_{-i}}[\lambda(\sigma_{-i}|\sigma_{c_i})-\lambda(\sigma_{-i}|\sigma_{d_i})]U_i(\sigma_{-i},\sigma_{c_i})\geq 0\Rightarrow\sum_{c_{-i}\in C_{-i}}[\overline{\lambda}(c_{-i}|c_i)-\overline{\lambda}(c_{-i}|d_i)]U_i(c)\geq 0$$

Thus if λ is a stable matching distribution, then so is $\overline{\lambda}$.

Proof of lemma 5.9: let $\overline{\lambda}\in\Delta(C)$. We only need to show that if $\overline{\lambda}$ is α-invariant then it is Γ-equivalent to a correlated strategy of Γ/α. Indeed, the converse is clear by linearity of $\overline{\lambda}\to\alpha_i*\overline{\lambda}$. Furthermore, letting $C/\alpha_i=C_i/\alpha_i\times C_{-i}$, it is enough to show that if $\overline{\lambda}$ is α_i-invariant then there exists $\lambda\in\Delta(C/\alpha_i)$ such that (i) $\overline{\lambda}$ is Γ-equivalent to λ and (ii) if $\overline{\lambda}$ is α_j-invariant, then so is λ. Indeed, as the number of players is finite, a simple induction then proves the property. So let us assume $\alpha_i*\overline{\lambda}=\overline{\lambda}$:

$$\alpha_i*\overline{\lambda}(c_{-i},c_i)=\sum_{d_i\in C_i}\alpha_i(c_i|d_i)\overline{\lambda}(c_{-i},d_i)=\overline{\lambda}(c_{-i},c_i)\ \ \forall c_i\in C_i,\forall c_{-i}\in C_{-i}$$

This means that, for all c_{-i} in C_{-i}, the vector $[\overline{\lambda}(c_{-i},c_i)]_{c_i\in C_i}$ is α_i-invariant. Therefore: (a) $\overline{\lambda}(c_i\times C_{-i})=0$ if c_i is α_i-transient and (b) for any minimal α_i-absorbing set B_i, $[\overline{\lambda}(c_{-i},c_i)]_{c_i\in B_i}$ is proportional to $[\sigma_{B_i}(c_i)]_{c_i\in B_i}$, where σ_{B_i} is the unique α_i-invariant strategy with support in B_i. More precisely, letting $\lambda(c_{-i},\sigma_{B_i})=\sum_{c_i\in B_i}\overline{\lambda}(c_{-i},c_i)$, we have:

$$\overline{\lambda}(c_{-i},c_i)=\lambda(c_{-i},\sigma_{B_i})\times\sigma_{B_i}(c_i)\ \ \forall c_i\in B_i,\forall c_{-i}\in C_{-i}$$

λ defines an element of $\Delta(C/\alpha_i)$ and the above equality means that $\overline{\lambda}$ is Γ-equivalent to λ. Finally it is straightforward to check that if $\overline{\lambda}$ is α_j-invariant, then so is λ. This completes the proof.

Proof of proposition 5.10: we prove proposition 5.10 for correlated equilibrium distributions. The other proofs are similar. Assume that λ is not a c.e.d.. Then there exist i in N and σ_i,τ_i in C_i/α_i such that σ_i has positive probability under λ but τ_i is a strictly better response than σ_i to $\lambda(.|\sigma_i)$. If $c_i\in C_i$ belong to the support of σ_i, player i is indifferent between c_i and σ_i against $\lambda(.|\sigma_i)$ (proposition 4.5), hence τ_i is a strictly better response than c_i to $\lambda(.|\sigma_i)$. Finally, $\overline{\lambda}(c_i\times C_{-i})>0$ and $\overline{\lambda}(.|c_i)$ is Γ-equivalent to $\lambda(.|\sigma_i)$. Therefore τ_i is a strictly better response than c_i to $\overline{\lambda}(.|c_i)$ hence $\overline{\lambda}$ is not a c.e.d.

Proof of proposition 5.13, point (i): first recall that the same strategies and strategy profiles are eliminated in all full dual reductions. So we only need to prove that the results hold for some full dual reduction. Step 1: Assume that $\mu(c)=0$ for all c.e.d. μ of Γ. Then it follows from [10, page 432 and Proposition 2] that there exists a dual vector α such that $g(c,\alpha)<0$. Since $g(d,\alpha)<0$ for all d in C, this implies that if c has positive probability in some correlated strategy μ then $g(\mu,\alpha)<0$. Step 2: we may assume α full (otherwise, replace α by some strictly convex combination of α and some full dual vector). If σ belongs to C/α, then σ is α-invariant thus $g(\sigma,\alpha)=0$ by (5). Hence c cannot have positive probability in σ. Since this holds for all σ in C/α, c has been eliminated in the full dual reduction Γ/α. Finally, c_i must have been eliminated for some i, otherwise c would not have been eliminated.

Proof of proposition 5.16: consider first an elementary iterative dual reduction Γ_e of Γ. Since Γ_e is elementary, Γ_e has a strict c.e.d. with full support σ^e. Since Γ has a unique c.e.d., Γ_e has a unique c.e.d. too, thus σ^e is actually a Nash, hence a strict Nash equilibrium. So σ^e is pure. But σ^e has full support. So σ^e is the only strategy profile. Finally, σ^e must be Γ-equivalent to σ, hence $\Gamma_e=\Gamma_r$.
Consider now a full dual reduction Γ/α of Γ. By proposition 5.13, the strategies that do not participate in σ are eliminated in Γ/α. For each i in N, the strategies of player i that participate to σ_i jeopardize each other and thus must be grouped in a single mixed strategy. Finally, the unique strategy profile of Γ/α must be equivalent to σ, hence $\Gamma/\alpha=\Gamma_r$.

Proof of proposition 5.18. Proof of (i): let $c \in C$. By optimality of $\alpha_1 * c_1$, $U_1(\alpha_1 * c_1, c_2) \geq v$, where v is the value of the game. Similarly, $U_2(c_1, \alpha_2 * c_2) \geq -v$. Since $U_1(c)+U_2(c) = 0$, $\sum_{i=1,2}[U_i(c_{-i}, \alpha_i * c_i) - U_i(c)] \geq 0$. That is, $g(c, \alpha) \geq 0$. Since this holds for all c in C, α is a dual vector.
Proof of (ii): assume that there exists $\sigma_i \in \Delta(C_i)$ such that $\alpha_i * c_i = \sigma_i$ for all c_i in C_i. Then the only α_i-invariant strategy is σ_i. Therefore, $C_i/\alpha_i = \{\sigma_i\}$.
Proof of (iii): The above implies that any two-player zero-sum game whose set of strategy profiles is not a singleton can be further reduced. Together with claim 5.17, this implies that in any elementary iterative dual reduction of Γ, there is a unique strategy profile. This strategy profile induces a Nash equilibrium in Γ. Therefore it must be (equivalent to) a product of optimal strategies of Γ.

Definition of symmetric games and proof of proposition 5.21: let Γ be a game in which all players have the same number m of pure strategies. Let $c_{i,k}$ denote the k^{th} strategy of player i. Thus $C_i = \{c_{i,1}, ..., c_{i,m}\}$. For all i in N, let k_i be an integer in {1,...,m}. Let $(c_{i,k_i})_{i \in N}$ denote the profile of strategy in which, for all i, player i plays his k_i^{th} strategy. Γ is a *symmetric game* if for all permutations p of the set of players,

$$U_i((c_{j,k_{p(j)}})_{j \in N}) = U_{p(i)}((c_{j,k_j})_{j \in N})$$

This means that if, for all i, player i plays as player $p(i)$ used to play, then the payoff of player i in the new configuration is the payoff of player $p(i)$ in the old configuration. We now prove the proposition:

Step 1: let us say that a deviation vector α of a symmetric game is symmetric if $\alpha_i(c_{i,k'}|c_{i,k}) = \alpha_j(c_{j,k'}|c_{j,k})$ for all i, j in N and all k, k' in {1,2,..,m}. It is clear that if Γ is a symmetric game and α a symmetric dual vector, then Γ/α is a symmetric game. So it is enough to show that there exists a symmetric full dual vector.

Step 2: let α denote a deviation vector. For all permutations p of the set of players, let α^p denote the deviation vector such that:

$$\alpha^p_{p(i)}(c_{p(i),k'}|c_{p(i),k}) = \alpha_i(c_{i,k'}|c_{i,k}) \qquad \forall i \in N$$

Let $\bar{\alpha}$ denote the symmetrized deviation vector given by:

$$\bar{\alpha} = \frac{\sum_p \alpha^p}{n!}$$

where n is the number of players and the summation is taken over all permutations p of the set of players.
It is easy to check that $\bar{\alpha}$ is symmetric and that if α is a dual vector then so are all the α^p, hence so is $\bar{\alpha}$. Furthermore if $\alpha_i(d_i|c_i)$ is positive then so is $\bar{\alpha}_i(d_i|c_i)$ (since in the summation defining $\bar{\alpha}$, $\alpha^p = \alpha$ when p is the identity permutation). Thus if α is a full dual vector then $\bar{\alpha}$ is a symmetric full dual vector.

Proof of proposition 6.1, point (ii): Assume that c_i is not eliminated in full dual reductions and let α be a full dual vector. For $0 < \lambda \leq 1$, define the dual vector α^λ by: $\alpha_i^\lambda(c_i|c_i) = \lambda\alpha_i(c_i|c_i)$, $\alpha_i^\lambda(d_i|c_i) = \alpha_i(d_i|c_i) + (1 - \lambda)\alpha_i(c_i|c_i)$ and all other $\alpha_j^\lambda(d_j|c_j)$ as in α. Since α is full and α and α^λ are positive in the same components, α^λ is full too. Therefore, there exists an α_i^λ-invariant strategy σ_i^λ such that $\sigma_i^\lambda(c_i) > 0$. We claim that if $\lambda' \neq \lambda$, σ^λ is not $\alpha_i^{\lambda'}$-invariant (proof below). This implies that if $\lambda' \neq \lambda$, α^λ and $\alpha^{\lambda'}$ induce different full dual reductions. Therefore there exists an infinity of different full dual reductions. Finally, to prove the claim, note that if σ^λ is $\alpha_i^{\lambda'}$-invariant, then $\sum_{e_i \in C_i - c_i} \alpha_i^{\lambda'}(c_i|e_i)\sigma_i^\lambda(e_i) = [1 - \alpha_i^{\lambda'}(c_i|c_i)]\sigma_i^\lambda(c_i)$. But if $\lambda' \neq \lambda$:

$$\sum_{e_i \in C_i - c_i} \alpha_i^{\lambda'}(c_i|e_i)\sigma_i^\lambda(e_i) = \sum_{e_i \in C_i - c_i} \alpha_i^{\lambda}(c_i|e_i)\sigma_i^\lambda(e_i) = [1 - \alpha_i^\lambda(c_i|c_i)]\sigma_i^\lambda(c_i) \neq [1 - \alpha_i^{\lambda'}(c_i|c_i)]\sigma_i^\lambda(c_i).$$

Proof of proposition 6.7: assume that Γ is not elementary and let α be a nontrivial dual vector. Assume that some strategies of player 1 (for instance) are grouped together. That is, there exists a minimal α_1-absorbing set B_1 with at least two elements. Let B_2 and B_2' be minimal α_2-absorbing sets. Let σ_{B_1} denote the α_1-invariant strategy with support in B_1. Define σ_{B_2} and $\sigma_{B_2'}$ similarly. By proposition 6.2, σ_{B_1} is a Nash equilibrium strategy both of $\Gamma_{B_1 \times B_2}$ and of $\Gamma_{B_1 \times B_2'}$. Since Γ is $*$-generic, this implies $B_2 = B_2'$. Therefore, there is a unique minimal α_2-absorbing set, B_2. That is, C_2/α_2 is a singleton. Moreover, since Γ is locally generic, B_1 and B_2 have same cardinal. Thus B_2 has at least two elements. Therefore, by the above reasoning, the strategy set of player 1 in Γ/α is also a singleton and we are done.

Proof of lemma 7.1 : for conciseness, we refer to [8], [9] or [2] for the definitions and the notations used below. Assume that there exists a c.e.d. μ with full support. By [8, theorem 2], if μ is acceptable, then any pure strategy is acceptable, hence any pure strategy is predominant. Thus, it is enough to show that μ is acceptable. The trick is that, because μ has full support, it is possible to find trembles that will mimick μ, so that whoever the players trembling, a nontrembling player always faces the same conditional probabilities given his signal than in μ.

More precisely, assume that there exists some ϵ-correlated strategy such that:

$$\forall S \subset C, \forall e_S \in C_S, \forall c \in C, \nu^\epsilon(c, e_S) = K(S,\epsilon)\mu(c_{-S}, e_S) \tag{8}$$

where K is a positive constant that depends only on S and on ϵ (but not on c_{-S}). That is, given any coalition S of trembling players, any vector e_S of trembles assigned to S, and any strategy profile c, the probability in ν^ϵ that (c_{-S}, e_S) will be played as a result of the players being recommended c, the players of $C - S$ not trembling, and the players of S trembling to e_S, is proportional to the probability of (c_{-S}, e_S) in μ. The total probability in ν^ϵ that S (and only S) trembles and (c_{-S}, e_S) is played is then: $\sum_{d_S \in C_S} \nu^\epsilon((c_{-S}, d_S), e_S) = K'(S,\epsilon)\mu(c_{-S}, e_S)$, where K' is a positive constant which only depends on S and ϵ. It follows that, if $i \notin S$ and $c_i \in C_i$, the expected strategy of the other players in ν^ϵ, given c_i and given that S trembles, is the same that the expected strategy of the other players in μ given c_i. A fortiori, the expected strategy in ν^ϵ given (only) c_i is the same that the expected strategy in μ given c_i, to which c_i is a best response. Thus, ν^ϵ is an ϵ-equilibrium.

It remains to show that it is possible to find a sequence of ϵ-correlated strategy checking (8) and such that $\nu^\epsilon(c, \emptyset)$ tends to $\mu(c)$ as ϵ goes to zero. Such a sequence may be build by taking for all c in C and for some suitable positive normalization constant A:

$$\nu^\epsilon(c, \emptyset) = A \times \mu(c)$$

and, inductively, if the cardinal of $S \subset C$ is $m+1$:

$$\nu^\epsilon(c, e_S) = \frac{\epsilon}{1-\epsilon} A_m \times \mu(c_S, e_S)$$

with

$$A_m = \min_{d \in C} \min_{T \in S: Card T = m} \min_{e_T \in C_T} \nu^\epsilon(d, e_T)$$

A The linear programming proofs of existence of correlated equilibria

In this appendix, we review and connect the proofs of existence of correlated equilibria given in [3], [10] and [7].

A.1 Hart & Schmeidler's proof

Consider the following two-player, zero-sum, auxiliary game G_{HS}: the maximizer chooses a strategy profile $c = (c_1, .., c_n) \in C$; the minimizer chooses a player i in N and a couple of strategy (c_i', d_i) in $C_i \times C_i$. The payoff is $U_i(c) - U_i(d_i, c_i)$ if $c_i' = c_i$ and 0 otherwise. In mixed strategies the maximizer chooses a correlated strategy

$\mu \in \Delta(C)$ and the minimizer a probability distribution ν on triples $(i, c_i, d_i) \in N \times C_i \times C_i$; the expected payoff is then:

$$g_{hs}(\mu, \nu) = \sum_{c \in C} \mu(c) \sum_{i \in N} \sum_{d_i \in C_i} \nu(i, c_i, d_i)[U_i(c) - U_i(c_{-i}, d_i)] \tag{9}$$

As in the auxiliary game G of section 3, μ guarantees 0 if and only if μ is a correlated equilibrium distribution of the original game. Thus, to prove the existence of correlated equilibrium distributions, it is enough to show that the value of G_{HS} is nonnegative. To do so, Hart and Schmeidler could have used the existence of invariant distributions for finite Markov chains:[16]

Lemma A.1 *Let M be a $m \times m$ stochastic matrix (i.e. nonnegative with columns summing to unity); there exists a probability vector $x = (x_j)_{j=1,\dots,m}$ such that $Mx = x$.*

Instead, they used the following lemma:

Lemma A.2 (Hart&Schmeidler) *Let $(a_{jk})_{1 \leq j,k \leq m}$ be nonnegative numbers. There exists a probability vector $x = (x_j)_{j=1,\dots,m}$ such that, for any vector $u = (u_j)_{j=1,\dots,m}$,*

$$\sum_{j=1}^{m} x_j \sum_{k=1}^{m} a_{jk}(u_j - u_k) = 0$$

Proposition A.3 *Lemmae A.1 and A.2 are equivalent*

Proof. (i) in lemma A.2, one may assume $\sum_j a_{jk} = 1$ (indeed, one may increase arbitrarily the coefficients a_{kk} to ensure that each row sums to some positive constant and then divide all the coefficients by this constant to normalize); (ii) by linearity the property holds if and only if it holds for all basis vectors (i.e. with one component equal to 1 and all the others zero); (iii) This is equivalent to $\sum_j x_j a_{ji} = x_i$ $(= \sum_j a_{ji} x_i)$ for all i; that is, $A^T x = x$ where A^T denote the $m \times m$ square matrix whose (i, j) entry is a_{ji}. (iv) Thus lemma A.2 boils down to lemma A.1 applied to $M = A^T$. Reciprocally, lemma A.1 is a special case of lemma A.2 ■

Incidentally, Hart&Schmeidler prove their lemma using the Minimax theorem; so proposition A.3 yields a game-theoretic proof of the existence of invariant distributions for finite Markov chains.[17]

A.2 Other proofs

Nau and McCardle's proof is very similar. They also introduce (implicitly) the payoff matrix of G_{HS}. A strategy profile c is said jointly coherent if $g(c, \alpha) = 0$ for all dual vectors α. Nau and McCardle show through lemma A.1, and essentially as in section 3, that there exists a jointly coherent strategy profile. Finally, they prove through a variant of Farkas lemma that a strategy profile is jointly coherent if and only if it has positive probability in some correlated equilibrium distribution.[18] Thus correlated equilibrium distributions exists.

[16]Let λ be a positive constant. If λ is small enough, any strategy of the minimizer in G can be emulated in G_{HS}, up to the scaling factor λ, by letting: $\nu(i, c_i, d_i) = \lambda \alpha_i(d_i|c_i)/n$ if $d_i \neq c_i$. Conversely, any strategy ν of the minimizer in G_{HS} can be emulated in G by letting $\alpha_i(d_i|c_i) = \nu(i, c_i, d_i)$ if $c_i \neq d_i$; it follows that the value of G is nonnegative if and only if the value of G_{HS} is nonnegative. Thus the proof of section 3 must go through.

[17]I owe this remark to B. Von Stengel, who first showed me a proof of lemma A.1 based on linear duality. Such a proof can also be found in [6, ex. 9, p.41]

[18]In the framework of section 3, this corresponds to the following result: in a finite, two-player zero-sum game, a pure strategy is a best-response to all optimal strategies of the other player if and only if it has positive probability in some optimal strategy. This follows from the strong complementarity property of linear programs

Myerson's proofs is essentially the proof of section 3. The only difference is that instead of introducing an auxiliary zero-sum game, Myerson introduces an auxiliary linear programm and then uses linear duality. Deviation vectors appear as vectors of dual variables, hence the terms dual vector and dual reduction. Myerson's linear programm corresponds to the maximisation programm of the maximizer in the auxiliary game of section 3.

References

[1] Aumann, R. (1974), "Subjectivity and Correlation in Randomized Strategies", *Journal of Mathematical Economics* **1**, 67-96

[2] Dhillon, A. and J.F. Mertens (1996), "Perfect Correlated Equilibria", *Journal of Economic Theory* **68**, 279-302

[3] Hart, S., and D. Schmeidler (1989), "Existence of Correlated Equilibria", *Mathematics Of Operations Research*, **14**, 18-25

[4] Karr A.F. (1990), "Markov Processes" in *Handbooks in Operations Research and Management Sciences*, D.P. Heyman, M.J. Sobel Eds, Elsevier Science Publishers (North Holland), vol. 2, chap. 2, 95-123

[5] Moulin, H and J-P. Vial (1978), "Strategically Zero-Sum Games: The Class of Games Whose Completely Mixed Equilibria Cannot be Improved Upon", *International Journal of Game Theory*, vol. 7, Issue 3/4, 201-221

[6] Mertens, J.F., S. Sorin and S. Zamir (1994) "Repeated games, Part A, Background material", CORE discussion paper n9402, Université Catholique de Louvain

[7] Myerson, R.B. (1997), "Dual Reduction and Elementary Games", *Games and Economic Behavior*, **21**, 183-202

[8] Myerson, R.B. (1986), "Acceptable and Predominant Correlated Equilibria", *International Journal of Game Theory*, vol. 15, issue 3, 133-154

[9] Myerson, R.B. (1991), *Game Theory: Analysis of Conflict*, Harvard University Press, Cambridge, MA.

[10] Nau, R.F. and K.F. McCardle (1990), "Coherent Behavior in Noncooperative Games", *Journal of Economic Theory*, **50**, 424-444

[11] Rosenthal, R.W. (1974), "Correlated Equilibria in Some Classes of Two-Person Games", *International Journal of Game Theory*, vol. 3, Issue 3, 119-128

[12] Sorin, S. (1998), "Distribution Equilibrium I: Definition and Examples", cahier du THEMA n9435, Université de Cergy-Pontoise

[13] Viossat, Y. (2002), "Distribution Equilibrium II: properties", unpublished manuscript

[14] Viossat, Y. (2003), "Binding games", in preparation.

[15] von Stengel, B. (2002), "Computing Equilibria for Two-Person Games", in *Handbook of Game Theory*, R.J. Aumman and S. Hart Eds, Elsevier Science Publishers (North Holland), vol. 3, chap. 45, 1723-1759