Dual Reduction and Elementary Bayesian Games

See GEB 21:183-202 (1997): https://doi.org/10.1006/game.1997.0573
https://home.uchicago.edu/~rmyerson/research/eldual2023notes.pdf
https://home.uchicago.edu/~rmyerson/research/eldual2025bayes.pdf

Here we extend dual reduction from my 1997 GEB paper to general Bayesian games.
Bottom line:
Dual reduction generalizes dominated-strategy elimination: "imperfect" equilibria are eliminated by reducing the game, not by testing each equilibrium separately.
With dual reduction, any Bayesian game can be represented by a reduced model where all incentive constraints are satisfied strictly in almost all incentive-compatible mechanisms, so that the weak incentives of "imperfect" equilibria do not appear.
But this may involve some reinterpretation of Harsanyi's Bayesian model, where a (reduced) type may represent a player's *socially accessible* private information, and a player's action may represent a randomized strategy that can be conditioned on any parts of the player's private information that are not socially accessible.

A finite ***Bayesian game*** is any $\Gamma = (I, (T_i)_{i\in I}, (C_i)_{i\in I}, p, (u_i)_{i\in I})$ with
nonempty finite sets: $I=\{\text{players}\}$, $T_i=\{\text{i's types}\}$ $\forall i\in I$, $C_i(t_i)=\{\text{i's actions if } t_i\}$ $\forall t_i\in T_i$,
$T = \times_{i\in I} T_i$, $p\in\Delta(T)$ probability distribution of the set of type profiles, and
$u_i(c,t)\in\mathbb{R}$ utility payoffs $\forall i\in I, \forall t\in T, \forall c\in C(t) =\{c=(c_j)_{j\in I}: c_j\in C_j(t_j)\ \forall j\in I\}$.

We may write $t = (t_j)_{j\in I} = (t_{-i},t_i) \in T$, $c = (c_i)_{i\in I} = (c_{-j},c_j) \in C(t)$,
$T_{-i} = \times_{j\neq i} T_j$, $C_{-i}(t_{-i}) = \{c=(c_j)_{j\neq i}: c_j\in C_j(t_j)\ \forall j\neq i\}$.
Unlike Harsanyi (1967), we allow a player's action set to depend on the player's type.

Let $M(\Gamma)$ denote the set of direct ***coordination mechanisms*** μ that specify probabilities $\mu(c|t)\geq 0\ \ \forall c\in C(t),\ \forall t\in T,$ such that $\sum_{c\in C(t)} \mu(c|t) = 1\ \forall t\in T.$

For any player i and any type t_i in T_i, let

$U_i(\mu,t_i) = \sum_{t_{-i}\in T_{-i}} \sum_{c\in C(t)} \mu(c|t)p(t)u_i(c,t),$ [terms in prior EU_i from when i's type is t_i]

$\hat{U}_i(\mu,c_i,d_i,s_i,t_i) = \sum_{t_{-i}\in T_{-i}} \sum_{c_{-i}\in C_{-i}(t_{-i})} \mu(c_{-i},d_i|t_{-i},s_i)p(t)u_i(c,t)\ \ \forall c_i\in C_i(t_i),\ \forall d_i\in C_i(s_i),\ \forall s_i\in T_i.$

[...altered terms where t_i reports s_i and then does c_i when d_i is recommended]

Any μ in $M(\Gamma)$ is ***incentive compatible*** (***IC***) iff it satisfies the ***incentive constraints***:

$U_i(\mu,t_i) \geq \sum_{d_i\in C_i(s_i)} \hat{U}_i(\mu,\delta_i(d_i),d_i,s_i,t_i)\ \ \forall \delta_i:C_i(s_i)\to C_i(t_i),\ \forall s_i\in T_i,\ \forall t_i\in T_i,\ \forall i\in I.$

Any equilibrium of any communication system is equivalent to an IC mechanism. [*Rev P*]

The above incentive constraints are ***trivial*** when $s_i=t_i$ & $\delta_i(d_i)=d_i\ \forall d_i$, because any mechanism μ always satisfies $U_i(\mu,t_i) = \sum_{d_i\in C_i(t_i)} \hat{U}_i(\mu,d_i,d_i,t_i,t_i)\ \forall t_i\ \forall i.$

A Bayesian game Γ is an ***elementary*** game iff there exists some μ^0 in M that satisfies all nontrivial incentive constraints as strict inequalities (>). So $\forall i,\ \forall t_i,\ \forall s_i,\ \forall \delta_i,\ \forall c_i,\ \forall e_i$:

$U_i(\mu^0,t_i) > \sum_{d_i\in C_i(s_i)} \hat{U}_i(\mu^0,\delta_i(d_i),d_i,s_i,t_i)$ if $s_i\neq t_i$, & $\hat{U}_i(\mu^0,c_i,c_i,t_i,t_i) > \hat{U}_i(\mu^0,e_i,c_i,t_i,t_i)$ if $e_i\neq c_i$.

Fact 1: If Γ is elementary, then almost all IC mechanisms satisfy all nontrivial incentive constraints strictly. (*If μ does not then $(1-\varepsilon)\mu+\varepsilon\mu^0$ does.*)

(Then problems of weak incentives in "imperfect" equilibria can be avoided.)

A mechanism μ in $M(\Gamma)$ is incentive compatible if and only if, for some vector ψ,

$U_i(\mu,t_i) \geq \sum_{d_i\in C_i(s_i)} \psi_i(d_i,s_i|t_i)\ \ \forall s_i\in T_i,\ \forall t_i\in T_i,\ \forall i\in I;$ and

$\psi_i(d_i,s_i|t_i) \geq \hat{U}_i(\mu,c_i,d_i,s_i,t_i)\ \ \forall c_i\in C_i(t_i),\ \forall d_i\in C_i(s_i),\ \forall s_i\in T_i,\ \forall t_i\in T_i,\ \forall i\in I.$

If μ is IC, these constraints can be satisfied with $\psi_i(d_i,s_i|t_i) = \max_{e_i\in C_i(t_i)} \hat{U}_i(\mu,e_i,d_i,s_i,t_i).$

Consider this ***primal linear programming problem***: (with artificial variables π)
minimize $\sum_i \sum_{ti} \pi_i(t_i)$ over $\mu \geq \mathbf{0}$ & π & ψ such that

$\pi_i(t_i) + \sum_{t-i} \sum_c \mu(c|t)p(t)u_i(c,t) - \sum_{di \in Ci(si)} \psi_i(d_i,s_i|t_i) \geq 0 \ \ \forall s_i \in T_i, \forall t_i \in T_i, \forall i \in I;$ [$\alpha_i(s_i|t_i)$]

$\psi_i(d_i,s_i|t_i) - \sum_{t-i} \sum_{c-i} \mu(c_{-i},d_i|t_{-i},s_i)p(t)u_i(c,t) \geq 0 \ \ \forall c_i \in C_i(t_i), \forall d_i \in C_i(s_i), \forall s_i \ \forall t_i \ \forall i;$ [$\alpha_i(c_i|d_i,s_i|t_i)$]

$\sum_{c \in C(t)} \mu(c|t) = 1 \ \ \forall t \in T.$ [$\beta(t)$]

The solutions to this LP are the incentive-compatible mechanisms μ, which exist (Nash) and yield optimal value 0, with $\pi = \mathbf{0}$ and $\psi_i(d_i,s_i|t_i) = \max_{ci \in Ci(ti)} \hat{U}_i(\mu,c_i,d_i,s_i,t_i)$.

(Trivial constraints with $s_i = t_i$ & $d_i = c_i$ imply $\pi_i(t_i) \geq \sum_{di \in Ci(ti)} \hat{U}_i(\mu,d_i,d_i,t_i,t_i) - U_i(\mu,t_i) = 0$.)

The ***dual LP problem*** is: maximize $\sum_t \beta(t)$ over $\alpha \geq \mathbf{0}$ & β such that

$\beta(t) + \sum_{i \in I} \sum_{si \in Ti} \alpha_i(s_i|t_i)p(t)u_i(c,t) +$
$\quad - \sum_{i \in I} \sum_{si \in Ti} \sum_{di \in C(si)} \alpha_i(d_i|c_i,t_i|s_i)p(t_{-i},s_i)u_i((c_{-i},d_i),(t_{-i},s_i)) \leq 0 \ \ \forall c \in C(t), \forall t \in T;$ [$\mu(c|t)$]

$\sum_{ci \in Ci(ti)} \alpha_i(c_i|d_i,s_i|t_i) - \alpha_i(s_i|t_i) = 0 \ \ \forall d_i \in C_i(s_i), \forall s_i \in T_i, \forall t_i \in T_i, \forall i \in I;$ [$\psi_i(d_i,s_i|t_i)$]

$\sum_{si \in Ti} \alpha_i(s_i|t_i) = 1 \ \ \forall t_i \in T_i, \ \forall i \in I.$ [$\pi_i(t_i)$]

Let $\bar{A}(\Gamma)$ denote the set of all vectors α that are optimal dual solutions, with some β.

We may define the ***aggregate α-deviation value*** at (c,t) to be

$D(c,t,\alpha) = \sum_{i \in I} \sum_{si \in Ti} \sum_{di \in C(si)} \alpha_i(d_i|c_i,t_i|s_i)p(t_{-i},s_i)u_i((c_{-i},d_i),(t_{-i},s_i)) - \sum_{i \in I} p(t)u_i(c,t).$

Then the first dual constraint is equivalent to: $\beta(t) \leq D(c,t,\alpha) \ \ \forall c \in C(t), \forall t \in T.$

Any optimal dual solution has $\beta(t) = \min_{c \in C(t)} D(c,t,\alpha) \ \forall t \in T$, and $\sum_{t \in T} \beta(t) = 0$.

Fact 2: Given any dual solution α in $\bar{A}(\Gamma)$, for any mechanism μ in $M(\Gamma)$:

$$\sum_{i\in I} \sum_{ti\in Ti} \sum_{si\in Ti} \sum_{di\in Ci(si)} \sum_{ci\in Ci(ti)} \alpha_i(c_i|d_i,s_i|t_i)\mathbf{(}\hat{U}_i(\mu,c_i,d_i,s_i,t_i) - \max_{ei\in Ci(ti)} \hat{U}_i(\mu,e_i,d_i,s_i,t_i)\mathbf{)}$$

$$+ \sum_{i\in I} \sum_{ti\in Ti} \sum_{si\in Ti} \alpha_i(s_i|t_i)\mathbf{(}\sum_{di\in Ci(si)} \max_{ei\in Ci(ti)} \hat{U}_i(\mu,e_i,d_i,s_i,t_i) - U_i(\mu,t_i)\mathbf{)}$$

$$= \sum_{i\in I} \sum_{ti\in Ti} \mathbf{(}\sum_{si\in Ti} \sum_{di\in Ci(si)} \sum_{ci\in Ci(ti)} \alpha_i(c_i|d_i,s_i|t_i)\hat{U}_i(\mu,c_i,d_i,s_i,t_i) - U_i(\mu,t_i)\mathbf{)}$$

$$= \sum_{t\in T} \sum_{c\in C(t)} \mu(c|t)D(c,t,\alpha) \geq \sum_{t\in T} \beta(t) = 0.$$

The inequality $\geq$ above must become equality $=$ when μ is incentive compatible.
(*Expected net gains of unilateral α-deviations from μ must have a nonnegative sum.*)

Fact 3: A dual solution α with $\alpha_i(c_i|d_i,s_i|t_i)>0$ exists if & only if there exists some $\delta_i:C_i(s_i)\rightarrow C_i(t_i)$ such that $\delta_i(d_i)=c_i$ and every incentive-compatible μ satisfies:

$$U_i(\mu,t_i) = \sum_{di\in Ci(si)} \hat{U}_i(\mu,\delta_i(d_i),d_i,s_i,t_i)$$

$$= \hat{U}_i(\mu,c_i,d_i,s_i,t_i) + \sum_{bi\in Ci(si),bi\neq di} \max_{ai\in Ci(ti)} \hat{U}_i(\mu,a_i,b_i,s_i,t_i).$$

(*Follows from strict complementary of linear programming solutions.*)

A ***trivial*** dual solution has, $\forall i\ \forall t_i\ \forall c_i$, $\alpha_i(c_i|c_i,t_i|t_i)=1$ and all other $\alpha_i(c_i|d_i,s_i|t_i)=0$.
Fact 4: Nontrivial dual solutions in $\bar{A}(\Gamma)$ exist if and only if Γ is not elementary.

Given any dual solution $\alpha\in\bar{A}(\Gamma)$ and any $i\in I$, let $\mu*\alpha_i$ denote the mechanism such that,

$$(\mu*\alpha_i)(c|t) = \sum_{si\in Ti} \sum_{di\in Ci(si)} \alpha_i(c_i|d_i,s_i|t_i)\, \mu(c_{-i},d_i|t_{-i},s_i) \quad \forall c\in C(t),\ \forall t\in T.$$

Then $\sum_{si\in Ti} \sum_{di\in Ci(si)} \sum_{ci\in Ci(ti)} \alpha_i(c_i|d_i,s_i|t_i)\hat{U}_i(\mu,c_i,d_i,s_i,t_i) = U_i(\mu*\alpha_i,t_i)$.
So by Fact 2, for any mechanism μ in $M(\Gamma)$, $\sum_{j\in I} \sum_{tj\in Tj} (U_j(\mu*\alpha_j,t_j) - U_j(\mu,t_j)) \geq 0$.
But if μ is an incentive-compatible mechanism μ then $U_i(\mu*\alpha_i,t_i) = U_i(\mu,t_i)\ \ \forall t_i\in T_i$.
In this sense, α_i defines a strategy for (mis)reporting i's type and then reacting to any recommended action that player i would be willing to apply with any IC mechanism.

In this α_i-manipulative strategy, $\alpha_i(c_i|d_i,s_i|t_i)$ denotes the probability that i's type t_i would imitate s_i with a plan to do c_i if d_i is subsequently recommended, and the overall probability of t_i imitating s_i would be $\alpha_i(s_i|t_i) = \sum_{c_i \in C_i(t_i)} \alpha_i(c_i|d_i,s_i|t_i) \ \ \forall d_i \in C_i(s_i)$.

In any incentive-compatible mechanism, each $t_i \in T_i$ is willing to imitate other types s_i with probabilities $\alpha_i(s_i|t_i)$; but these imitated types might similarly α_i-imitate others...

Consider the Markov chain on T_i with transition probabilities $\alpha_i(s_i|t_i)$ from $t_i \in T_i$ to $s_i \in T_i$, and let T_i/α_i denote the set of minimal nonempty absorbing sets for this Markov chain.

These minimal absorbing sets will be i's "reduced-types" in our α-reduced game.

Then for any t_i in T_i and any τ_i in T_i/α_i, let $\omega_i(\tau_i|t_i,\alpha_i)$ denote the probability that the Markov chain ends in the absorbing set τ_i from an initial condition t_i.

These ω_i probabilities are the unique solution to the equations:

$\omega_i(\tau_i|t_i,\alpha_i) = \sum_{s_i \in T_i} \alpha_i(s_i|t_i)\omega_i(\tau_i|s_i,\alpha_i) \ \forall t_i \in T_i$;

if $t_i \in \tau_i$ then $\omega_i(\tau_i|t_i,\alpha_i)=1 \ \ \forall t_i$; and

if $\rho_i \in T_i/\alpha_i$ but $\rho_i \neq \tau_i$ and $s_i \in \rho_i$ then $\omega_i(\tau_i|s_i,\alpha_i)=0 \ \ \forall s_i, \forall \rho_i$.

So a plan for each type t_i to report reduced-types τ_i with probabilities $\omega_i(\tau_i|t_i,\alpha_i)$ would be invariant under α_i-imitations.

We will interpret $\omega_i(\tau_i|t_i,\alpha_i)$ as the probability of type t_i reporting reduced-type τ_i in the α-reduced game.

Given any reduced-type $\tau_i \in T_i/\alpha_i$, let $\Omega(\tau_i) = \{t_i \in T_i | \omega_i(\tau_i|t_i,\alpha_i)>0\}$.

Let $[\alpha_i|\tau_i]$ denote the α_i probabilities conditioned on the Markov chain converging to τ_i:

$[\alpha_i|\tau_i](c_i|d_i,s_i|t_i) = \alpha_i(c_i|d_i,s_i|t_i)\omega_i(\tau_i|s_i,\alpha_i)/\omega_i(\tau_i|t_i,\alpha_i) \;\; \forall t_i \in \Omega(\tau_i)$.

When α_i imitations imply that a reduced-type report τ_i could come from any t_i in $\Omega(\tau_i)$, a mediator can only recommend a behavioral strategy σ_i in $\times_{ti\in\Omega(\tau i)} \Delta(C_i(t_i))$.

A recommendation is invariant under the α_i manipulation leading to τ_i only if it satisfies $[\alpha_i|\tau_i]$-stationarity: $\sigma_i(c_i|t_i) = \sum_{si\in\Omega(\tau i)} \sum_{di\in Ci(si)} \sigma_i(d_i|s_i)[\alpha_i|\tau_i](c_i|d_i,s_i|t_i) \;\; \forall c_i, \forall t_i \in \Omega(\tau_i)$.

The support of such an invariant behavioral strategy must be absorbing for $[\alpha_i|\tau_i]$, in the sense that: if $\sigma_i(d_i|s_i)>0$ and $[\alpha_i|\tau_i](c_i|d_i,s_i|t_i)>0$ then $\sigma_i(c_i|t_i)>0 \;\; \forall t_i \in \Omega(\tau_i), \forall s_i \in \Omega(\tau_i)$.

So we say that a function γ_i is an ***action subspace*** on $\Omega(\tau_i)$ iff $\emptyset \neq \gamma_i(t_i) \subseteq C_i(t_i) \;\; \forall t_i \in \Omega(\tau_i)$.

We say that an action subspace γ_i is ***absorbing*** for $[\alpha_i|\tau_i]$ iff,

if $d_i \in \gamma_i(s_i)$ and $[\alpha_i|\tau_i](c_i|d_i,s_i|t_i)>0$ then $c_i \in \gamma_i(t_i) \; \forall t_i \in \Omega(\tau_i), \forall s_i \in \Omega(\tau_i)$.

An intersection of absorbing subspaces is absorbing, $\neq\emptyset$ on all $\Omega(\tau_i)$ if $\neq\emptyset$ on any in τ_i.

So the minimal absorbing action subspaces select disjoint sets of actions for each t_i in τ_i.

Let $C_i/[\alpha_i|\tau_i]$ denote the set of minimal absorbing action subspaces for $[\alpha_i|\tau_i]$.

$C_i/[\alpha_i|\tau_i]$ will be the set of reduced-actions for reduced-type τ_i in our α-reduced game.

Given any $\gamma_i \in C_i/[\alpha_i|\tau_i]$, let $\{\sigma_i(\bullet|\gamma_i,\tau_i,t_i,\alpha_i)\}_{ti\in\Omega(\tau i)}$ be the probability distributions on the sets $\{\gamma_i(t_i)\}_{ti\in\Omega(\tau i)}$ that uniquely satisfy the $[\alpha_i|\tau_i]$-stationarity equations:

$$\sigma_i(c_i|\gamma_i,\tau_i,t_i,\alpha_i) = \sum_{si\in\Omega(\tau i)} \sum_{di\in\gamma i(si)} \sigma_i(d_i|\gamma_i,\tau_i,s_i,\alpha_i)\alpha_i(c_i|d_i,s_i|t_i)\omega_i(\tau_i|s_i,\alpha_i)/\omega_i(\tau_i|t_i,\alpha_i)$$

$$\forall c_i \in \gamma_i(t_i), \forall t_i \in \Omega(\tau_i);$$

$$\sum_{ci\in\gamma i(ti)} \sigma_i(c_i|\gamma_i,\tau_i,t_i,\alpha_i) = 1 \;\; \forall t_i \in \Omega(\tau_i).$$

Interpret $\sigma_i(c_i|\gamma_i,\tau_i,t_i,\alpha_i)$ as probability of action c_i for type t_i in reduced-action γ_i for τ_i.

Given any dual solution α for the game Γ, let us define the ***α-reduced game*** Γ/α such that the set of players is still I, the set of reduced-types for each player i is T_i/α_i, the set of reduced-actions for any reduced-type τ_i in T_i/α_i is $C_i/[\alpha_i|\tau_i]$, the probability distribution over reduced-types is q such that

$$q(\tau) = \sum_{t\in T} p(t)(\prod_{j\in I} \omega_j(\tau_j|t_j,\alpha_j)) \quad \forall\tau\in\times_{j\in I} T_j/\alpha_j, \text{ and}$$

the utility function for each player i is v_i such that, $\forall\tau\in\times_{j\in I} T_j/\alpha_j$, $\forall\gamma\in\times_{j\in I} C_j/[\alpha_j|\tau_j]$:

$$v_i(\gamma,\tau)\, q(\tau) = \sum_{t\in T} \sum_{c\in C(t)} p(t)(\prod_{j\in I} \omega_j(\tau_j|t_j,\alpha_j))(\prod_{j\in I} \sigma_j(c_j|\gamma_j,\tau_j,t_j,\alpha_j))u_i(c,t).$$

Any coordination mechanism λ for the reduced game Γ/α induces a mechanism μ^λ for the original game such that

$$\mu^\lambda(c|t) = \sum_\tau \sum_\gamma \lambda(\gamma|\tau)(\prod_{j\in I} \omega_j(\tau_j|t_j,\alpha_j)\, \sigma_j(c_j|\gamma_j,\tau_j,t_j,\alpha_j)) \quad \forall c\in C(t), \forall t\in T.$$

(Here τ is summed over the set of reduced-type profiles $\times_{i\in I} T_i/C_i$, and then γ is summed over the set of reduced-action profiles $\times_{i\in I} C_i/[\alpha_i|\tau_i]$ for the reduced-types in τ.)

Theorem 1: If α is a dual solution for Γ, then any incentive-compatible mechanism λ for the α-reduced game Γ/α induces an incentive-compatible mechanism μ^λ for Γ.

(*By Fact 2, a player in Γ expects no loss from α-deviations when others act as in Γ/α...*)

Theorem 2: For any finite Γ, iterative dual reduction yields an elementary reduced game. (*Proof: Use Fact 4...*)

We let $\mu * \alpha_i$ denote the mechanism such that, $\forall t \in T$, $\forall c \in C(t)$,

$$(\mu * \alpha_i)(c|t) = \sum_{si \in Ti} \sum_{di \in Ci(si)} \alpha_i(c_i|d_i,s_i|t_i)\mu(c_{-i},d_i|t_{-i},s_i).$$

Then we get $\mu^\lambda * \alpha_i = \mu^\lambda$ for all $i \in I$. The proof is as follows:

$$(\mu^\lambda * \alpha_i)(c|t) = \sum_{si \in Ti} \sum_{di \in Ci(si)} \alpha_i(c_i|d_i,s_i|t_i)\mu^\lambda(c_{-i},d_i|t_{-i},s_i)$$
$$= \sum_{si \in Ti} \sum_{di \in Ci(si)} \alpha_i(c_i|d_i,s_i|t_i) \sum_\tau \sum_\gamma \lambda(\gamma|\tau)\sigma_i(d_i|\gamma_i,\tau_i,s_i,\alpha_i)\omega_i(\tau_i|s_i,\alpha_i) \times$$
$$\times (\textstyle\prod_{j \neq i} \omega_j(\tau_j|t_j,\alpha_i)\sigma_j(c_j|\gamma_j,\tau_j,t_j,\alpha_i))$$
$$= \sum_\tau \sum_\gamma \lambda(\gamma|\tau) \sum_{si \in \Omega(\tau i)} \sum_{di \in Ci(si)} \sigma_i(d_i|\gamma_i,\tau_i,s_i,\alpha_i)\alpha_i(c_i|d_i,s_i|t_i)\omega_i(\tau_i|s_i,\alpha_i) \times$$
$$\times (\textstyle\prod_{j \neq i} \omega_j(\tau_j|t_j,\alpha_i)\sigma_j(c_j|\gamma_j,\tau_j,t_j,\alpha_i))$$
$$= \sum_\tau \sum_\gamma \lambda(\gamma|\tau)\sigma_i(c_i|\gamma_i,\tau_i,t_i,\alpha_i)\omega_i(\tau_i|t_i,\alpha_i)(\textstyle\prod_{j \neq i} \omega_j(\tau_j|t_j,\alpha_i)\sigma_j(c_j|\gamma_j,\tau_j,t_j,\alpha_i))$$
$$= \mu^\lambda(c|t).$$

More generally, a mechanism μ for Γ is ***strategically measurable for player i in the reduced game*** Γ/α iff there is some $\eta(\bullet|\bullet)$ such that μ can be written

$$\mu(c|t) = \sum_{\tau i \in Ti/\alpha i} \sum_{\gamma i \in Ci/[\alpha i|\tau i]} \eta(c_{-i},\gamma_i|t_{-i},\tau_i)\, \sigma_i(c_i|\gamma_i,\tau_i,t_i,\alpha_i)\, \omega_i(\tau_i|t_i,\alpha_i) \quad \forall c \in C(t), \forall t \in T.$$

Fact 5: Given $\alpha \in \bar{A}(\Gamma)$, if $\mu \in M(\Gamma)$ is strategically measurable for i in Γ/α then $\mu * \alpha_i = \mu$.

Proof: $(\mu * \alpha_i)(c|t) = \sum_{si \in Ti} \sum_{di \in Ci(si)} \alpha_i(c_i|d_i,s_i|t_i)\mu(c_{-i},d_i|t_{-i},s_i)$

$$= \sum_{si \in Ti} \sum_{di \in Ci(si)} \sum_{\tau i} \sum_{\gamma i} \alpha_i(c_i|d_i,s_i|t_i)\eta(c_{-i},\gamma_i|t_{-i},\tau_i)\sigma_i(d_i|\gamma_i,\tau_i,s_i,\alpha_i)\omega_i(\tau_i|s_i,\alpha_i)$$
$$= \sum_{\tau i} \sum_{\gamma i} \eta(c_{-i},\gamma_i|t_{-i},\tau_i) \sum_{si \in \Omega(\tau i)} \sum_{di \in Ci(si)} \alpha_i(c_i|d_i,s_i|t_i)\sigma_i(d_i|\gamma_i,\tau_i,s_i,\alpha_i)\omega_i(\tau_i|s_i,\alpha_i)$$
$$= \sum_{\tau i} \sum_{\gamma i} \eta(c_{-i},\gamma_i|t_{-i},\tau_i)\, \sigma_i(c_i|\gamma_i,\tau_i,t_i,\alpha_i)\omega_i(\tau_i|t_i,\alpha_i) = \mu(c|t).$$

Proof sketch for Theorem 1:

We are given λ as an incentive-compatible mechanism for the α-reduced game Γ/α.

If μ^λ were not incentive-compatible for Γ, then there would be some player i who could gain by manipulating reports or reactions in some way.

Let $\tilde{\mu}$ be a mechanism for Γ that is induced when everyone other than i is expected to behave according to λ in the reduced game but player i uses a Γ-optimal strategy for reporting and reacting to the λ-mediator.

By the Fact 2, $\sum_{j\in I}\sum_{t_j\in T_j}(U_j(\mu*\alpha_j,t_j) - U_j(\mu,t_j)) \geq 0$ for any mechanism μ.

But for every player j who acts according to the reduced game, we have $\tilde{\mu}*\alpha_j = \tilde{\mu}$.

Thus, for the deviating player i alone, we have $\sum_{t_i\in T_i}(U_i(\tilde{\mu}*\alpha_i,t_i) - U_i(\tilde{\mu},t_i)) \geq 0$.

So $\tilde{\mu}*\alpha_i$ is also a mechanism in which player i uses a Γ-optimal strategy for reporting and reacting, while everyone else still behaves according to λ in the reduced game.

That is, the Γ-optimality of i's behavior is preserved by the α_i transformation, as long as everyone else behaves according to the strategic restrictions of the reduced game Γ/α.

But iterative transformation by α_i would make i's reporting & reacting strategies approach behavior that is feasible in the α-reduced game, where honesty and obedience to λ are optimal for j (as λ is IC for all players in the reduced game).

For more proof details, see also:

https://home.uchicago.edu/~rmyerson/research/eldual2025more.pdf

Density of games with transitive dual solutions.

For any game Γ, we say that a dual solution α is *maximal* iff its set of strictly positive nontrivial components $\alpha_i(c_i|d_i,s_i|t_i)>0$ is maximal among all dual solutions in $\bar{A}(\Gamma)$.

(The *trivial components* of α are the components $\alpha_i(c_i|c_i,t_i|t_i)$ for any $i\in I$, $t_i\in T_i$, $c_i\in C_i(t_i)$.)

We may define the *dual magnitude* of a game to be the number of nonzero nontrivial components in a maximal dual solution for the game. (So it is 0 if Γ is elementary.)

We say that the game Γ has ***transitive dual solutions*** iff, in a maximal dual solution α,
$\forall i\in I$, $\forall \{t_i,s_i,r_i\}\subseteq T_i$, $\forall c_i\in C_i(t_i)$, $\forall d_i\in C_i(s_i)$, $\forall e_i\in C_i(r_i)$:
if $\alpha_i(c_i|d_i,s_i|t_i)>0$ & $\alpha_i(d_i|e_i,r_i|s_i)>0$ then $\alpha_i(c_i|e_i,r_i|t_i)>0$.

Any $(I, (T_i)_{i\in I}, (C_i)_{i\in I}, p)$ as above may be called a *framework* for Bayesian games, which can then be defined by specifying a vector of utilities for all players in all outcomes.

Given this framework, the game defined by a utilities vector u may be denoted $\Gamma(u)$.

We say that the type distribution p has *full support* iff $p(t) > 0$ for all $t\in T$.

Theorem 3: Given any framework $(I, (T_i)_{i\in I}, (C_i)_{i\in I}, p)$ for Bayesian games, suppose that the type distribution p has full support. Then the set of utility functions u such that $\Gamma(u)$ has transitive dual solutions is dense in the utility space.

Proof: Given any nonempty open set in the utility space, we can select u in this set so that $\Gamma(u)$ has a dual magnitude (a nonnegative integer) that is minimal in the set.

Then we can show that this $\Gamma(u)$ with locally minimal dual magnitude will have transitive dual solutions...

Let α be any maximal dual solution for $\Gamma(u)$, which has locally minimal dual magnitude.

$\Gamma(u)$ also has an IC mechanism μ that is *maximally* incentive compatible, satisfying with strict inequalities all incentive constraints that are strict in any IC mechanism for $\Gamma(u)$, including all incentive constraints from Fact 3 for the zero components in α.

If $\Gamma(u)$ did not have transitive dual solutions, then we could find some $i\in I$ and nontrivial strictly positive components $\alpha_i(c_i|d_i,s_i|t_i)>0$ and $\alpha_i(d_i|e_i,r_i|s_i)>0$ such that $\alpha_i(c_i|e_i,r_i|t_i)=0$.

(If either of these positive components was trivial, they could not violate transitivity.)

With Fact 3, the two positive dual variables would give us:

$U_i(\mu,t_i) = \hat{U}_i(\mu,c_i,d_i,s_i,t_i) + \sum_{b_i\in C_i(s_i),b_i\neq d_i} \max_{a_i\in C_i(t_i)} \hat{U}_i(\mu,a_i,b_i,s_i,t_i)$,

$U_i(\mu,s_i) = \hat{U}_i(\mu,d_i,e_i,r_i,s_i) + \sum_{b_i\in C_i(r_i),b_i\neq e_i} \max_{a_i\in C_i(s_i)} \hat{U}_i(\mu,a_i,b_i,r_i,s_i)$.

But with $\alpha_i(c_i|e_i,r_i|t_i)=0$ and μ maximally incentive compatible, we would get

$U_i(\mu,t_i) > \hat{U}_i(\mu,c_i,e_i,r_i,t_i) + \sum_{b_i\in C_i(r_i),b_i\neq e_i} \max_{a_i\in C_i(t_i)} \hat{U}_i(\mu,a_i,b_i,r_i,t_i)$.

Let $\delta_i:C_i(s_i)\rightarrow C_i(t_i)$ be such that $\delta_i(d_i)=c_i$ and $\delta_i(b_i)\in\text{argmax}_{a_i\in C_i(t_i)} \hat{U}_i(\mu,a_i,b_i,s_i,t_i)\ \forall b_i\neq d_i$.

For any small $\varepsilon>0$, let us construct $\tilde{u}$ to be the same as u except that, $\forall t_{-i}$, $\forall b\in C(t_{-i},s_i)$:

$p(t_{-i},s_i)\tilde{u}_i(b,(t_{-i},s_i)) = (1-\varepsilon)p(t_{-i},s_i)u_i(b,(t_{-i},s_i)) + \varepsilon\, p(t_{-i},t_i)u_i((b_{-i},\delta_i(b_i))),(t_{-i},t_i))$.

Then any μ that is maximally incentive compatible for $\Gamma(u)$ would also be incentive compatible for $\Gamma(\tilde{u})$, but its set of strictly satisfied incentive constraints would expand in $\Gamma(\tilde{u})$ to include all constraints that could justify positivity of $\alpha_i(d_i|e_i,r_i|s_i)$ in Fact 3.

(In the ε term: t_i would strictly lose by reporting r_i & planning reaction $c_i=\delta_i(d_i)$ to e_i.)

So this construction would yield utility vectors $\tilde{u}$ that are in the open set but give $\Gamma(\tilde{u})$ a strictly smaller dual magnitude than $\Gamma(u)$, contradicting the selection of u.

Corollary: We can show that such dual transitivity implies also that:
if $\alpha_i(r_i|s_i)>0$ and $\alpha_i(s_i|t_i)>0$ then $\alpha_i(r_i|t_i)>0$.

Proof: By the "if" part, for any e_i in $C_i(r_i)$, we can find $d_i \in C_i(s_i)$ and $c_i \in C_i(t_i)$ such that $\alpha_i(d_i|e_i,r_i|s_i)>0$ and $\alpha_i(c_i|d_i,s_i|t_i)>0$.

Then dual transitivity implies $\alpha_i(c_i|e_i,r_i|t_i)>0$, and so
$\alpha_i(r_i|t_i) = \sum_{b_i \in C_i(t_i)} \alpha_i(b_i|e_i,r_i|t_i) > 0$.

Concluding note:

Dual reduction identifies incentive constraints that are hard to satisfy with strict perfection, and it models them as inseparable alternatives in a reduced game.

A reduced-type may represent a pooling of inseparable types, omitting aspects of the player's private information that are socially inaccessible.

A reduced-action may randomize over inseparable actions in way that is strategically conditioned on this socially-inaccessible information.

Iterative dual reduction of all such inseparable actions and inseparable types yields an elementary reduced game where all incentive constraints can be satisfied strictly.

Thus, dual reduction allows us to analyze games without any knife-edge imperfection issues, because any such issues in the original game have been identified and embedded into the structure of the reduced game.

https://home.uchicago.edu/~rmyerson/research/eldual2025bayes.pdf

Examples with no private information, strategic-form games (as in the 1997 paper):

C_1: \ C_2:	c_2	d_2
c_1	3, 2	0, 0
d_1	0, 0	2, 3

All incentive constraints can be satisfied strictly with $\mu(c_1,c_2)=0.5=\mu(d_1,d_2)$.
So this game is elementary, and it has no nontrivial dual solutions.

C_1: \ C_2:	c_2	d_2
c_1	5, 5	0, 5
d_1	5, 0	1, 1

Dual solutions include $\alpha_1(d_1|c_1)=1$, $\alpha_1(c_1|d_1)=0$, $\alpha_2(d_2|c_2)=1$, $\alpha_2(c_2|d_2)=0$.
$C_i/\alpha_i = \{\{d_i\}\}$. In the reduced game, the dominated actions c_1 and c_2 are eliminated.

C_1: \ C_2:	c_2	d_2
c_1	7, 0	2, 5
d_1	4, 3	6, 1

Dual solutions include $\alpha_1(d_1|c_1)=1$, $\alpha_1(c_1|d_1)=0.4$, $\alpha_2(d_2|c_2)=0.6$, $\alpha_2(c_2|d_2)=0.8$, and the reduced game has one absorbing set of actions $\{c_i,d_i\}$ for each player i.
The α-stationary strategies are the unique Nash equilibrium strategies: $(2/7)[c_1]+(5/7)[d_1]$, $(4/7)[c_2]+(3/7)[d_2]$.
The reduced game is 1×1 with the equilibrium payoffs (4.857, 2.143).

Example of a 3×3 strategic-form game (rock, scissors, paper):

C_1: \ C_2:	c_2	d_2	e_2
c_1	0, 0	1, 0	0, 1
d_1	0, 1	0, 0	1, 0
e_1	1, 0	0, 1	0, 0

This game has a correlated equilibrium that randomizes uniformly over the six nondiagonal outcomes of the game, which satisfies strictly six incentive constraints and so implies that every dual solution must have:

$\alpha_1(e_1|c_1)=0$, $\alpha_1(c_1|d_1)=0$, $\alpha_1(d_1|e_1)=0$, $\alpha_2(e_2|c_2)=0$, $\alpha_2(c_2|d_2)=0$, $\alpha_2(d_2|e_2)=0$.

But this game is not elementary. There is a dual solution with:

$\alpha_1(d_1|c_1)=1$, $\alpha_1(e_1|d_1)=1$, $\alpha_1(c_1|e_1)=1$, $\alpha_2(d_2|c_2)=1$, $\alpha_1(e_2|d_2)=1$, $\alpha_2(c_2|e_2)=1$.

The dual reduction is the 1×1 game where each player j's only option is to randomize uniformly over $\{c_j,d_j,e_j\}$, yielding expected payoffs (1/3, 1/3).

Thus, dual reduction suggests that the correlated equilibrium that we described above may be imperfect in some sense.

Sensitivity analysis: Suppose that we change the diagonal payoffs from (0,0) to $(\varepsilon,\varepsilon)$ for some number ε.

If $\varepsilon < 0$, the game is elementary, as the above correlated equilibrium satisfies all incentive constraints strictly.

If $\varepsilon > 1$, the game is elementary, as $\mu(c_1,c_2)=\mu(d_1,d_2)=\mu(e_1,e_2)=1/3$ satisfies all incentive constraints strictly.

If $0 < \varepsilon < 1$, no incentive constraint can be satisfied strictly, and a dual solution where all nontrivial dual variables are strictly positive is: $\alpha_1(e_1|c_1)=\alpha_1(c_1|d_1)=\alpha_1(d_1|e_1)=\alpha_2(e_2|c_2)=\alpha_2(c_2|d_2)=\alpha_2(d_2|e_2)=\varepsilon$, $\alpha_1(d_1|c_1)=\alpha_1(e_1|d_1)=\alpha_1(c_1|e_1)=\alpha_2(d_2|c_2)=\alpha_1(e_2|d_2)=\alpha_2(c_2|e_2)=1-\varepsilon$.

These games have ***nontransitive dual solutions*** only in two cases: when $\varepsilon=0$, and when $\varepsilon=1$.

A similar 3×3 sender-receiver game (one sender with types, one receiver with actions):

p:	T_1: \ C_2:	a_2	b_2	c_2
1/3	r_1	0, 0	1, 0	0, 1
1/3	s_1	0, 1	0, 0	1, 0
1/3	t_1	1, 0	0, 1	0, 0

An IC μ has $\mu(b_2|r_1)=\mu(c_2|r_1)=0.5$, $\mu(a_2|s_1)=\mu(c_2|s_1)=0.5$, $\mu(a_2|t_1)=\mu(b_2|t_1)=0.5$.
Its six strictly satisfied constraints imply that any dual solution has
$\alpha_1(s_1|r_1)=\alpha_1(t_1|s_1)=\alpha_1(r_1|t_1)=0$ and $\alpha_2(c_2|a_2)=\alpha_2(a_2|b_2)=\alpha_2(b_2|c_2)=0$.

But this game is not elementary. A dual solution has
$\alpha_1(t_1|r_1)=\alpha_1(r_1|s_1)=\alpha_1(s_1|t_1)=1$ and $\alpha_2(b_2|a_2)=\alpha_2(c_2|b_2)=\alpha_2(a_2|c_2)=1$.
In the reduced game, the sender has one reduced-type (equally likely to be r_1, s_1, or t_1), and the receiver has one reduced-action (randomizing uniformly over $\{a_2,b_2,c_2\}$).
Thus, dual reduction suggests that the incentive-compatible mechanism μ that we described above may be imperfect in some sense.

Sensitivity analysis: Suppose we change the diagonal elements from (0,0) to $(\varepsilon,\varepsilon)$, for some number ε.
If $\varepsilon<0$, the game is elementary, as the above IC μ then satisfies strictly all nontrivial incentive constraints.
If $\varepsilon>1$, it is elementary, with $\mu'(a_2|r_1)=\mu'(b_2|s_1)=\mu'(c_2|t_1)=1$ satisfying strictly all nontrivial incentive constraints.
If $0<\varepsilon<1$, no incentive constraint can be satisfied strictly, and a dual solution where all nontrivial incentive constraints have strictly positive dual variables is:
$\alpha_1(s_1|r_1)=\alpha_1(t_1|s_1)=\alpha_1(r_1|t_1)=\alpha_2(c_2|a_2)=\alpha_2(a_2|b_2)=\alpha_2(b_2|c_2)=\varepsilon$,
$\alpha_1(t_1|r_1)=\alpha_1(r_1|s_1)=\alpha_1(s_1|t_1)=\alpha_2(b_2|a_2)=\alpha_1(c_2|b_2)=\alpha_2(a_2|c_2)=1-\varepsilon$.
These games have nontransitive dual solutions only in two cases: when $\varepsilon=0$, and when $\varepsilon=1$.

A 3×4 example with one sender and one receiver:

p:	T_1: \ C_2:	a_2	b_2	c_2	d_2	
1/3	r_1	3, 0	0, 3	0, 3	3, 0	[*bad type*]
1/3	s_1	9, 9	8, 8	0, 0	0, 0	[*left good type*]
1/3	t_1	0, 0	0, 0	8, 8	9, 9	[*right good type*]

Dual solutions include $\alpha_1(s_1|r_1) = \eta$, $\alpha_1(t_1|r_1) = 1-\eta$ for $1/3 \leq \eta \leq 2/3$,
$\alpha_2(b_2|a_2) = 1$, $\alpha_2(c_2|d_2) = 1$, with all other nontrivial components of α being 0.

For the symmetric solution η=1/2, the reduced game looks like:

q:	1's reduced-type:	$\{b_2\}$	$\{c_2\}$
0.5	$\{s_1\}\sim(2/3)[s_1]+(1/3)[r_1]$	5.33, 6.33	0, 1
0.5	$\{t_1\}\sim(2/3)[t_1]+(1/3)[r_1]$	0, 1	5.33, 6.33

With η=1/3, an asymmetric reduced game on one end would be

q:	1's reduced-type:	$\{b_2\}$	$\{c_2\}$
4/9	$\{s_1\}\sim 0.75[s_1]+0.25[r_1]$	6, 6.75	0, 0.75
5/9	$\{t_1\}\sim 0.6[t_1]+0.4[r_1]$	0, 1.2	4.8, 6

With η=2/3, an asymmetric reduced game on the other end would be

q:	1's reduced-type:	$\{b_2\}$	$\{c_2\}$
5/9	$\{s_1\}\sim 0.6[s_1]+0.4[r_1]$	4.8, 6	0, 1.2
4/9	$\{t_1\}\sim 0.75[t_1]+0.25[r_1]$	0, 0.75	6, 6.75

All these reduced games are elementary, with strict mechanism $(\{s_1\}\rightarrow\{b_2\},\{t_1\}\rightarrow\{c_2\})$.
The *belief probability* of r_1 given $\{\tau_1\}$ is $(1/3)\alpha_1(\tau_1|r_1)/((1/3)\alpha_1(\tau_1|r_1)+(1/3)(1))$.
Belief probabilities of the bad type r_1 must be ≥ 0.25 to deter a_2 & d_2.